

Product Innovation, Product Diversification, and Firm Growth: Evidence from Japan's Early Industrialization

Serguey Braguinsky, Atsushi Ohyama, Tetsuji Okazaki, Chad Syverson*

June 2020

Abstract

We explore how firms grow by adding products. In contrast to most earlier work on the topic, our conceptual and empirical framework allows for separate treatment of product innovation (vertical differentiation) and diversification (horizontal differentiation). The market context is Japan's cotton spinning industry at the turn of the last century. We find that introducing innovative products outside of the previously feasible is a key to firm growth. It provides opportunities to capture high-end vertically differentiated product markets when successful while also facilitating the firm's growth through horizontal expansion in product space. However, this process involves a high degree of uncertainty, so firms tend to introduce innovative products on experimental basis. In long-term outcomes, the right tail of the firm size distribution becomes dominated by firms that were able to expand in both directions: moving first into technologically challenging vertically differentiated products, and then later applying their newly acquired high-end technical competence to horizontal expansion of their product portfolios.

* Braguinsky: University of Maryland, Robert H. Smith School of Business, Osaka University Institute of Economic and Social Research, and NBER, e-mail: sbraguinsky@rhsmith.umd.edu; Ohyama: Institute of Innovation Research, Hitotsubashi University, e-mail: ohyama@iir.hit-u.ac.jp; Okazaki: Tokyo University, Faculty of Economics, e-mail: okazaki@e.u-tokyo.ac.jp; Syverson: University of Chicago Booth School of Business and NBER, e-mail: chad.syverson@chicagobooth.edu. Braguinsky and Syverson were supported in part by NSF grant #1632833. Braguinsky was also supported in part by International Joint Research Promotion Program (Osaka University). Ohyama was supported in part by JSPS KAKENHI Grant #15K03409 and #15KK0078. Okazaki was supported in part by JSPS KAKENHI Grant #190H00587.

1. Introduction

We now know an important mechanism for firms' growth is expansion of the scope of product varieties they offer. The recent research literature has used multiple frameworks to analyze this process theoretically and empirically. One set emphasizes supply-side innovation, where development of new, innovative products delivers a whole (sub)market to successful innovators (e.g., Klette and Kortum, 2004; Klepper and Thompson, 2006), or productivity determines firms' range of products (Bernard, Redding, and Schott, 2010). Another set focuses on demand-based heterogeneity where firms expand the number of varieties in differentiated product markets (e.g., Bernard, Redding, and Schott, 2010; Khandelwal, 2010; Hottman, Redding, and Weinstein, 2016). In both cases, firm heterogeneity is captured by endowed primitives (productivity, demand appeal, or both) that influence firms' size, scope, and growth rates. Furthermore, this work has typically abstracted from the specific identities of and interrelations among a given firm's product offerings. Instead, product variety scope is typically summarized by the number of symmetrically related products the firm makes.

However, product additions are unlikely to all be the same in reality. Adding a highly innovative product may affect growth differently than adding one similar to those the firm has already been producing. First, demand and cost effects may be different, as customers' willingness to substitute and firms' economies of scope are likely to depend on the degree of product similarity. Furthermore, and perhaps more importantly in the long run, bolder innovation may result in firms acquiring technical and/or marketing knowledge applicable to subsequent development of other product varieties. Thus it is not just the number of product varieties a firm makes that summarizes its heterogeneity but also the nature and interrelation of its differentiated products, such as the vertical distance between new and existing varieties.¹

The process of product variety expansion and firm growth through this channel is not well understood. In this paper, we employ rich historical firm-level data from the Japanese cotton spinning industry at the turn of the last century to take a step toward opening the "black box" of what and how firms do to expand their product varieties and grow. The data, explained in more

¹ "If a firm previously producing air brakes of various kinds...enters the production of electronic equipment, it is certainly diversifying its productive activities...although it may reduce the varieties of air brakes produced. ... Clearly...for a study of the growth of firms the type of diversification and the reasons for it are of more relevance than the 'amount' of diversification." (Penrose, 1959, p. 96)

detail below and in the appendix, are uniquely suited to examine growth through product variety expansion and its effects on long-term firm outcomes. One key feature of the data is that we can distinguish between introductions of products that lie outside the firm's current technological frontier from those inside the frontier, and we can directly observe the entrepreneurial actions aimed at such introductions.

We show how innovation along the dimension of vertical differentiation involves facing a "supply-side constraint" and a high degree of uncertainty, because it involves installing a new type of capital (machines) and using previously unfamiliar inputs and production processes. Vertically innovating firms coped with these hurdles through entrepreneurial experimentation (Kerr et al., 2014; Cusolito and Maloney, 2018, Ch. 4) and by recruiting top-notch engineering human capital. We further show how such vertical innovation (resulting from experimentation in which we can obtain plausibly exogenous variation) led to later horizontal product diversification within the firm's technology frontier.

This first-vertical-then-horizontal sequencing shapes firms' paths of product variety expansion and growth. We characterize the specific mechanisms through which firms applied the knowledge they gained through vertical innovation in order to spur their horizontal expansions and overall growth. Given the relationship and sequencing between vertical and horizontal product expansion, it is perhaps not surprising that we also find that firms that did not vertically innovate did not do much horizontal product diversification either. Those firms' much more modest growth was limited to intensive-margin expansion in sales of their existing products.

Some aspects of our findings align with past studies, but others provide novel insights into the relationship between product variety expansion and firm growth. First, as predicted by the oft-used Klette and Kortum (2004) model, more new product introductions are associated with higher average firm growth. Beyond this, however, our results highlight the importance of accounting for the identity of a newly developed product and its impact on subsequent product variety expansion. Vertically differentiated innovations are special. They have spillovers into horizontal expansions that do not operate in the reverse. We also show that vertical innovation tends to happen not incrementally, where firms move from more to less familiar products (for example as envisaged by Stokey, 1988), but rather through leapfrogging that is followed by "bridging" the firm's existing and newly innovated products (in the style of Callander, 2011).

Second, similar to the analyses in which product/firm appeal plays a critical role in determining a range of product varieties (Bernard, Redding, and Schott, 2010; Hottman, Redding, and Weinstein, 2016), we find that high-growth firms have more product flexibility and improve the quality of their products to better respond to demand changes. While those studies typically treat product/firm appeal as a primitive, our analysis suggests that appeal arises endogenously through experimenting with innovative products. Understanding the relationship between product variety expansion and firm growth requires incorporating the dynamic aspects of product variety expansion.

Third, our findings render support for the notion of complementarity between flexibility of the production system and the number of product varieties (Roberts, 2004). Firms that did not experiment with technologically challenging innovative products were confined to a narrow product variety range and infrequent product changeovers. Experimenting firms, on the other hand, changed their product portfolios frequently and expanded their overall scope. Our findings indicate product upgrading experimentation generated such complementarity.

Finally, in our data, only a limited number of firms attempt to push out the technology frontier and extend the set of the industry's feasible offerings. These firms exhibit "awareness" in the sense of Karni and Vierø (2013) that sets them apart. In the appendix, we use our data to describe the selection process through which firms that push out the technology frontier through product innovation discover and harness new growth opportunities.

We organize our analysis as follows. In the next section we briefly describe our data and historical context (with more details provided in the appendix). Section 3 documents basic trends in industry and firm growth through product varieties expansion. These include the tight relationship between growth in average industry firm size and product variety; that much of this variety growth come from existing firms rather than entry, and that—crucially—product variety growth in all directions was tied to firms' experimental attempts to climb the quality ladder, but not their efforts at horizontal product expansions. Sections 4 and 5 examine in more detail the connections among product upgrade experimentation (that is, trying to climb the quality ladder), its antecedents, and firm growth. Section 6 probes the specific mechanisms through which product upgrade experimentation led to growth, finding that its effects on firms' production flexibility and market-perceived quality as being particularly important. In Section 7 we conduct several robustness checks and tests of alternative hypotheses. Section 8 concludes.

2. Data

2.1 Data and historical context

Our main data come from monthly bulletins (Geppo, 1893-1914) published by Japan's Cotton Spinners Association (hereafter, "Boren" for short, using its Japanese acronym). These report, for every Japanese cotton spinning firm, the quantities of each product the firm made that month. (See Photo 1 in appendix A.1 for a photocopy of an original of one such report.) The data start in May 1893 and extend through December 1914. The industry exhibited phenomenal growth during this period (see, e.g., Braguinsky et al., 2015; Figure A1 in appendix A.2 depicts the dynamics of industry-wide output and the number of firms in our sample) and caught up with the worldwide technological frontier by expanding into high-end products.

We match this product variety data with a firm-level database that contains monthly measures of inputs (spindles in operation, factory operatives, raw cotton) in physical units, firm-specific output prices (for select counts), wage rates, and the existence and size of industry-wide output cuts imposed by the association in periods of slow demand. We aggregate monthly-frequency data to semi-annual level to correspond to the frequency of observations on important supply constraint-related characteristics, such as each firm's machine capacity (number of installed spindles), employment of educated engineers, and board composition available at semi-annual frequency from shareholders reports and other sources. More details are provided in appendix A.1.

Spun cotton yarn is differentiated by thickness, measured by "count," which gives the length of a type of thread in yards that would weigh one pound. Thus higher counts correspond to finer yarn. Higher count yarn is more comfortable to the touch and as such is of higher quality. Yarn is also differentiated by the direction in which it is twisted during the spinning process (S-twist vs. Z-twist). Some yarns are produced by twisting two single-yarn threads together using separate equipment called doubling frames. These yarns were translated into Japanese at the time as "doubled yarn."² Both single and doubled yarn can be processed further through a process called gassing, which involves passing yarn quickly through gas burners to burn away fluff and

² In the English term it is "twisted yarn" (see, e.g., Woodhouse, 1921). We will use the Japanese "doubled yarn" terminology, in part because it is important for our purposes to distinguish also single yarns by the direction of twist as above. There is no distinction between Z- and S-twist in case of doubling because the second twist is always applied in the direction opposite to the first twist.

make the product glossy. The result is called “gassed yarn.” The demand for various types of cotton yarn comes from weavers, and the degree of substitution across yarn varieties in producing a particular garment or textile is generally quite low. Our product-by-firm data reports count, twist (if single yarn), and doubling or gassing (if done).

We use the degree of technical difficulty in producing a product as our conceptual and empirical measure of vertical product differentiation. This implies finer (thinner) and more processed threads are of higher quality. They generally require more versatile and/or specialized capital stock (machines designed for thinner counts, doubling and gassing frames, etc.), higher-quality raw cotton as input, as well as superior technology and production organization. This classification is consistent with overall technological trajectories in the Japanese cotton spinning industry. As Tables A3-A4 in appendix A.2 show, at the beginning of our sample (10 years after the inception of the industry) lower-count yarns—counts 20 and below—were almost entirely domestically produced. Users of higher-count yarns still largely relied on imports from Britain because of Japanese cotton spinning firms’ lack of technical proficiency. When expanding their product varieties, Japanese firms thus faced a problem whether to add a lower-count cotton yarn that was easy to produce but faced limited demand, or a higher-count cotton yarn that was difficult to produce but potentially faced a much less crowded market.

The thickest count recorded in the data during our sample is 2.5 (S-twist), and the finest is 100 (gassed). The latter doesn’t first appear until made by a single firm in January 1903. Product scope evolved as the industry did. At the beginning of the sample, industry firms made only about 30 products. This grew to over 100 by the end.³ The total number of different products that show up at least once in the data from 1893-1914 is 201. To create a set that can be consistently applied throughout the sample period, we aggregated these into 35 different product varieties—10 varieties (different counts or count ranges) each of S-twist and Z-twist single yarns, 10 varieties of doubled yarn, and five varieties of gassed yarn. Our references to a “product variety” below regard one of these 35 varieties unless otherwise stated.

For the purposes of this study, we often dichotomize varieties into the low end and high end of the product space, with counts 20 and below (and not doubled or gassed) low-end

³ In the process, some previously disaggregated data on very thick counts start being reported in a more aggregated way (e.g., counts 10 and below are lumped together after a certain point in time), while the data on finer counts remain reported in a more disaggregated fashion.

varieties. Counts above 20 are high-end products. Among our 35 total varieties, 16 are high-end according to this classification and 19 low-end. We chose this threshold because counts above the mid-20s were generally impossible to produce without specialized machines and inputs designed for those varieties.⁴ Thus, while demand constraints may not have been that important for high-end products (competition was mostly from imports, and Japanese textile makers would have been willing to substitute to domestic sources at similar FOB prices), supply-side (machine and technology) constraints were very real. The demand limitations faced by low-end products, on the other hand, are highlighted by the differential treatment they received during mandatory output cuts periods. In periods of slow demand, Boren (the industry association) imposed mandatory output cuts on its member firms. However, output of counts above 20 were largely exempt from those cuts (see appendix A.3 for details). This differential treatment also allows us to utilize mandatory output cuts as an instrumental variable in our firm growth analysis, as we explain below.

To investigate firms' product expansion patterns, we distinguish between new-to-firm vertical product differentiations aimed at pushing out the firm's technology frontier (we also refer to these as "product upgrades" below) and horizontal differentiation aimed at diversifying the product portfolio inside the existing technological frontier ("product diversifications"). While closely related to the distinction between high-end and low-end product varieties above, the definition of an "upgrade" is stricter than just making a high-end product. A new product addition is an upgrade if it a) involves a high-end product *and* b) the firm must not have previously made at scale a product of an even higher count. A new product addition is called a "diversification" if, regardless of the new product's count, the firm had previously made a product of a higher count.

2.2 Experimental production and product upgrading

⁴ As is nearly inevitable with classifications, there is some "gray zone." In this particular case, some classifications adopted in Japan at the time classified single mechanically spun yarn into coarse yarn (counts below 20 as well as 20-22 count yarn), medium yarn (counts 23-44), and fine yarn (counts 45 and above). Our definition considers both medium and fine yarns as high-end, and we also include counts 21 and 22 as such. Reclassifying these two counts as low-end products makes no difference to our results. This is perhaps not surprising given the relatively small amount of production of counts 21 and 22. While for example 20-count output alone represented roughly 30 percent of the industry's total output during our sample, counts 21 and 22 together accounted on average for only about 1.1 percent of industry output.

The notion of experimental product introduction plays an important role in our analysis below. We explain how we define it here. It is the addition to the firm's portfolio in semi-annual period t a variety that is both novel and produced at a modest quantity. More precisely, for a product variety j to be defined as (the beginning of) an experiment at t , it has to satisfy all of the following conditions: (i) the firm did not produce variety j in $t-1$, (ii) variety j accounts for less than three percent of the firm's total output in t , and (iii) variety j had never accounted for more than three percent of the firm's total output in any period before t .⁵ Thus, a product variety remains in the "experimental stage" from introduction (if at a scale below the three-percent threshold) until it reaches this threshold for the first time (if at all). Once a firm produces a variety at a scale above the threshold, it is no longer considered experimental even if output later falls below the threshold again. This definition allows for potentially multiple experiments in the same product line. If a firm temporarily discontinues an experimental product before it reaches the threshold and re-introduces it below the threshold after at least one period of non-production, this is a second experiment in that product line. (A further cycle would be a third experiment, and so on.) We define an experiment as "successful" if the product rises above the three percent threshold after production for a continuous set of periods; otherwise, the experiment has "failed."

Panel A of Table 1 shows that among the 685 firm-product varieties introduced during our sample, 439 (64 percent) were initially "experimental," and of those, 271 never reached the threshold scale.⁶ Hence, about 60 percent of product lines introduced on experimental basis completely "failed," that is, never scaled above the threshold despite possibly multiple experiments. Of these experimental varieties, 76 (11 percent) involved upgrade products. (Recall that an experiment is an upgrade if it the firm produces a high-end product with a count higher than any count the firm had produced above the three percent threshold before.) Half of these never reached the threshold. In Panel B we look at all experimental product launches, counting multiple experiments in the same product line as separate episodes. The fraction of experiments that failed is 73 percent for all products, and 64 percent for upgrade experiments. We can thus see that *experimental product development often fails*.

--- Table 1 around here ---

⁵ We tried other reasonable thresholds, such as four percent or two percent, and the results were qualitatively similar. See Figure A2 and its discussion in appendix A.4 for the distribution of initial scales of newly introduced products.

⁶ Note that if a firm operated only two or three product lines, those lines are less likely be counted as "experimental" by construction. Hence, the above probably represents a lower bound on the actual number of experiments.

The small scale that defines an experimental product and the high failure rates indicate that product experimentation *per se* cannot have a large direct effect on firm growth. However, experimentation might offer firms valuable technical or marketing knowledge, especially if the experiment involves upgrade products that require mastering new technologies. In our analyses below we use the cumulative number of upgrade experiments conducted by firm i through time t as a proxy for technical knowledge accumulated through experimentation. Similarly, we also compute a firm's cumulative number of (horizontal) product diversification experiments to use as an additional regressor in some specifications.

2.3 Machine orders and high-end versus low-end machines

Japan did not produce its own cotton spinning machines during our sample. It imported all machinery from Britain, mostly from Platt Brothers of Oldham (Saxonhouse, 1974). The Platt collection in Preston, Lancashire, U.K. contains British textile machine manufacturers' order books (including but not limited to Platt Brothers) reflecting orders from cotton spinning firms worldwide. We collected data on all Japanese orders placed from the inception of the industry in the early 1880s through 1914.⁷

All machines were custom made. Each order contains the placement date, shipping dates (usually multiple, as machines were commonly shipped in several installments), type of frames ordered (ring, mule, or doubling), number of frames and spindles per frame, the range of counts the frames were designed to spin, and other technical characteristics (description of cotton input, hank roving to be fed into the machines, rotation speed of the spindles, etc.).

We matched these orders with Japanese archival sources (*Enkakukiji*, 1901; *Sankosho*, 1903-1914; and individual company reports) that provide semiannual firm-level data on the total installed capacity: number of spindles, separately by ring, mule, and doubling frames. The details of the matching process are described in appendix A.1.6. While we could not match all changes in firms' capacity as recorded by Japanese sources with corresponding orders from British textile machine manufacturers, we were able to do so for 105 of the 118 firms in our sample. For these firms, we constructed an alternative measure of firm capacity by summing up the number of

⁷ These orders had been previously examined and processed by Gary Saxonhouse and archived at the ICSPR after he passed away (Wright, 2011). However, with no originals, it turned out to be impossible to match ICSPR archived data to the Japanese data at the firm level, as most firm names were either missing or incorrectly assigned. We will make our newly collected data, including digital photos of the original orders, publicly available.

spindles in all orders placed by a given firm, while also taking into account transfers through acquisitions, second-hand purchases and sales, decommissioning, and destruction in earthquakes and fires. In the end, we accomplished an average match rate of 99 percent between this “reconstructed” number of spindles and the number actually recorded in the corresponding firms’ balance sheets. Most of the remaining 13 firms for which we do not have orders data were short-lived and/or very small. They constitute just two percent of total industry capacity at any given point in time. The bottom line is that we have machine capacity breakdowns for almost all of our sample. We can therefore link machines’ technical characteristics to the product varieties produced by firms that owned those machines.

As mentioned, high-end products and low-end products required different machinery. We employ the matched orders data to distinguish between machines designed predominantly for high-end products and those for low-end products. More precisely, we define machines as high-end if their specifications indicated the ability to spin counts of 23 or higher. We classify any other machines as low-end.⁸ We also classify most doubling frames (for which our main source of data is Japanese firms’ balance sheets) as high-end machines, with the exception of a few cases where firms had no high-end frames and thus applied doubling to low-end products. Among the 105 firms for which we have a breakdown of machine capacity, 42 had at least one high-end machine at some point in our sample. However, the capacity distribution was highly skewed, with a median number of spindles of 16,128 and a mean of 36,763.

2.4 Mandatory output cuts as a source of exogenous variation

As mentioned, the industry association Boren periodically imposed output cuts for low-end yarns, applied uniformly to all firms, during periods of slow demand. The timing of these

⁸ See Photos 2 and 3 and the corresponding description in appendix A.1 for two specific examples. There is ambiguity about some machines ordered early (before the start of our product-variety data), when Japanese firms could not yet produce counts above 20. For example, an order placed by Osaka Spinning Company in October 1888 lists the counts to be spun from 10-20 using Japanese cotton, but “also up to No. 32 if the Japanese decide to mix imported cotton with Japanese cotton.” (Osaka Spinning was working to develop Indian imported cotton at the time, but actual imports had not started yet.) We classify these machines as high-end due to their versatility, even though they were different from subsequent high-end machines specifically designed for high-end products and for use with even higher-quality U.S. and Egyptian cotton. That said, re-classifying those early high-end machines as low-end instead does not affect our main findings.

cuts was not under the control of any given firm.⁹ Because high-end products were exempt from these mandated cuts, they gave firms extra incentive to experiment with high-end products.

In appendix A.3 we present the details of output cuts imposed at different points in time, compiled from Shoji (1930). There were two major types of curtailments. In one case, Boren would impose a certain number of days in a month as “mandated holidays.” Machines producing low-end (count 20 and below) yarn were not allowed to operate on those days. The other type of curtailments required a certain fraction of spindling frames used to spin low-end yarn to be idled. Compliance was enforced by inspections and by Boren officials securing idled frames by putting physical seals on spinning rails (Shoji, 1930, p. 156).¹⁰ The fraction of such sealed equipment ranged from 20 percent to 40 percent (Table A6 in the appendix).

To differentiate these mandatory output curtailments by their severity, we constructed a variable which takes a value of zero during months with no curtailments and equals the fraction of idled low-end machine spindles (or its equivalent calculated from the number of mandated holidays) during months with mandatory output cuts. The average of this value over a given semi-annual observation period, shown in the third column in appendix Table A6, represents our measure of the degree of relative output cuts imposed on low-end products.¹¹ In subsequent analyses, we interact this measure with the installation of high-end machines ordered from England to obtain an exogenous variation in incentives to conduct upgrade experiments.

3. Product Variety Expansion Patterns and Firm Growth

3.1 Product variety expansion

3.1.1 Decomposition analysis

Figure 1 plots the dynamics of average firm size as well as the average number of high-end and low-end product varieties per firm.¹² Average varieties per firm is mostly flat from the

⁹ Association voting rules required consensus from all member firms for output cuts to be imposed (Shoji, 1930), so it was impossible for any firm (or a coalition thereof) to exploit the policy for unilateral advantage.

¹⁰ Although Shoji (1930) does report some instances of temporary non-compliance by a handful of firms, generally speaking, the enforcement appears to have been quite effective.

¹¹ As can be seen from Table A6, there were some mandated cuts among high counts in 1910-12. Because we are interested in the *relative* stringency of low-end versus high-end output cuts, when constructing our measure during that period we subtracted the fraction of high-end cuts from the fraction of low-end output cuts.

¹² We employ conversion coefficients similar to those developed in Braguinsky et al. (2015) to aggregate threads of various counts into a single, “20-count equivalent” measure of total firm output in physical units. All the results below are robust to using machine inputs (number of spindles in operation) or labor inputs (number of worker-hours) as alternative measures of firm size.

start of our data until about 1899, but then starts increasing, first in high-end product space and later in low-end products. There is a particularly sharp increase in both product types around 1907. From 1907 onward, there is a high correlation between the total number of product varieties per firm and average output per firm. A salient feature of industry-level product variety expansion in Figure 1 is that low-end products are not replaced by high-end products. Rather, a sharp increase in the number of high-end product varieties accompanies an increase in the number of low-end product varieties. We examine this pattern closely below.

--- Figure 1 around here ---

We now employ more formal decomposition analysis to quantify sources of change in the number of product varieties per firm we observe in Figure 1. We first decompose the market-share-weighted average number of products at time t as

$$\bar{y}_t \equiv \sum_{i=1}^{N_t} s_{it} y_{it} = \sum_{i=1}^{N_t} s_{it} \bar{y}_i + \sum_{i=1}^{N_t} s_{it} \tilde{y}_{it},$$

where N_t is the number of firms operating at time t ; y_{it} and s_{it} are firm i 's number of products and market share at time t , respectively; $\bar{y}_i = \frac{1}{T_i} \sum_{t=1}^{T_i} y_{it}$ is the average number of products produced by firm i over the whole period it is observed in the sample (T_i), and $\tilde{y}_{it} = y_{it} - \bar{y}_i$. The change in the weighted average number of products between t and $t+1$ is then

$$\begin{aligned} \bar{y}_{t+1} - \bar{y}_t = & \left[\sum_{i=1}^{N_{t+1}} s_{it+1} \bar{y}_i - \sum_{i=1}^{N_t} s_{it} \bar{y}_i \right] + \left[\sum_{i \in C} s_{it} (\tilde{y}_{it+1} - \tilde{y}_{it}) \right] + \left[\sum_{i \in C} \tilde{y}_{it+1} (s_{it+1} - s_{it}) \right] \\ & + \left[\sum_{i \in EN} s_{it+1} \tilde{y}_{it+1} \right] - \left[\sum_{i \in EX} s_{it} \tilde{y}_{it} \right] \end{aligned}$$

where C , EN , and EX indicate continuing firms, entrants, and exiting firms, respectively.

We call the first term on the right hand side the “composition effect.” It captures the change in the average number of products due to the difference in the composition of firms between t and $t+1$. This measures the difference, over their lifetimes, in the average number of products of new entrants versus exiting firms. The second term is the “expansion effect” of continuing firms. It captures within-firm changes in product variety counts between t and $t+1$, holding fixed firms’ base period market shares. The third term is the “allocation effect,” measuring the contribution of changes in the market shares of continuing firms between t and $t+1$. Finally, the fourth and fifth terms measure the contribution of entrants and exiting firms, respectively, coming from deviations in the first observation (for entrants) and the last

observation (for exiting firms) from their own long-term average number of products. A positive (negative) number in the fourth term means that entrants' average number of products varieties is greater (less) when they enter than in later periods of their operation. A positive (negative) number in the fifth term means that exiting firms' average number of products is greater (less) right before they exit than in their earlier periods of operation. We label the sum of the last four terms in the decomposition the “overall within effect.”

Table 2 presents the results. The main takeaways are as follows. First, the decomposition shows there was an increase of 9.0 in the average number of all products per firm between 1893 and 1914. Of this, the within effect accounts for a greater share (5.6 products) than the composition effect (3.3 products). However, when we divide the sample before and after the spike in product varieties in 1907, there are stark differences. From 1893-1906, the overall increase of 2.5 products per firm is more modest, and the composition effect (1.6) contributes more than the within effect (0.9). Decomposition of the within effect also shows that it is entirely driven by allocation—changes in market shares. We see a complete reversal from 1907-1914. While the absolute magnitude of the composition effect remains roughly the same (1.7), the within effect (4.7) becomes a dominant contributor to the total growth in average products per firm of 6.4. Almost all of this within effect now comes from the expansion of continuing firms, while the allocation effect is much smaller.

Thus, during the first subperiod (1893-1906), which is the period of large-scale entry, followed by a shakeout and initial industry consolidation, the growth in number of product varieties was driven by new entry and by increasing market shares of continuing firms that produced more product varieties. In the second subperiod (1907-1914), while new entry still contributed to product variety growth at about the same magnitude as before, the number of product varieties produced by continuing firms swelled. This growth came to dominate the overall expansion of the number of product varieties produced by the industry. Separate decompositions of high- and low-end products present essentially the same picture in both cases, though with some nuanced differences between them.

--- Table 2 around here ---

Thus, decomposition analysis shows that the big boost received by industry growth after 1907 from the expansion of the number of product varieties seen in Figure 1 resulted almost

entirely from an increase in the number of both high-end and low-end products produced within continuing firms.

The findings from the decomposition analysis raise questions. Were the firms whose expansion drove the growth in high-end and low-end product varieties the same or different firms? If they were the same, did they first expand their low-end product varieties and then move on to high-end products, as implied by the learning-by-doing theory of Stokey (1988), for example? Or, did they invert that order?

3.1.2 Same or different firms?

Table A7 in appendix A.5 shows that as high-end machine capacity increases from virtually zero at the start of our sample, it spreads across firms of different sizes but remains heavily concentrated among the largest firms. Perhaps not surprisingly, this leads to the same pattern in the number of high-end product varieties. More interestingly, however, the number of low-end product varieties, which do not require high-end machines for their production, also becomes heavily concentrated in the same set of firms. Thus the same firms accounted for the expansion of both high- and low-end product varieties. Moreover, those were also the firms that invested in high-end machines and grew to become the top firms in the industry (see also appendix Table A7). This leads to questions about the relationships between product upgrading and diversification and their ties to firm growth. This is what we explore next.

3.2 Upgrade experiments and product varieties expansion

We saw above that almost two-thirds of firms' new products were initially launched on a small scale (experimental production), and 60-70 percent of those experiments were not scaled up. Most experiments could therefore not directly contribute to the proliferation of new product varieties. It turns out, nevertheless, that upgrade (although not diversification) experiments did contribute to product varieties expansion. We establish this fact here, and we examine the mechanisms behind it in the next section. More specifically, we show how accumulated past upgrade experiments are tied to expansion in product varieties of all types, even those not directly targeted by such experiments.

The estimation equation is:

$$y_{it} = \alpha + \beta_1 cuml_upgrade_exp_{it-1} + \beta_2 cuml_divers_exp_{it-1} + \beta_3 X_{it-1} +$$

$$\gamma_i + \delta_t + \zeta_\tau + \varepsilon_{it}, \quad (1)$$

where y_{it} it is the change from $t-1$ to t in the total number of products produced by firm i , and separately in the number of high-end and low-end products; $cuml_upgrade_exp_{it-1}$ and $cuml_divers_exp_{it-1}$, represent the cumulative numbers of upgrade and diversification experiments conducted by firm i by time $t-1$, respectively; X_{it-1} is a vector of controls, and ε_{it} is the error term. We include firm fixed effects γ_i and semi-annual period fixed effects δ_t . In addition, we include a set of dummies ζ_τ to nonparametrically control the number of periods τ the firm has been in the sample (the first equal to one for every firm's earliest period in the sample, the second equal to one for firms' second periods, and so on). We include these to ensure that our key variables of interest, the cumulative numbers of upgrade and diversification experiments, capture accumulated past experience with experiments and not simply how long the firm has happened to be in the data. We exclude each firm's first and last periods because they often cover less than a full six months.¹³

The results, shown in Table 3, make clear that a firm's past upgrade product experiments, but not diversification experiments, are tied to growth in its number of products of all types (total, as well as high- and low-end specifically). Columns (1)-(3) include the number of products (all, high-end and low-end, respectively) produced by firm i at time $t-1$ as a control in X_{it-1} .¹⁴ An additional past upgrade experiment is associated with adding product varieties of all types in the following period. To give some sense of the magnitude of this relationship, conditional on past upgrade experiments being positive, the 25th-percentile of cumulative past upgrade experiments is one, while the 75th-percentile is five experiments. Hence, the interquartile differential is tied to about 1.7 (0.43x4) more new product varieties added during any given semi-annual period.

It is worth emphasizing that past upgrade experiments are associated with not just future growth of high-end products, but low-end products as well. Indeed, the magnitudes of the

¹³ We do the same in all subsequent regressions below, unless explicitly stated otherwise.

¹⁴ The total number of available product varieties (and hence the potential number of products any firm can produce) is bounded from above in our data, so we need to control for the level of product diversification already attained. To check if the estimation results are sensitive to the potential "subtraction bias" (inclusion of the lagged number of products, which also enters the dependent variable with the minus sign, as an explanatory variable), we conducted an ordered logit estimation with the dependent variable being a dummy taking values of minus 1, zero, and plus one if the firm respectively reduced, did not change, or increased the number of product varieties from $t-1$ to t . The results were very similar; details are available upon request.

experiment-associated high- and low-end product growth are similar. Because upgrade experiments never involve low-end products by construction, the results in column (3) suggest there may be substantial spillovers from firms experimenting with product upgrading to their abilities to increase the number of seemingly unrelated low-end products.

--- Table 3 around here ---

Columns (4)-(6) of Table 3 include additional controls: an indicator for the firm installing new machines (whether high-end or low-end) during the observation period, and another indicator equaling one if the firm employed a university-educated engineer or had a board member who was a prominent cotton yarn or garments merchant (as a proxy for “connectedness” to markets; see Braguinsky et al., 2015). The point estimates on past upgrade experiments fall slightly but retain economic and statistical significance.

Among the covariates, high-end machine installation has a positive, economically large and statistically significant association with both high-end and low-end product varieties expansion. Other things equal, expanding high-end machine capacity relates to an average increase of new product introductions of 0.66 products per period, about 40 percent ($0.27/0.66$) of which are new high-end products and the remainder new low-end products. The tie between growth in low-end products and high-end machine expansion (which in principle could be used for low-end production but rarely were because of their expense) is more suggestive evidence that pushing the technology frontier helps firms grow within the frontier as well.

Low-end machine expansions, on the other hand, do not exhibit patterns characteristic of spillovers. They are statistically unrelated to the expansion of high-end product varieties. And while they do accompany growth in low-end products, the magnitude of this relationship is about half of the relation with high-end machine expansion.

It is also worth noting that the relationships between high-end machine expansion and product growth hold even conditioning on past upgrade experiments. As we will see below, however, adoption and expansion of high-end machines (but not of low-end machines) are nevertheless significantly related to additional experimentation.

4. Exploration of Mechanisms

The relationship between pushing the technology boundary and subsequent product variety expansion and firm growth raises questions about the mechanics behind it. As mentioned,

most upgrading experiments fail to scale up and to stay in a firm's portfolio, so the direct effect of upgrade experiments on the expansion of product varieties is limited. Further, upgrade experiments have no direct effect whatsoever on the expansion of low-end product varieties yet are strongly connected to them quantitatively. We address these questions in this section.

4.1 High-end machines, market ties, engineering human capital, and product experiments

High-end machines were a prerequisite for venturing into the high-end product space, and we saw above that having such machines was indeed associated with more product upgrade experiments. But delivering high-end products was difficult and entailed much uncertainty. At the beginning of our sample, Japanese mills were catching up to the world technology frontier. Introducing high-end machines was a costly investment decision. Moreover, successfully introducing high-end products often required more than just procuring high-end machines. It called for a combination of engineering talent, market knowledge, and product experimentation.¹⁵

In Table 4 we present summary statistics showing the relationships between having high-end machines and a number of firm-level outcomes: employment of university-educated engineers; the presence of cotton yarn and garments merchants (hereafter, “merchants” for short) on firms' boards of executives, as a proxy for market knowledge or “market ties”; and product experimentation.

--- Table 4 around here ---

Firms that installed high-end machines had merchants as their board members in 71 percent of the cases compared to 57 percent for other firms (a difference that is statistically significant). These firms also employed 1.6 university-educated engineers on average, ten times the level of firms that did not. A similar order-of-magnitude difference exists for the number of second-tier educated engineers—graduates from technical colleges, corresponding to today's

¹⁵ For example, Nihon Spinning, the first firm to produce gassed yarn in Japan, ordered its first machines from Platt Brothers in April 1894. Henry Ainley, a British engineer who met with three of Nihon Spinning's founders at the time, expressed surprise that the Japanese were even contemplating producing gassed yarn at that stage (Geppo, 1893, No. 5, p. 89). Indeed, it took two full years for the firm to actually start production. In-between, Japanese engineers, together with two British advisers who came from England to help, had to resolve a host of technological issues. These ranged from finding ways to procure heat-resistant brick in Japan to dealing with drafts that caused burners to flutter and damage the thread, all the while fighting suffocating heat inside gassing chambers (Kinugawa, 1964, Vol. 7, p. 13). See also the example of Amagasaki Spinning company detailed in the appendix.

Institutes of Technology. Thus, adoption of high-end machines was closely related to superior market ties and engineering human capital.

The numbers in Table 4 also show that firms with high-end machines conducted experimental product introductions at a much higher frequency, introducing an average of 1.2 products per year (0.6 per semiannual observation), as opposed to 0.5 products per year for those without. In contrast, as can be seen from the bottom three rows of the table, there was no difference in the frequency with which firms introduced new products above the three percent threshold. Thus, high-end machines facilitated experimentation with new products before scaling up in ways low-end machines could not, or at the very least, did not.

Notably, while high-end machines were unsurprisingly associated with a greater frequency of product upgrade experiments, they were also tied to a doubled propensity to conduct horizontal product diversification experiments. (The opposite is true regarding the propensity to introduce new products on a non-experimental basis, per the bottom rows.) Because high-end machines were not directly relevant to low-end products, this is once again consistent with adoption of high-end machines having spillover effects on all types of new product experimentation. Conditional on conducting experiments, firms with high-end machines had a lower fraction of successful experiments (new experimental product introductions that were subsequently scaled). Those with high-end machines scaled 22 percent of their experimental products compared to 35 percent (of a much smaller number) of experimental products of firms without high-end machines. This difference is economically and statistically significant and may reflect both higher uncertainty faced by high-end machine users and the fact that they derived useful knowledge from experimentation regardless of the outcome.

Given the size of the investment high-end machinery required and the potential heterogeneity in firms' benefits and costs from doing so, it is likely that the differences observed in Table 4 between firms with and without high-end machines are not purely causal. We take installations as given immediately below but later employ an instrument based on industry-wide mandatory output cuts of low-end products as a source of exogenous variation in firms' incentives to use their machines to conduct upgrade experiments.¹⁶

¹⁶ We examine further in appendix A.6 how firms select into procuring high-end machines. A key finding is that among firms that did not yet have high-end machines, those that would later purchase them—"future adopters"—tended to already conduct more experiments and have more merchant board members than firms that never installed high-end machines. Future adopters were somewhat larger in terms of output, but they were neither more diversified

4.2 Predictors of vertical upgrade experiments

We now examine more closely the complementary factors to upgrade experimentation. In Panel A of Table 5 we present the results of regressions where the dependent variable is the number of upgrade experiments started by a firm in a given semi-annual period. The explanatory variables of interest are whether the firm had installed high-end machines, whether it expanded its high-end (low-end) machine capacity, the presence of merchants on firms' boards of executives, and whether the firm employed a degreed engineer. We also control for firm age and semi-annual time dummies. Because the dependent variable is a count variable consisting of zeroes and small integers, we employ Poisson regression estimates. (We obtained similar results using negative binomial regression, as well as OLS.)

In the estimation results in column (1), the coefficient on the indicator for having high-end machines indicates that high-end machines were associated with starting about 1.1 additional upgrade experiments per semi-annual period. This echoes the pattern in Table 4 above. In column (2) we look at high- and low-end machine capacity expansions. Adding high-end machines, but not low-end machines, was associated with an additional 1.5 new upgrade experiments per period. Thus, both the presence and expansion of high-end machines were associated with firms experimenting more with vertically upgraded products. Repeating the same exercise on the subsample of firms that already had high-end machines in column (5) yields a coefficient that implies about 1.1 new upgrade experiments upon high-end machines expansions, a number that is still statistically significant despite the much smaller number of observations.

In column (3) we look at engineering human capital and market ties. Firms that employed a university-educated engineer started a marginally significant 0.7 more upgrade experiments per period. Having a merchant on the board of executives was associated with an additional 1.5 upgrade experiments per semi-annual period. Thus our proxy for having market ties is similarly related to incremental experimentation as high-end machine capacity expansions, and is twice that of having a university-educated engineer. These results are not directly comparable to the firm-fixed-effects estimations in Table 3 above because this regression uses across-firm

nor more likely to expand low-end machine capacity before adopting high-end machines. Thus again we see that diversification and expansion of low-end machines happened *after* firms successfully expanded into high-end product space.

variation. However, it is worth noting the contrast between the strong correlation between having market ties and upgrade experimentation here and the lack of a significant relationship between market ties and product varieties expansion seen above.¹⁷ Merchants on executive boards accompanied experimentation but played little explanatory role in subsequent product varieties expansion and, as we will see below, overall firm growth. This suggests market ties were important in helping firms cope with initial market uncertainty when experimentally introducing unfamiliar products, but subsequent expansion was less dependent on such ties and instead relied more on the experience and technological knowledge embodied in engineers.¹⁸

--- Table 5 around here ---

The relationships in columns (1)-(3) and (5) may reflect the influence of third factors that jointly determine both the explanatory variables and experimentation. To gain exogenous variation in experimentation for our causal analysis below, we construct an instrument that uses the interaction between contemporaneous installment of new high-end machines and the industry's imposition of mandatory output cuts for low-end products, as described in Section 2.4. Here, we use indicator variables for each condition for illustrative purposes; in our firm growth specifications below, we use more continuous measures of each.

The logic of this interaction as a valid instrument for product upgrade experimentation is as follows. The relevance condition requires that the coincidence of a firm's high-end machinery installation and mandated output cuts in a given period be correlated with the number of upgrade experiments conducted by the firm that period. The connection between having and installing high-end machines and conducting experiments is intuitively obvious. As for mandated output cuts, because they exempted high-end products, firms had a lower opportunity cost of shifting labor and capital inputs from (now restricted) low-end products to high-end opportunities.

The exclusion condition requires that the interaction be uncorrelated with any factors that drive experimentation other than through the relevance mechanism above. The mandated output cuts were imposed at the industry level and were sourced in the wake of unexpected demand

¹⁷ Regression specifications in Table 3 include firm fixed effects because we are interested in examining within-firm impact of knowledge accumulation, as proxied by cumulative past upgrade experiments. Nevertheless, if we drop firm fixed effects from these specifications, the coefficient on the indicator for having a merchant as a board member is still about half the magnitude of the coefficient on the indicator for having a university-educated engineer, in sharp contrast to the estimation results in Table 5.

¹⁸ As discussed below and in appendix A.11, the presence of merchants on executive boards was also closely related to experiments that led to the selection of firms into purchasing high-end machines in the first place.

shocks like the Boxer rebellion of 1900 or the stock market crash of 1907. It is unlikely therefore that the cuts would correspond to any firm- or even industry-wide shifts in the desire to experiment with new products. Regarding high-end machinery installations, one might be concerned that they were driven in part by a firm's expected return to high-end product experimentation. However, by using machinery that started operating only in period t to construct the instrument, we are taking advantage of the considerable lags between ordering, delivery, and installation of machinery in the industry. The minimum order-to-installation lag is one year and could often be longer, for random reasons; see Saxonhouse (1974). Therefore it is implausible that firms would—somehow anticipating a fall in demand and a resulting mandated output cut at least one year in the future—place a machine order with the intention of it arriving and becoming operational coincident with the output cut. We believe it was nearly impossible to time the arrival of the new high-end machines to coincide with the imposition of output cuts, even if firms were to consider doing so.¹⁹

We include this interaction term in columns (4) and (6) of Table 5. As can be seen, it is positively and significantly related to firms starting more upgrade experiments, verifying the relevance condition as an instrument. The effect is especially large among firms that already have high-end machines. There, high-end machine capacity expansions that coincided with mandatory output cuts led to those firms starting 1.7 new experiments in such periods.

4.3 Factors in diversifying experiments and selection into high-end machine adoption

Panel B of Table 5 shows the results of running the same specifications as in Panel A, except with the dependent variable being the number of *diversification* experiments started in a given period.

With the notable exception of the interaction of new high-end machine installation and mandatory output cuts (which of course applied to low-end products), the factors that explain diversification experiments are similar to those for product upgrade experiments in Panel A. Firms with high-end machines start on average about one more diversification experiment in each period than do firms with no high-end machines. And once again high-end, but not low-end, machine capacity expansions are associated with more experimentation. This conforms to the

¹⁹ High-end machine expansion also happened through acquisitions. There too it took at least a year, often more, to consummate an acquisition.

results in Table 3, which showed low-end product expansion was related to the same factors as high-end product expansion.

Engineers play a more prominent role than merchants in explaining diversification experiments. This is consistent with market ties playing a bigger role in the initial push to expand the technology frontier, while engineering and technical knowledge matters more for following through with production at scale. As we will see in the next section, the interaction between accumulated past upgrade experimentation with the expansion of low-end product varieties is a key determinant of firm output scale growth.

5. Product upgrading, product diversification and firm growth

We have established that the relationship between pushing the technology boundary and subsequent product variety expansion is mediated by experiments in moving to products up the quality ladder, followed by application of the knowledge gained toward diversifying low-end product offerings. In this section, we examine if these factors also translated into higher total output growth. Specifically, we measure how the interaction of a firm's cumulative upgrade experiments and the fraction of low-end products in the firm's portfolio affects output growth. This interaction summarizes the experiment-learn-apply process.

The basic estimating equation is:

$$\ln(y_{it+1}) - \ln(y_{it}) = \alpha + \beta_1 \text{cuml_upgrade_exp}_{it} + \beta_2 \text{low_frac}_{it} + \beta_3 \text{cuml_upgrade_exp}_{it} \times \text{low_frac}_{it} + \beta_4 X_{it} + \gamma_i + \Delta_t + \varepsilon_{it}, \quad (2)$$

where y_{it} is firm i 's output at time t , $\text{cuml_upgrade_exp}_{it}$ is the cumulative number of upgrade experiments conducted by the firm up to time t , low_frac_{it} is the fraction of low-end products in the firm's total number of products, and X_{it} is a set of control variables, including indicators for firm i employing a university-educated engineer or having a merchant on its board of executives, the growth rates of both high-end and low-end machine capacity (spindles) between t and $t+1$, and logged output at time t . This panel regression with firm fixed effects γ_i captures firm-specific knowledge accumulation through upgrade experiments and the transfer of that knowledge into expansion of the firm's low-end product portfolio. The parameters of interest are β_1 and, especially, that on the interaction term, β_3 . A positive β_3 is consistent with a complementarity between product upgrading experimentation and product diversification.

Table 6 reports the results. In column (1) we simply examine whether the cumulative number of upgrade experiments raises firms' growth rates. The coefficient on the cumulative number of upgrade experiments is positive but not statistically significant at conventional levels. Employing a degreed engineer is associated with a 9.6 percent increase in the firm growth rate, other things equal. This corresponds to about half of the gap in growth rates between the 25th and the 75th quartiles in the sample. In contrast, having a merchant on the firm's board of executives is unrelated to firm growth.

In the specification in column (2) we add the interaction of the cumulative number of upgrade experiments and the fraction of low-end products. While the coefficient on cumulative upgrade experiments is economically and statistically indistinguishable from zero, the coefficient on the interaction is positive and significant. Given that the mean number of cumulative upgrade experiments in the sample is 1.14 and the mean fraction of low-end products is 0.84, the coefficient of 0.04 implies an additional upgrade experiment at the mean low-end product fraction is associated with a 3.4-percentage-point higher output growth rate. Similarly, increasing the fraction of low-end products by one standard deviation (0.27) is associated with a 1.2 percentage point higher growth rate at the mean number of cumulative upgrade experiments. These results indicate the link between upgrade experiments and growth indeed appears to act through a mechanism tying experimentation with vertically differentiated products to the extent of the horizontal differentiation of the firm's product offerings.

--- Table 6 around here ---

While looking within firm as above controls for several possible confounds, it is still possible that a firm may time its product upgrading experiments with other unobservable factors related to firm growth. To obtain a better sense of the causal connection between upgrade experimentation and firm growth, as discussed above we use the interaction of high-end machine arrival and installation with mandatory output cuts to gain a source of exogenous variation in upgrade experimentation. Here, we employ continuous measures of each. For machine installation, we use the growth rate of high-end machine capacity due to new installations during period t . For output cuts, we use the size of the required production drop in low-end products. The resulting instrument is the product of these two values, and it varies both intertemporally and in the cross section.

Because the endogenous variable is cumulative high-end experiments but the logic of our exogenous variables works contemporaneously (it predicts current-period experiments), we estimate our IV specification in two stages. First, similar to Table 3 above, we use a Poisson regression to obtain a predicted number of upgrade experiments conducted by a firm in a particular period. The model uses the aforementioned instrument as well as other exogenous variables like firm age. Next, we sum these period-specific predicted values to construct the firm's predicted cumulative upgrade experiments. This constructed value is the key explanatory variable in our causal inference regressions.

We present the first stage regression results in column (1) of Table 7. The dependent variable is the number of upgrade experiments the firm starts in period t , while the explanatory variables include the instrument, which is itself an interaction, and the two main effects. The coefficient on the interaction is positive and significant. Firms that expanded high-end machines during times of output cuts were in fact pushed toward conducting (more) upgrade experiments.

--- Table 7 around here ---

As a placebo test to help verify the validity of the instrument, we conducted a similar regression using the growth rate in low-end machine capacity instead. The results are in column (2). As expected, there is no significant relationship between upgrade experiments and the interaction in this case.

We use the constructed number of upgrade experiments to estimate the effect of such experiments on firms' output growth. The results are in columns (3) and (4) of Table 7. The coefficient on the predicted cumulative number of upgrade experiments in column (3) is positive but small and statistically insignificant, similar to column (1) in Table 6. However, when we include in column (4) the interaction of constructed cumulative upgrade experiments and the fraction of low-end products, the interaction's coefficient is twice as large as the corresponding coefficient in Table 6 and more precisely estimated.

That there appear to be growth effects of product upgrade experimentation raises the question of why firms in our sample did not engage it in all the time (or at least, more often). A couple points are worth noting. First, there could be substantial firm heterogeneity in expected growth effects. In our empirical results above, we condition on the firm having installed high-end machinery. As seen in Table 4, these firms might be systematically different from those that did not. One of these differences could be the expected return to upgrade experimentation. To

examine this further, we also conducted panel firm-fixed effect estimation similar to that in Table 6 but separately for firms that already had high-machinery installed. Remarkably, the coefficients measuring growth effects of specifically upgrade experiments conducted by firms that already had high-end machines were very similar in magnitude and statistical significance to those from the second-stage regression in Table 7 despite very different estimation methodologies (details are available upon request).

Second, experimentation likely also has costs, even among firms with high-end machines. A particular cost relevant in the context of the instrumental variable is the opportunity cost of the labor and capital inputs the firm uses to conduct experiments. Part of this opportunity cost is the value of the low-end product output (and any associated future growth tied to this) that those inputs would create in absence of experimentation. Effectively, the instrument uses instances where the firm faces an exogeneously sourced temporary reduction in those opportunity costs. This pushes the net benefit of upgrade experimentation into positive territory, inducing experimentation and the gross growth benefits it creates. We are measuring the average growth benefit across firms whose upgrade experimentation benefit-cost calculus is close enough to the margin for the instrument to spur them into action.

6. Transmission channels

How exactly does technical knowledge acquired through experimentation with upgrade products lead to the expansion of low-end product varieties and firm output growth? We consider two possible channels: increased flexibility of the firm's production system, including in low-end products, and increased demand-side appeal of low-end products.

6.1 Production system flexibility as a transmission channel

Recall that besides being differentiated by counts, product varieties in our data are also distinguished by the direction of twist (S-twist and Z-twist single yarn) and whether yarn is further doubled and/or gassed. These different types, even within the same count category, serve different needs of fabric weavers. During our sample, S-twist yarn was often associated with weft while Z-twist was often associated with warp.²⁰ S-twist and Z-twist yarn also differ in

²⁰ In the process of weaving textiles from yarn, the longitudinal warp yarns remain stationary on a frame or loom, while the transverse weft is drawn through and inserted over and under the warp.

strength and softness (Z-twist is stronger, but S-twist is softer to the touch). Doubled and gassed yarn serve still other weaving purposes. As demand conditions change, firms that can flexibly switch across different twist directions within a given count or across adjacent counts are better able to respond to such changing conditions, and thus have better growth opportunities. We examine if the knowledge capital accumulated through upgrading experiments contributed to firms developing more flexible production systems.

To construct an empirical measure of production system flexibility, we use monthly data to construct two measures of the frequency of a firm's product portfolio rebalancing. The first counts how often the firm changes the "lead direction" of its yarn of the same count. For example, suppose that both firm A and firm B produce 16-count yarn. Suppose further that firm A produces 80 percent of its 16-count output in S-twist and 20 percent in Z-twist in both t and $t-1$. Firm B, on the other hand, produces 80 percent of its 16-count output in S-twist and 20 percent in Z-twist in $t-1$, but then changes to 80 percent in Z-twist and 20 percent in S-twist in t . We say that firm B rebalanced its portfolio of 16-count yarn between $t-1$ and t , while firm A did not. More generally, for each count category, we define the "lead direction" to be the way in which the majority of yarn in the count category was processed (S-twist, Z-twist, doubled or gassed). We use monthly data to count the number of times the firm changed this lead direction in any given semi-annual period. Adjusting the lead direction is time-consuming (see appendix A.7), so we infer that firms that did such adjustments more frequently had a more flexible production system.

A second measure of portfolio rebalancing counts the number of times a firm changed its "lead" count category within each direction (S-twist, Z-twist, doubled and gassed yarn). Suppose firms A and B produce two count categories of S-twist yarn, 16 and 20 counts. Firm A produces 80 percent of its output of S-twist yarn as 16-count and 20 percent as 20-count in both t and $t-1$. Firm B, on the other hand, produces 80 percent of its output of S-twist yarn as 16-count and 20 percent as 20-count in $t-1$, but then switches in t to 80 percent as 20-count and 20 percent as 16-count. Once again, we say that firm B rebalanced its portfolio of S-twist yarn between $t-1$ and t , and firm A did not. This operation is also time-consuming, so our second measure of a firm's production system flexibility is the number of times it changed the "lead count" category in a given semi-annual period.

To address the role played by knowledge capital accumulated through product upgrade experiments as cleanly as possible, we limit our product flexibility measure to low-end products (up to 20 count). We examine how the two measures of product portfolio rebalancing at the low-end of the product variety spectrum were associated with the cumulative number of product upgrade experiments, which by construction involve only high-end products, and product diversification experiments, which can involve low-end products. (All the findings presented below are robust to using product portfolio rebalancing measures over all product varieties.)

Table 8A reports estimation results from within-firm panel regressions. The estimation equation is

$$y_{it} = \alpha + \beta_1 cuml_upgrade_exp_{it-1} + \beta_2 cuml_divers_exp_{it-1} + \beta_3 X_{it} + \gamma_i + \Delta_t,$$

where y_{it} are the two measures of firm-specific production system flexibility described above. The explanatory variables of interest are again the cumulative numbers of upgrade and diversification experiments conducted by firm i by time $t-1$. We include as control variables X_{it} indicators for high-end and low-end machine expansions in $t-1$ as well as for the presence of a university-educated engineer and a merchant on the board of firm i at time t . The regressions also include semi-annual time dummies and firm fixed effects. The firm fixed effects mean we are measuring the within-firm relationship between conducting more upgrade/diversification experiments and its production system flexibility in the low-end product space.

The results strongly support the conjecture that knowledge accumulated through product upgrade experiments contributed to greater production system flexibility in low-end product varieties. After firms conducted vertical upgrading experiments, they rebalanced their portfolios of low-end product varieties more frequently, both across twist directions within a given count category and across count categories within a twist direction. The relationship was stronger for within-count lead direction changes than for within-direction lead count changes (e.g., switching from 20-count S-twist to 20-count Z-twist rather than switching from 16-count S-twist to 20-count S-twist).

The coefficient on the cumulative number of upgrade experiments in the within-count portfolio rebalancing regression (the first column in Table 8A) is 0.24, while the mean number of within-count portfolio rebalancing events is 0.5. Hence an additional upgrade experiment was associated with an increase in portfolio rebalancing within counts by 50 percent of its mean. The coefficient on the cumulative number of upgrade experiments in the across-count portfolio

rebalancing estimations (the second column in Table 8A) is 0.16, implying that an additional upgrade experiment is associated with an increase in portfolio rebalancing across counts by about 13 percent of its mean (equal to 1.2).

Remarkably, experience conducting diversification experiments had no such effects, even though diversification experiments involve low-end products. Once again, vertical product expansions appear to have spillovers into production capabilities that horizontal expansions do not. These spillovers do not appear to be due the direct operation of the (indeed more flexible) high-end machines themselves. There is no obvious advantage of high-end machines for within-count direction rebalancing, and we control for high-end machine expansion in the regressions. Accumulated technical knowledge from product upgrade experiments, not the presence of high-end machines per se, appears responsible for more flexible within-count portfolio rebalancing.²¹

--- Table 8A around here ---

Table 8B confirms that these transmission channels of upgrade experiments contributed to firm growth. A rise in portfolio rebalancing from t to $t+1$, both within- and across-count, was economically and statistically significantly associated with accelerations in firm growth. One extra within- (across-) count portfolio rebalancing is associated with 1.0 (1.2) percentage point higher growth rates. Because the distribution of rebalancing events across firms and time is highly skewed, a better sense of the magnitudes might be obtained by calculating the effect of a firm moving from doing no portfolio rebalancing to the mean rebalancing level conditional on rebalancing being positive. This conditional mean of within- (across-) count portfolio rebalancing events is 2.33 (2.35). Moving from no within-count (across-count) portfolio rebalancing to its conditional mean is thus associated with a 2.33 (2.82) percentage point faster growth rate of total output. If we include both upgrade experiments and portfolio rebalancing in the same growth regressions (not shown), the coefficients on cumulative past upgrade

²¹ This link can be put in a broad historical context. At the start of our sample, the lion's share of all yarn produced by Japanese firms was S-twist. This changed dramatically over time. Figure A4 in appendix A.7 plots the dynamics of the fraction of non-S twist varieties in the total number of product varieties. The fraction of non-S-twist varieties was less than 30 percent of the total in both high-end and low-end product varieties early on, but producers of high-end product varieties quickly switched to almost exclusively Z-twist, doubled, and gassed yarn. The fraction of non-S-twist among low-end product varieties, on the other hand, was at about 40 percent at the turn of the 20th century and gradually increased to about 60 percent toward the end of the sample, with firms producing high-end products leading the way. The mean fraction of non-S-twist varieties in the total number of low-end varieties over the sample was 0.48 for firms that produced at least one high-end product, but only 0.28 for firms that did not produce high-end products. This difference is statistically significant. Experience with high-end products enabled the broadening of the product varieties range, creating more flexible production systems at the low end (cf. Roberts, 2004, pp. 37-38).

experiments in Table 6 and those on the number of portfolio rebalancing events in Table 8B remain the same. The two relationships therefore operate independently of each other.²²

--- Table 8B around here ---

6.2 Quality of low-end products as a transmission channel

In this section we examine how a firm's experience in upgrading experimentation affects the demand for its low-end products. To measure this demand appeal, we utilize the Khandelwal (2010) method of estimating quality for horizontally differentiated products, which essentially boils down to looking at relative market shares after adjusting for price differences. While we do not have firm-level price data for most product varieties, we do for a key 20-count yarn for about 40 percent of observations.²³

As described above, count and finish completely summarized the physical aspect of the product. However, some horizontal differentiation existed due to reasons of geography or idiosyncratic buyer-supplier relationships. Most firms did in fact have their own registered brands, including for 20-count yarn (e.g., Kanegafuchi Spinning's "*Rangyo* (Indigo Fish)" brand, Settsu Spinning's "*Kujyaku* (Peacock)" brand, etc.). That said, given the physical equivalence of the products, within-count substitution across brands was very high and there was very modest equilibrium price variation.²⁴ This makes it hard to precisely estimate the effect of price on demand, but we tried anyway.

²² Because upgrade experiments and the associated more flexible production system allow firms to better respond to changing demand, these should also be associated with higher capacity utilization rates. We know firms' installed spindle capacities as well as the number of spindles they had in operation each period, so we computed capacity utilization and confirmed that it positively covaries with our product system flexibility measures. Each additional lead change (both within- and across-count) is associated with about 2.4 percent increase in capacity utilization rates. Details are available upon request.

²³ 20-count yarn was the industry's modal output, accounting for 27.6 percent of production over the sample, but its importance goes beyond this. When the industry was still in its infancy, Japanese firms could not produce counts higher than 16 because they were limited to poor quality domestic and Chinese cotton. Indian cotton imports arrived at the end of the 1880s, and in 1890 the Osaka Spinning Company exported the first experimental batch of 20-count yarn to China. This marked the transition to a globally competitive industry (Kinugawa, 1961, Vol. 4, Ch. 1). Even as new high-end products were developed, 20-count yarn remained the most important low-end product. It was one of only two yarn counts listed on the Osaka Three Articles Exchange (alongside 16 count), and it totally dominated trade volume (95 million yen gross trade volume in 1914 as opposed to a mere 1320 yen for the 16 count).

²⁴ Table A11 in appendix A.8 shows the coefficient of variation of 20-count prices was just 2.8 percent across the entire sample, and the interquartile dispersion coefficient was a meager 1.6 percent. Even the 90-10 percentile dispersion coefficient was only a bit over 3 percent. All of these were an order of magnitude smaller than, for instance, the corresponding statistics of wages of female production workers in the industry.

Because 20-count yarn is also just at the borderline between low-end and high-end products, we can instrument for its price using a plausible cost shifter that would be difficult to obtain for other product types in our data. Specifically, we use our portfolio rebalancing measure introduced in the previous section, here focusing on counts around 20, and interact it with the degree of industry-wide mandatory low-end product output cuts. As noted above, rebalancing the portfolio entailed adjustment costs, and firms that had lower costs could adjust more frequently. Firms with more flexible production systems could therefore respond more easily to mandated cuts by shifting their production to high-end products. This implies, especially in the face of having to meet short-run fixed costs, that more flexible firms would feel less pressure to reduce the 20-count price during periods of slow demand.²⁵ This gives us a supply-side source of price variation that is plausibly uncorrelated with quality (demand appeal). Thus our instrument for the price of 20-count yarn interacts our across-count portfolio rebalancing measure for counts from 17 to 48 (that is, excluding very low and very high counts) with the time-varying degree of mandatory output cuts.

We present the first-stage estimates in the first column of Table 9A. Lower portfolio adjustment costs, reflected in more frequent portfolio rebalancing by a firm, have a small, negative correlation with the firm's 20-count price during normal times. However, rebalancing is strongly positively associated with the firm's 20-count price during periods of mandatory output cuts. This holds controlling for both firm and time fixed effects. Going from no portfolio rebalancing to its mean conditional on being positive and mandated output cuts being in place (0.38) is associated with 1.2 percent higher 20-count price. That is small in absolute size but, given the price variation described above, is two-thirds of the interquartile dispersion during periods of output cuts.

Recall that in the previous section we also constructed a second measure of portfolio rebalancing, across different directions (S-twist, Z-twist) within the same count. While this measure was associated with lower adjustment costs and firm growth (see Tables 8A and 8B above), there is no reason why lower adjustment costs within the same count should be relevant for keeping up the price of the 20 count during mandatory output cuts. Based on this logic, we conducted a “placebo test” by looking at the relationship between the 20-count price and the

²⁵ The average price of 20-count yarn was significantly lower (between 2-9 percent) during mandatory output cut periods than in adjacent periods without cuts.

interaction of within-count portfolio rebalancing and mandatory output cuts. The estimation results are in the second column of Table 9A. The logic holds; the relationship is statistically and economically indistinguishable from zero.

The second stage regression of the firm's (logged) share of industry-wide 20-count output on the instrumented logged price is shown in Table 9B. The estimated own-price elasticity is negative and large in magnitude, about -5.4. This is line with high substitutability across horizontally differentiated brands. That said, it is worth noting that the standard errors on these elasticity estimates are high, most likely because of low price variation in the sample discussed above. For the sake of comparison, in appendix A.8, Table A12 we present the results of an OLS estimation of firms' market shares of 20 count on own price. This "naïve" regression produces an own-price elasticity of -3.1, so the IV estimation, while imprecise, moves the estimated elasticity in the theoretically predicted direction.

We use the demand estimates to construct the Khandelwal (2010)-style quality (demand appeal) measure for each firm's 20-count product in each period. This involves adding the estimated coefficients for the corresponding firm and time fixed effect to the demand residual.²⁶

We then use this demand appeal metric as a dependent variable and examine whether past experimentation is related to it. As before, we capture the temporal complementarity between product upgrading and horizontal diversification by interacting the cumulative number of upgrade experiments and the fraction of low-end products in the total number of product varieties produced by a firm. The estimation equation is:

$$y_{it} = \alpha + \beta_1 cuml_upgrade_exp_{it} + \beta_2 low_frac_{it} + \beta_3 cuml_upgrade_exp_{it} \times low_frac_{it} + \beta_4 X_{it} + \Delta_t + \varepsilon_{it}, \quad (3)$$

where y_{it} is the quality (demand appeal) measure of firm i 's 20 count at time t constructed as above; $cuml_upgrade_exp_{it}$ is the cumulative number of upgrade experiments conducted by the firm up to time t ; low_frac_{it} is the fraction of low-end products in the total number of products the firm makes; and X_{it} is a set of control variables including indicators for the firm employing a university-educated engineer or having a merchant as a board member, the growth rates of high-end and low-end machine capacity between t and $t+1$, and firm age. As in the growth regression (2), the parameters of interest are β_1 and β_3 . If experimentation with vertical

²⁶ In our setting, Khandelwal's "outside good" is an amalgam of cotton yarns of all other counts.

product upgrading (the complementarity between experimentation with vertical product upgrading and product diversification) positively affects the demand appeal of 20-count yarn, we expect β_1 (β_3) to be positive.

--- Tables 9A-9C around here ---

Table 9C presents the results. A firm's upgrade experiments and their interaction with subsequent low-end product diversification are positively related to the firm-specific demand appeal in 20-count yarn. The estimated magnitude is largest when we instrument for upgrade experimentation, similarly to Table 7 above. Interestingly, the coefficient on the indicator for a firm employing a university-educated engineer is also positive and statistically significant in all specifications. Engineering talent apparently is independently associated with demand appeal.

Historical materials provided us with an opportunity to check the external validity of our estimation by comparing our quality measure to quality rankings of 20-count yarns published by the Osaka Three Articles Exchange in 1907. As can be seen from Figure A5 in appendix A.8, the two quality indices are by and large consistent with each other, with a correlation of 0.69.

7. Robustness

7.1 Survival rates

Our results indicate that firms that installed high-end machines and conducted product upgrade experiments grew faster. While this might seem to reflect a superior outcome for firms, it is possible that upgrade experimentation actually increased the variance of outcomes rather than raising their mean. If so, we would observe faster growth conditional on survival, but this would be balanced against a lower chance of surviving for experimenting firms. While such a variance increase might still be preferable for firms (limited liability companies are essentially a call option, after all), this is a qualitatively different mechanism than one where upgrade experimentation simply raises expected growth. To explore this possibility, we investigate survival patterns in our data.

In Table 10, we present summary statistics on the status of firms at the end of our sample in 1914. A total of 33 firms survived to the end of the sample. Of these, 19 (58 percent) had high-end machines. This compares to 42 of 105 firms for which we have machine data (40 percent) having high-end machines over the entire sample. Firms with high-end machines

therefore had a substantially higher probability of survival. This is inconsistent with product upgrade experimentation primarily increasing the variance of firms' growth rates.

Moreover, among the 72 firms that left the sample, we can distinguish exits by acquisition (53 firms) and liquidation (19 firms). Firms that exit by acquisition are more likely to have high-end machines than firms that shut down, as seen in the table. High-end machines not only were associated with a greater chance of acquisition, but they also improved shareholders' returns conditional on being acquired. We have data on acquisition prices for 46 acquisition cases. In 18 of these cases, the acquired firm had high-end machines. We computed the "salvage fraction" of shareholders paid-in capital by dividing the acquisition price by the shareholders paid-in capital. The mean salvage fraction was 1.04 for acquired firms with high-end machines but only 0.70 for those that did not (a statistically significant difference at the 5 percent level).

--- Table 10 around here ---

7.2 Mergers and acquisitions

Firms can expand the range of their product varieties and scale up production through mergers and acquisitions (M&A). Here we examine the possibility that M&A generate some of the patterns of complementarity between product upgrading and diversification we document above.

Using the industry M&A data from our earlier work (Braguinsky et al., 2015), we find that new product introductions do increase at the time of acquisition events (the probability of a product addition increases by about 25 percent), but acquisitions themselves are rare events compared to product introduction events. As a result, just 3.6 percent of new product introductions in our sample coincide with an acquisition event. Thus, while there is some uptick in new product introductions at the time of acquisitions, acquisitions played a minor role in product variety expansion during our sample. Also, if we include an acquisition event dummy in growth regressions like those in Table 6 above, the coefficient is not statistically significant at conventional significance levels. In growth regressions similar to Table 7, acquisitions are positively related to firm growth (as to be expected), but the effect of temporal complementarity between product upgrading and diversification on firm growth remains qualitatively unchanged even after controlling for merger and acquisition events.

7.3 Market competition

Increased competition in high-end product markets over time might induce firms to renew their attention to low-end product markets. This presents a potential alternative explanation for the “first experiment with upgrades, then turn to diversification” temporal relationship seen several ways above. We examine several market competition measures from our data to look for suggestive evidence regarding this hypothesis.

Figure A6 in appendix A.9 depicts the evolution of the market shares of firms operating in low- and high-end product markets, the industry-level ratio of high-end to low-end output, and the average real price of 20-count yarns (as a proxy for the average price of low-end products). While the fraction of firms producing in low-end product markets is stable throughout our sample period after some initial volatility, the fraction of firms in high-end product markets increases sharply between 1899 and 1904 and fluctuates thereafter. Our decomposition analysis in Table 2, on the other hand, showed that most of the product variety expansion into low-end product markets took place after 1907, so the timing of product variety expansion into low-end product markets and the timing of rapid entry (and increased competition) in high-end product markets do not match. Also, after falling sharply in the 1890s, the average real price of 20 count yarns fluctuates around 60-70 yen after that, without any discernible upward or downward trend. Thus it is difficult to infer that low-end product markets became more attractive after 1907, when firms started expanding low-end product varieties and scaling their production.

We also conducted firm-level regression analysis with new product introductions as the dependent variable, and the number of firms in low-end and high-end product markets as well as the average real price of 20-count yarns as explanatory variables. None of these variables were related to new product introductions at conventional significance levels. Overall, we do not see much supportive evidence for the hypothesis that heightened competition in high-end product markets pushed firms toward low-end product markets.

7.4 The role of exports

We explored whether the motive to operate in export markets might have driven both upgrade experimentation and firm growth.

As has been observed in many other settings in the literature, firms that export in our sample are larger on average than non-exporting firms. However, cotton spinning firms that

export more (as a share of their output) produce fewer product varieties, especially high-end varieties. Hence it does not appear that exporting was a driving factor in our results above.

The nature of industry exports during our sample explains this. Japanese cotton spinners successfully drove out imports and started exporting low-end products to East Asian markets during the 1890s. During those years, and to a large degree after that as well, exports were concentrated in a few low-end products (especially S-twist 16 count and Z-twist 20 count yarn). Major exporting firms tended to be focused on scaling their output in and around these product varieties. More often than not, they chose to forgo opportunities in high-end markets.

8. Conclusions and Discussion

Using detailed historical panel data from the Japanese cotton spinning industry, we find firm growth is associated with increased number of product varieties, but the identity (type) of product varieties matters. High-growth firms followed a particular pattern of product variety expansion. They first went outside of their existing technological frontiers and experimentally introduced innovative products, sometimes successfully but often not. Either way, they subsequently increased the number of product varieties they produced inside their frontier. In other words, they first engaged in attempted vertical product differentiation and, later, horizontal differentiation.

The consistency of this pattern reflected spillovers that product upgrade experiments had onto the horizontal diversification of firms' product sets. Experimentation required firms to overcome technological constraints and cope with uncertainty by investing in new types of machines and hiring educated engineers. These newly developed inputs and their associated technical knowledge were broadly applicable—useful not just for upgrading but also for producing new varieties within the firms' existing technological capabilities. This process was the major driver of industry firms' growth. In contrast, conducting only horizontal product diversification (experiments) without upgrade experiments did not generally lead to sustained growth in either the number of product varieties produced nor output.

We identified at least two specific channels through which accumulated technical knowledge contributed to higher firm growth throughout the product space: increased flexibility of the production system and improved quality of low-end products. Both of these were strongly tied to product upgrade experimentation.

Standard models do not consider the relationship between product variety expansion and firm growth uncovered in our study. Models of endogenous growth through product variety expansion (e.g., Romer, 1990) predict that high-growth firms are those that keep introducing new product lines, all of which are of the same type. Quality ladder models (e.g., Grossman and Helpman, 1991; Klette and Kortum, 2004) predict that high-growth firms are those that keep generating innovative products which are upgrades over existing versions, but vertical differentiation happens only within, not across product lines, and innovation is often untargeted. In these models, technically more difficult innovative products are positively associated with consumer preferences, i.e., quality. In both these versions of endogenous growth theory, any past product introduction contributes to the accumulation of knowledge capital and serves as a determinant of future product expansion and firm growth. In our setting, in contrast, introducing technologically more challenging products relates much more strongly to growth than simple horizontal product proliferation. This highlights the importance of incorporating into analytical frameworks heterogeneity in the particular ways and directions product variety expansion occurs.

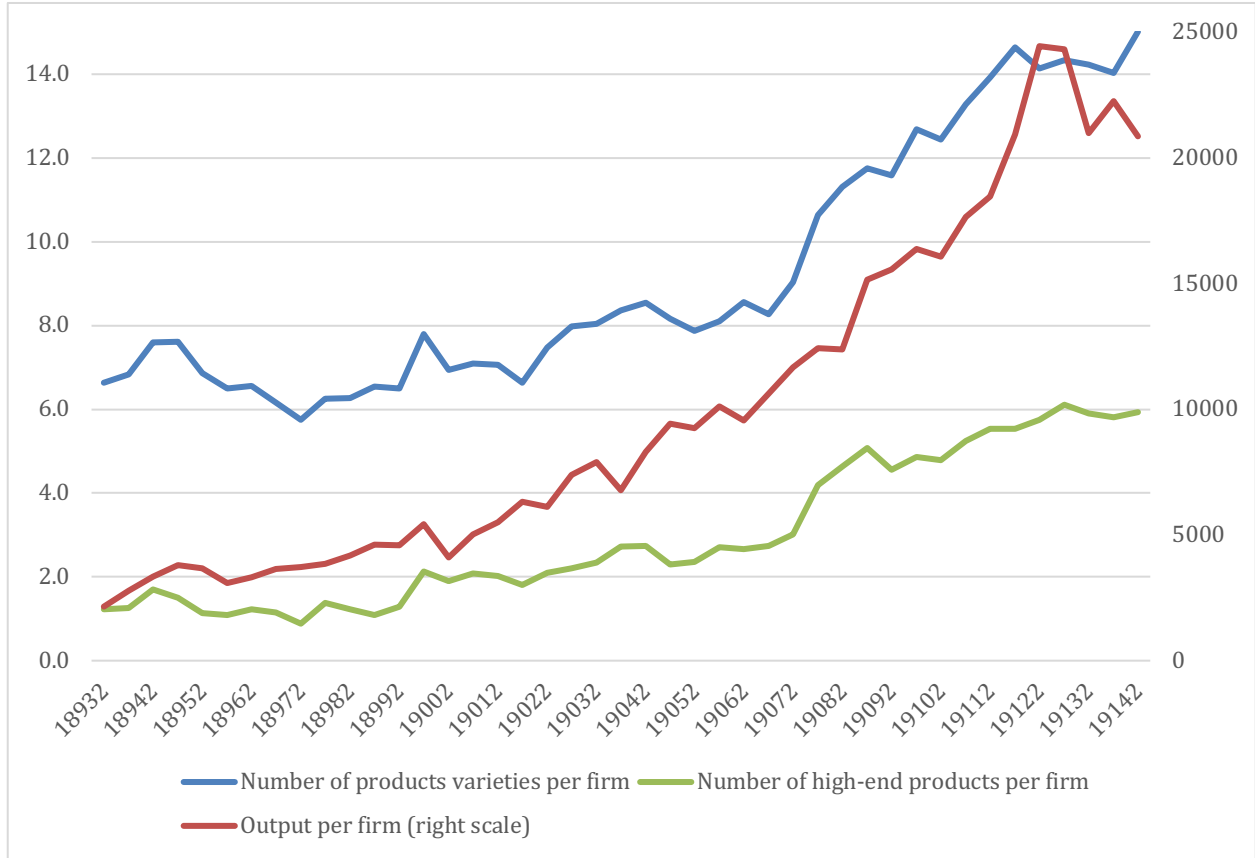
References

- Bernard, Andrew, Stephen Redding, and Peter Schott, 2010. “Multiple-Product Firms and Product Switching.” *American Economic Review* 100(1): 70-97.
- Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson, 2015. “Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry.” *American Economic Review* 105(7): 2086-2119.
- Callander, Steven, 2011. “Searching and Learning by Trial and Error.” *American Economic Review*, 101(6): 2277-2308.
- Cusolito, Ana Paula, and William F. Maloney, 2018. *Productivity Revisited. Shifting Paradigms in Analysis and Policy*. The World Bank Group.
- Enkakukiji (Outline of the Cotton Spinning Industry, in Japanese). 1901. Japan Ministry of Agriculture and Trade. Unpublished.
- Geppo (Monthly Bulletin of the Japanese Cotton Spinners’ Association, in Japanese). 1889–1920. Dainihon Boseki Rengokai: Osaka, Japan.
- Grossman, Gene and Elhanan Helpman, 1991. *Innovation and Growth in the Global Economy*. MIT Press, Cambridge, MA.
- Hottman, Colin, Stephen Redding, and David Weinstein, 2016. “Quantifying the Sources of Firm Heterogeneity.” *Quarterly Journal of Economics*, 131(3): 1291-1364.
- Karni, Edi, and Marie-Louise Vierø, 2013. “‘Reverse Bayesianism’: A Choice-Based Theory of Growing Awareness.” *American Economic Review*, 103(7): 2790-2810

- Kerr, William, Ramana Nanda, and Matthews Rhodes-Kropf, 2014. "Entrepreneurship as Experimentation." *Journal of Economic Perspectives*, 28(3): 25-48.
- Khandelwal, Amit, 2010. "The Long and Short (of) Quality Ladders." *Review of Economic Studies*, 77(4): 1450-1476.
- Kinugawa, Taiichi. 1964 (reprinted from the original 1942 edition). *Hompo Menshi Boseki Shi* (History of the Japanese Cotton Spinning Industry, in seven volumes, in Japanese). Hara Shobo: Tokyo, Japan.
- Klette, Tor, and Samuel Kortum, 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy*, 112(5): 986-1018.
- Klepper, Steven, and Peter Thompson, 2006. "Submarkets and the Evolution of Market Structure." *Rand Journal of Economics*, 37(4): 861-886.
- Penrose, Edith, 1959. *The Theory of the Growth of the Firm*. Blackwell, Oxford, U.K.
- Roberts, John, 2004. *The Modern Firm*. Oxford University Press, Oxford, U.K.
- Romer, Paul, 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5, Pt. II): S71-S102.
- Sankosho, 1903-1914. *Menshi Boseki Jijo Sankosho (Reference Materials on the State of the Spinning Industry*, in Japanese). Dainihon Boseki Rengokai: Osaka, Japan
- Saxonhouse, Gary, 1974. "A Tale of Japanese Technological Diffusion in the Meiji Period." *Journal of Economic History* 34(1): 149-165.
- Shoji, O., 1930. *Boseki Sogyo Tanshuku shi* (History of Operational Curtailments in Cotton Spinning, in Japanese). Osaka, Nihon Mengyo Kurabu.
- Stokey, Nancy L., 1988. "Learning by Doing and the Introduction of New Goods." *Journal of Political Economy* 96(4): 707-717.
- Takamura, Naosuke. 1971. *Nihon Bosekigyoshi Josetsu* (Introduction to the History of the Japanese Cotton Spinning, in Japanese), in 2 volumes. Hanawa Shobo: Tokyo, Japan.
- Tokei Nempo, 1915. *Osaka Shogyo Kaigisho Tokei Nenpo, 1914 (Osaka Chamber of Commerce Annual Statistical Report, 1914 issue, in Japanese)*. Osaka Chamber of Commerce, Osaka.
- Woodhouse, Thomas, 1921. *Yarn Counts and Calculations*. Oxfor Technical Manuals, R. & R. Clark, Edinburgh.
- Wright, Gavin, 2011. *Cotton Spinning Machinery Orders, British Textile Machinery Firms, 1878-1933*. ICPSR27141-v1. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2011-11-18. doi:10.3886/ICPSR27141.v1.

Figures and Tables

Figure 1. Dynamics of output and number of product varieties per firm



Source: Our calculations using the data described in the main text and in the appendix. The number of products per firm is weighted by firm's share in total industry output at any given time.

Table 1. Breakdown of product lines and products by experimental and nonexperimental

Panel A. Product lines				
	All (1)	Fraction in total	Never scaled (2)	Ratio: (2)/(1)
New product lines	685	1.000	271	0.396
Of which: never experimental	246	0.359		
initially experimental	439	0.641	271	0.617
Of which: Upgrade lines	76	0.111	39	0.513
Panel B. Experimental products				
	All	“Successful” (scaled)	“Failed” (not scaled)	Fraction “failed”
All experimental products	819	223	596	0.728
Of which: upgrade experiments	116	42	74	0.638
diversification	703	181	522	0.743
Fraction upgrades	0.142	0.188	0.124	

Source: Our calculations using the data described in the main text and in the appendix. “New” product lines are those that were not in the firm’s set of product varieties produced at entry or at the time of the first observation.

Table 2. Decomposition Analysis

Period	Total change	Composition	Within (Deviation from own average)			
			Total	Continuing firms	Entrants	Exiting firms
1893.2-1914.2	9.0	3.3	5.6	6.7	-0.2	-0.8
1893.2-1906.2	2.5	1.6	0.9	1.0	-0.3	0.1
1907.1-1914.2	6.4	1.7	4.7	5.6	0.0	-0.9

Source: Our calculations using the data described in the main text and in the appendix.

Table 3. Product variety expansion as a function of past upgrade experiments

VARIABLES	DV: # of all products at t , minus # of all products at $t-1$ (1)	DV: # of high-end products at t , minus # of high-end products at $t-1$ (2)	DV: # of low-end products at t , minus # of low-end products at $t-1$ (3)	DV: # of all products at t , minus # of all products at $t-1$ (4)	DV: # of high-end products at t , minus # of high-end products at $t-1$ (5)	DV: # of low-end products at t , minus # of low-end products at $t-1$ (6)
Cumulative number of upgrade experiments at $t-1$	0.433*** (0.093)	0.258*** (0.059)	0.202*** (0.074)	0.384*** (0.090)	0.239*** (0.058)	0.174** (0.073)
Cumulative number of diversification experiments at t	0.018 (0.039)	0.001 (0.009)	0.014 (0.034)	0.022 (0.036)	0.003 (0.008)	0.017 (0.032)
Dummy equal to 1 if high-end machine expansion during				0.659*** (0.214)	0.269** (0.128)	0.389** (0.149)
Dummy equal to 1 if low-end machine expansion during				0.216 (0.137)	0.002 (0.046)	0.213* (0.113)
Dummy =1 if university-educated engineer employed at				0.370* (0.190)	0.272* (0.138)	0.121 (0.142)
Dummy =1 if merchant a member of board at t				0.193 (0.151)	0.041 (0.064)	0.148 (0.114)
Number of all products at $t-1$	-0.405*** (0.041)			-0.408*** (0.040)		
Number of high-end products at $t-1$		-0.455*** (0.045)			-0.462*** (0.042)	
Number of low-end products at $t-1$			-0.385*** (0.049)			-0.385*** (0.049)
Constant	2.740*** (0.609)	0.686*** (0.226)	2.007*** (0.470)	2.375*** (0.626)	0.558** (0.239)	1.767*** (0.486)
Semiannual time and observation dummies	Included	Included	Included	Included	Included	Included
Firm FE	Included	Included	Included	Included	Included	Included
Observations	1,618	1,618	1,618	1,618	1,618	1,618
Within R-squared	0.221	0.241	0.225	0.235	0.252	0.234
Number of firms	99	99	99	99	99	99

Panel data estimation with firm fixed effects. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4. Firms with and without High-End Machines: Market Ties, Educated Engineers, and Experimental v. Non-Experimental Product Introductions

	High-End Machines	No High-End Machines
Fraction of firms with a merchant as a board member	0.71	0.57
Number of university-educated engineers employed	1.61	0.14
Number of technical college-educated engineers employed	4.76	0.72
New experimental product introductions: All	0.60	0.25
Of which, fraction that are successful	0.21	0.35
New experimental product introductions: Upgrade	0.10	0.03
Of which, fraction that are successful	0.36	0.35
New experimental product introductions: Diversification	0.50	0.22
Of which, fraction that are successful	0.18	0.35
New non-experimental product introductions: All	0.11	0.11
New non-experimental product introductions: Upgrade	0.06	0.01
New non-experimental product introductions: Diversification	0.05	0.09

Notes: The differences in means of new experimental product introduction are highly statistically significant between firms with and without high-end machines using double-sided *t*-test. The differences in means of fractions of successful experiments are not statistically significant for upgrading whereas they are statistically highly significant for diversification. The differences in the fraction of firms with a merchant as a board member and in the means of the number of university- and technical college-educated engineers are highly statistically significant between firms with and without high-end machines.

Table 5. Factors affecting the number of experiments a firm started during period t

Panel A. DV: Product upgrade experiments						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	All firms				Firms with high-end machines	
Dummy equal to one if had high-end machines in period t	1.066*** (0.378)					
Dummy equal to 1 if high-end machine expansion during period t		1.493*** (0.305)		0.823** (0.366)	1.098*** (0.314)	0.464 (0.352)
Interaction term between high-end machines expansion and mandated output cuts dummies during period t				1.012* (0.580)		1.703*** (0.609)
Dummy equal to 1 if low-end machine expansion during period t		-0.266 (0.359)		-0.386 (0.405)	-0.474 (0.459)	-0.458 (0.456)
Dummy =1 if university-educated engineer employed at t			0.715* ((0.408)	0.568 (0.426)		0.509 (0.497)
Dummy =1 if merchant a member of board at t			1.521*** (0.432)	1.455*** (0.414)		1.623*** (0.567)
Firm age	0.007 (0.026)	0.018 (0.022)	-0.009 (0.030)	-0.011 (0.031)	0.007 (0.028)	-0.033 (0.033)
Constant	-2.396*** (0.485)	-2.268*** (0.422)	-3.501*** (0.580)	-3.448*** (0.562)	-1.695*** (0.518)	-3.132*** (0.851)
Semi-annual time dummies	Included	Included	Included	Included	Included	Included
Observations	1,618	1,618	1,618	1,618	701	701
Panel B. DV: Product diversification experiments						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	All firms				Firms with high-end machines	
Dummy equal to one if had high-end machines in period t	0.941*** (0.209)					
Dummy equal to 1 if high-end machine expansion during period t		1.101*** (0.215)		0.835*** (0.315)	0.774*** (0.213)	0.624** (0.303)
Interaction term between high-end machines expansion and mandated output cuts dummies during period t				0.055 (0.437)		0.257 (0.431)
Dummy equal to 1 if low-end machine expansion during period t		0.221 (0.227)		0.130 (0.231)	0.079 (0.280)	0.019 (0.277)
Dummy =1 if university-educated engineer employed at $t-1$			0.846*** (0.182)	0.747*** (0.177)		0.681*** (0.210)
Dummy =1 if merchant a member of board at $t-1$			0.456** (0.180)	0.388** (0.162)		0.438 (0.267)
Firm age	-0.024 (0.017)	-0.017 (0.017)	-0.040** (0.018)	-0.040** (0.017)	-0.024 (0.020)	-0.047** (0.019)
Constant	-1.150*** (0.319)	-1.006*** (0.319)	-1.440*** (0.330)	-1.497*** (0.305)	-0.256 (0.383)	-1.014** (0.490)
Semi-annual time dummies	Included	Included	Included	Included	Included	Included
Observations	1,618	1,618	1,618	1,618	701	701

Poisson regression with the number of upgrade experiments started in period t as the dependent variable. Robust standard errors clustered at the firm level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Note: mandated output cuts measure does not vary within periods and is therefore absorbed by the semi-annual time dummies.

Table 6. Firm Growth and Complementarity between Product Innovation and Diversification:
Panel Estimation

	DV: Ln(output) at $t+1$, minus Ln(output) at t	
	(1)	(2)
Cumulative number of upgrade experiments at t	0.021 (0.013)	-0.001 (0.011)
Cumulative number of upgrade experiments x fraction of low-end products at t		0.040** (0.019)
Fraction of low-end products in total number of products at t	-0.018 (0.066)	-0.104 (0.080)
Dummy = 1 if university-educated engineer at t	0.096** (0.042)	0.101** (0.042)
Dummy = 1 if merchant board member at t	0.017 (0.025)	0.017 (0.026)
Logged installed high-end spindles in $t+1$, minus logged installed high-end spindles in t	0.013* (0.007)	0.013* (0.007)
Logged installed low-end spindles in $t+1$, minus logged installed low-end spindles in t	0.008 (0.014)	0.008 (0.014)
Ln(output) at t	-0.308*** (0.048)	-0.312*** (0.048)
Constant	2.497*** (0.356)	2.605*** (0.367)
Semiannual time dummies	Included	Included
Firm FE	Included	Included
Observations	1,608	1,608
R-squared	0.325	0.326
Number of firms	99	99

Fixed-effect panel estimations. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Note: because high-end and low-end capacity can take values of zero, we have applied the inverse hyperbolic sine transformation, $z = \log(y + \sqrt{1 + y^2})$, where y is the actual number of spindles to obtain “Logged installed high-end spindles” and “Logged installed low-end spindles” in the table above. We apply the same transformation in Table 7 and Table 8B below.

Table 7. Firm Growth and Complementarity between Product Innovation and Diversification: IV Estimation

VARIABLES	DV: number of upgrade experiments started at t		DV: Ln(output) at $t+1$, minus Ln(output) at t	
	First stage	“Placebo test”	Second stage	
	(1)	(2)	(3)	(4)
Cumulative number of upgrade experiments			0.012 (0.010)	-0.014 (0.013)
Fraction of low-end products in total number of products	-2.365*** (0.463)	-2.355*** (0.461)	-0.028 (0.035)	-0.079** (0.039)
Cumulative number of upgrade experiments x fraction of low-end products				0.080*** (0.023)
Fraction of output cuts enforced at t x Logged installed high-end spindles in $t+1$, minus Logged installed high-end spindles in t	2.542*** (0.645)			
Fraction of output cuts enforced at t x Logged installed low-end spindles in $t+1$, minus Logged installed low-end spindles in t		-0.089 (0.568)		
Logged installed high-end spindles in $t+1$, minus Logged installed high-end spindles in t	-0.132 (0.095)	0.039 (0.091)	0.019*** (0.007)	0.019** (0.007)
Logged installed low-end spindles in $t+1$, minus Logged installed low-end spindles in t	0.134* (0.069)	0.116 (0.081)	0.018 (0.017)	0.016 (0.017)
Dummy = 1 if university-educated engineer at t	-0.681 (0.524)	-0.627 (0.513)	0.064*** (0.021)	0.071*** (0.022)
Dummy = 1 if merchant board member at t	0.970** (0.383)	1.020*** (0.384)	0.025 (0.017)	0.022 (0.017)
Logged total firm output at t	0.577*** (0.180)	0.519*** (0.185)	-0.030*** (0.008)	-0.046*** (0.009)
Firm age	0.051 (0.032)	0.053 (0.033)	-0.006*** (0.002)	-0.009*** (0.002)
Constant	-4.899*** (1.530)	-4.555*** (1.543)	0.507*** (0.081)	0.682*** (0.096)
Semiannual time dummies	Included	Included	Included	Included
Observations	1,608	1,608	1,608	1,608
Log pseudolikelihood (Adj. R-squared)	-273.1	-277.3	0.170	0.177
Estimation	Poisson	Poisson	IV	IV

First stage: Poisson regression with robust standard errors clustered at the firm level. Second stage: OLS with robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Cumulative number of upgrade experiments is an instrumented variable in the IV estimations. Note: mandated output cuts measure does not vary within periods and is therefore absorbed by the semi-annual time dummies.

Table 8A. Product Upgrade Experiments and Production System Flexibility

VARIABLES	DV: Portfolio rebalancing	
	Within-count	Across-count
Cumulative number of upgrade experiments at $t-1$	0.244*** (0.082)	0.156* (0.086)
Cumulative number of diversification experiments at $t-1$	0.021 (0.016)	0.011 (0.024)
Dummy = 1 if university-educated engineer at t	0.186 (0.152)	-0.093 (0.292)
Dummy = 1 if merchant board member at t	0.086 (0.098)	0.156 (0.171)
Dummy equal to 1 if high-end machine expansion during period t	0.142 (0.155)	0.208* (0.122)
Dummy equal to 1 if low-end machine expansion during period t	0.192 (0.139)	-0.004 (0.101)
Total output (thousands of tons, adjusted to 20-count) during period t	-0.034 (0.035)	-0.138*** (0.038)
Constant	0.448* (0.253)	1.130*** (0.313)
Observations	1,605	1,605
Semiannual time dummies and firm dummies	Included	Included
Number of firms	99	99
R-squared	0.085	0.067

Table 8B. Portfolio Rebalancing and Growth

VARIABLES	DV: Ln(output) at $t+1$ minus Ln(output) at t	
Change in the number of within-count portfolio rebalancing from t to $t+1$	0.010** (0.004)	
Change in the number of across-count portfolio rebalancing from t to $t+1$		0.012*** (0.004)
Dummy = 1 if university-educated engineer at t	0.103** (0.041)	0.103** (0.041)
Dummy = 1 if merchant board member at t	0.011 (0.026)	0.008 (0.025)
Logged installed high-end spindles in t , minus Logged installed high-end spindles in $t-1$	0.014* (0.007)	0.014* (0.007)
Logged installed low-end spindles in t , minus Logged installed low-end spindles in $t-1$	-0.002 (0.015)	-0.002 (0.015)
Logged total output at t	-0.298*** (0.048)	-0.296*** (0.047)
Constant	2.411*** (0.340)	2.402*** (0.335)
Observations	1,608	1,608
Semiannual time dummies and firm dummies	Included	Included
Number of firms	99	99
R-squared	0.324	0.326

Fixed-effect panel estimations. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9A. 20-Count Demand Estimation, First Stage

Estimation	DV: Logged 20-count price	
	Instrumental regression	Placebo test
Number of across-count portfolio rebalancing (between 17-48 counts)	-0.002** (0.001)	
Number of across-count portfolio rebalancing (between 17-48 counts), interacted with mandatory output cuts	0.032*** (0.007)	
Number of within-count portfolio rebalancing (between 17-48 counts)		-0.001 (0.001)
Number of within-count portfolio rebalancing (between 17-48 counts), interacted with mandatory output cuts measure		0.005 (0.006)
Constant	4.897*** (0.027)	4.895*** (0.028)
Semiannual time dummies and firm dummies	Included	Included
Observations	743	743
R-squared	0.984	0.983

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Note: mandated output cuts measure does not vary within periods and is therefore absorbed by the semi-annual time dummies.

Table 9B. 20-Count Demand Estimation, Second Stage

	DV: Logged market share of 20 count
Instrumented logged 20-count price	-5.407 (6.136)
Constant	16.932 (30.047)
Semiannual time dummies and firm dummies	Included
Observations	743

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Instruments: number of across-count portfolio rebalancing and number of across-count portfolio rebalancing interacted with mandatory output cuts measure as explained in the main text.

Table 9C. Product Upgrade Experiments and Quality of 20 Count Products

VARIABLES	DV: Khandelwal (2010)-style measure of 20-count quality			
	(1)	(2)	(3)	(4)
Estimation	OLS		IV	
Cumulative number of upgrade experiments at t	0.067*** (0.026)	-0.037 (0.081)		
Cumulative number of upgrade experiments x fraction of low-end products at t		0.158 (0.114)		
Cumulative number of upgrade experiments at t (instrumented)			0.503*** (0.062)	-0.102 (0.217)
Cumulative number of upgrade experiments x fraction of low-end products at t (instrumented)				0.926*** (0.320)
Fraction of low-end products in total number of products at t	0.883*** (0.289)	0.801*** (0.294)	1.531*** (0.289)	1.027*** (0.364)
Dummy = 1 if university-educated engineer at t	0.990*** (0.118)	1.002*** (0.118)	0.764*** (0.116)	0.763*** (0.117)
Dummy = 1 if merchant board member at t	0.189* (0.112)	0.168 (0.112)	0.027 (0.107)	-0.056 (0.110)
Logged installed high-end spindles in $t+1$, minus Logged installed high-end spindles in t	0.055 (0.036)	0.051 (0.037)	0.050 (0.039)	0.046 (0.038)
Logged installed low-end spindles in $t+1$, minus Logged installed low-end spindles in t	0.296 (0.312)	0.288 (0.310)	0.206 (0.314)	0.203 (0.309)
Firm age	-0.018* (0.010)	-0.018* (0.010)	-0.052*** (0.011)	-0.060*** (0.012)
Constant	3.069** (1.413)	3.175** (1.415)	2.631* (1.400)	3.200** (1.420)
Semiannual time dummies	Included	Included	Included	Included
Observations	721	721	721	721
R-squared	0.492	0.493	0.522	0.527

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. “Khandelwal-style” quality measure is calculated at the firm-observation level, as the sum of firm and time fixed effects and the residuals from the regression of logged market share of 20 count on instrumented logged 20-count price, as detailed in the main text.

Table 10. Firm Survival

Numbers of:		Surviving firms	Exiting firms; of which:		Total
			By acquisition	Shut down	
Had high-end machines	Yes	19	22	1	42
	No	14	31	18	63
Total		33	53	19	105

Appendix to “Product Innovation, Product Diversification, and Firm Growth: Evidence from Japan’s Early Industrialization (will be posted online once the paper is published)

Serguey Braguinsky, Atsushi Ohyama, Tetsuji Okazaki, Chad Syverson²⁸
January 2020

A.1 Data construction

A.1.1. Product varieties data

Our data source for firm-level output by product varieties comes from bulletins (“Geppo”) published by the All-Japan Cotton Spinners’ Association (“Boren,” using its name’s abbreviation in Japanese). These bulletins were issued monthly, starting from 1889 (even earlier data are available from government statistics, starting from 1883), and they, in particular, contained firm-level input and output data for all members of the Association which was basically the universe of all firms operating in the industry at any time. The first time the breakdown of output by product varieties was added to the bulletins was May 1893. Since then, the data have been published continuously in dedicated tables in each issue. We have coded all such tables until December 1914 which is the end of our sample. Photo 1 presents a table showing output by product varieties by each firm for December 1906.

Photo 1. Output by product varieties table, December 1906.

Source: Geppo, No. 173, Jan. 25, 1907.

A.1.2. Engineers data

The information about university-educated engineers was obtained from restricted-use alumni lists, *Gakushikai Kaiinroku* compiled by Gakushikai (The University Graduates’ Society), the association of the alumni of Imperial Universities, containing information about addresses and workplaces of the graduates. Until 1897 Tokyo Imperial University was the only one. In 1897 Kyoto Imperial University was founded and its first cohort graduated in 1901. Two more Imperial Universities were founded in 1907

²⁸ Braguinsky: University of Maryland, Robert H. Smith School of Business, Osaka University Institute of Economic and Social Research, and NBER, e-mail: sbraguinsky@rhsmith.umd.edu; Ohyama: Institute of Innovation Research, Hitotsubashi University, e-mail: ohyama@iir.hit-u.ac.jp; Okazaki: Tokyo University, Faculty of Economics, e-mail: okazaki@e.u-tokyo.ac.jp; Syverson: University of Chicago Booth School of Business and NBER, e-mail: chad.syverson@chicagobooth.edu.

and 1911 but there were no graduates of the last one available to the industry at the end of our sample (1914) as yet.

The above information was verified and supplemented, especially for earlier years, from chapters dedicated to the history of each firm in Kinugawa (1964) and from published company histories (Kanebo, 1988; Unitika, 1989; Toyobo, 1986; Fujibo, 1998, Shikibo, 1968, Kurabo, 1953). Engineers educated in British universities, in particular, were identified from these industry history sources and added to the list of graduates of Japanese Imperial Universities.

For technical college graduates, we used annual *Ichiran (Catalogs) (Tokyo Koto Kogyo Gakko Ichiran, Kyoto Koto Kogyo Gakko Ichiran, Osaka Koto Kogyo Gakko Ichiran, Nagoya Koto Kogyo Gakko Ichiran, Kumamoto Koto Kogyo Gakko Ichiran, and Sendai Koto Kogyo Gakko Ichiran)*, which contain the lists of alumni with their current workplaces, and picked up all graduates of mechanical engineering and dyeing departments who worked in one of the firms in our sample in any given year. The first technical college was established in Tokyo in 1881, the second one in Osaka, in 1896. By the end of our sample there was the total of six technical colleges that already had alumni working in the industry; all those alumni data were coded and added to the database of educated engineers employed by cotton spinning firms.

A.1.3. Board members and merchants data

About 90 percent of firms in our sample (and all significant firms) were public (joint stock) companies, obligated to issue shareholders' reports every half a year (see Braguinsky et al., 2015, and Agarwal et al., 2020, for details). We have photocopied and processed 1,443 reports on 106 firms (Kokajo, 1883-1914), all such reports that we could find surviving until the present day.²⁹ Each report, in particular, contains a list of all shareholders and board members ("torishimariyaku") of the company issuing it. For privately held firms as well as in cases where some shareholders reports of incorporated firms were missing in the archival data, we supplemented this with information from the All-Japan Registry of Firms Executives ("Yakuinroku"), the first issue of which was published in 1893. We coded and reconciled the information on the members of the boards across these two sources, and also cross-checked the data by using information in company histories (Kinugawa, 1964; Kanebo, 1988; Unitika, 1989; Toyobo, 1986; Fujibo, 1998, Shikibo, 1968, Kurabo, 1953), and in Geppo.

The data published in "Yakuinroku" (above) were used to extract the names and addresses of board members of the four incorporated cotton yarn-related trade companies (Naigai Wata, Nihon Menka, Nitto Menshi and Mitsui Bussan), while the data contained in the 1943 brief history of Osaka Three Article Exchange ("Sampin Shoshi") were used to extract the names of board members and traders registered at this most important exchange dealing with cotton and cotton yarn in each year. The data in four available editions of *Nihon Zenkoku Shoko Jinmeiroku*, a nationwide registry of traders and manufacturers (for 1892, 1898, 1907, and 1914), were utilized to extract the names and addresses of individual merchants likely to play the most prominent role in cotton spinners' output markets; that is, traders in cotton yarn and cotton yarn-woven garments who paid business taxes exceeding a certain threshold (10-15 yen, depending on the year). We considered all these individuals to be potential providers of market knowledge ("market ties") for cotton spinning firms; hence, we matched their names and addresses to the names and addresses of cotton spinning firms' executives. Through this process (which involved both a computer algorithm and a manual re-check), we identified 55 executives of incorporated cotton yarn-related trade companies and board members and registered traders at the Osaka Three Articles Exchange, as well as 172 significant individual traders who were also board members in cotton spinning firms and created panel data reflecting their presence as an executive board member in firm i at time t . The variable "merchant as a board member" used in the main text takes value of one if at least one such merchant was identified to be a member of the executive board of firm i at time t and zero otherwise.

²⁹ We are grateful to Osaka University library and Prof. Takeshi Abe, its academic director at the time; to Yoshiyuki Murakami, head of the company history division of Toyo Boseki at the time (now retired); and to Kobe University Kanematsu Collection for their cooperation in allowing us to photocopy company reports used in this research.

A.1.4. Machine capacity data

The data on machine capacity (the number of spindles installed by firm i at time t) were compiled and validated by cross-checking from three sources. The first source is the already mentioned semiannual shareholders' reports (Kokajo, 1883-1914), which often contained inventory of all property owned by the firm, including details about machines (number of spindles, their basic type—ring, mule or doubling—and the name of the manufacturer, such as Platt Brothers of Oldham, etc.). In some cases, however, those detailed inventories were not included, or the reports themselves may be missing. “Enkakukiji” (1901) contains the details of all installation and decommissioning of machines by cotton spinning firms from the inception of the industry and until 1901, while “Sankosho” (1903-1914), a semiannual bulletin published by Boren parallel to Geppo, contains the details about machines (number of spindles and their basic type, although not the name of the manufacturer) starting from 1903. Combining and manually cross-checking the data from these three sources, we were able to construct the full panel with the number of spindles installed by each firm at each point in time during our sample.

A.1.5. Machine orders from British manufacturers

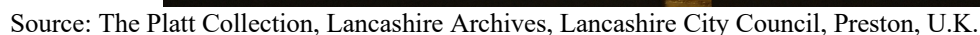
We collected and processed archival data on machine orders by Japanese cotton spinning firms placed with British textile machine manufacturers and preserved in Lancashire archives from the earliest such orders and until 1914. As mentioned in the main text, the same data collected by the late Gary Saxonhouse are archived on the ICPSR website (Wright, 2011). We had to collect and re-process these data once again because there were no original photos and firm names were missing from the ICPSR files, making it impossible to perform firm-level matching to Japanese sources of capacity data described above.³⁰ The total number of orders we collected and processed was 430. In Photos 2 and 3 below we present two representative orders placed with Platt Brothers of Oldham around the same time (Spring-Summer 1896) by two different Japanese firms. As can be seen, machines were custom-made, and each order contains the number of frames ordered, the number of spindles per frame (so that the total number of new spindles being ordered can be calculated), as well as detailed technical characteristics, additional hand-written notes taken by British engineers presumably during consultations with the client Japanese firms, the dates the order was taken and when machines were booked (shipped), in multiple installments.

The technical characteristics of machines specified in the orders allow us to differentiate between types of products they were designed to produce, types of inputs required, and other technological nuances. For example, the order in Photo 2 lists “Numbers to be Spun Av” as 42s. This is a “high-end” product in our classification and, hence, the machines that came with this order are high-end machines. We also know from the product varieties data (as well as from industry and firm history) that at the time, Amagasaki was beginning to scale its output of 42 count doubled yarn first introduced on experimental basis only a few years earlier. We can also see from the same photo that cotton input is described as Good American and that hank roving is listed as 6 ½ double. Thus, we can indeed see that high-end machines designed for vertically upgraded products required dedicated inputs and technologies (compare to Photo 3 discussed immediately below). In the particular case of this order, we also have a description of how it actually happened in the company history (Unitika, 1989, p. 13). Amagasaki's chief engineer, Kyozo Kikuchi (one of a few educated engineers employed by Japanese cotton spinning firms at the time), personally went to England to discuss the details of the order with the British engineers, traveling before that to the U.S. to examine the technology for producing 42 count doubled yarn, including comparing between Good and Middling cotton—and reaching the conclusion about using Good cotton.

Photo 3 presents another typical order, by Fukushima Spinning, for 4,992 spindles (13 frames with 384 spindles in each frame). In this order, “Numbers to be Spun Av” are 10's-20's, with average 14's. In line with this, cotton input is listed as Indian, Japanese, and Chinese while hank roving is 2 5/8 single. We therefore classify the machines in this order as “low-end machines” since they were not

³⁰ Our newly collected and processed data, including photos of the originals, will shortly become publicly available from NBER Industry, Productivity, and Digitalization Data Library (<http://www.nber.org/data>).

Photo 2. High-end machines order placed by Amagasaki Spinning in July 1896

[illegible]

50

A.1.6. Hand-matching machine capacity panel data with technical characteristics of machines

There were three major sources of changes to machine capacity owned by Japanese firms over the sample that are identifiable in the data. The first such source was the installation of new machines ordered from Britain. To illustrate how we matched orders to capacity changes in the Japanese firms that placed them, we work off Amagasaki Spinning 1897 capacity change.

Amagasaki Spinning's spindle capacity was reported by Enkakukiji (1901) to be 27,036 spindles in the first half of 1897 (verified from Unitika, 1989; matched also with previous orders). The company shareholders report No. 15 (for the first half of 1898) indicates that the capacity had increased by 18,176 spindles to the total of 45,212 spindles by the end of the second half of 1897. Where did these new machines come from? The order presented in Photo 2 provides us with the bulk of the answer: in 1896 the firm placed an order with Platt Brothers of Oldham, for 14,544 "high-end" spindles (36 frames, 404 spindles in each frame). Another order (not shown), for 2,828 spindles, also for high-end machines (numbers to be spun 30s-42s, average 36s) was placed with Platt Brothers even earlier, in December 1894 and booked the next year, but was apparently installed together with the 14,544 spindles above, as part of the same expansion (verified also through company history). Out of the remaining 804 new spindles, 404 are accounted for by an order (not shown) placed in 1896 with another British manufacturer, Dobson & Barrow, also to spin 42 count yarn.³¹ This leaves 400 spindles (2.2 percent) of the total capacity increase in late 1897 "unaccounted for," in the sense that we do not have a record of the corresponding order. From the firm inventory of property (not shown) we know that Amagasaki Spinning also added a 400-spindle frame manufactured by Brooks & Doxey at about the same time. Even though the actual order record is missing, due to its small size and timing, we can safely assume that this was another sample trial order, which in all probability should be similar to orders from Platt Brothers and Dobson & Barrow above. Hence, in this case, we have definitively matched 97.8 percent of the increase in installed capacity as reflected in property inventories to the corresponding orders, and we feel very confident that the remaining 2.2 percent of new capacity were of a similar design.

Applying the above procedure, we were able to hand-match all 430 orders whose records we found in Lancashire archives to corresponding changes in firms' machine capacity using the timing,³² number of spindles, and the information about which of possibly multiple mills operated by the firm the order was slotted for (e.g., in Photo 2 machines are slotted for Amagasaki Mill No. 3—and company reports tell us that No. 3 was the mill constructed at the time to expand the output of the 42 doubled count yarn—the designated number of counts to be spun noted in the order is consistent with that).

The second source of changes to machine capacity owned by Japanese firms over the sample was acquisitions. Fifty eight firm-by-firm acquisitions were consummated during the period of our sample in the industry, involving 69 different mills (plants) (see Braguinsky et al., 2015; the number of acquisitions in there is listed as 73, involving 95 plants but 15 acquisitions involving 26 plants happened in 1915-20 which is outside of our current sample). As a result of these acquisitions, two-thirds of total industry capacity changed hands over the period, hence, it was imperative to trace machines from their original owners to subsequent (at times multiple) owners. Fortunately, all the acquisitions are well documented in company histories and company reports; the latter also reflect new machines acquired as a result of firm-by-firm acquisitions in the property inventories. Hence, we were able to create a script that reassigned machines (and technical characteristics thereof) through (possibly multiple consecutive) acquisitions to

³¹ A hand-written annotation (apparently by a Dobson & Barrow employee) reads that "This is a sample trial order so let everything be finished with the utmost care, as further large order depends on the same being satisfactory." Amagasaki Spinning never placed another order with Dobson & Barrow, preferring to continue working with Platt Brothers. The note, however, gives a glimpse of the competition among British textile machinery makers.

³² As can be seen from the Amagasaki example above, one-two or even more years would normally elapse between the placing of the order and the time the machines would be fully installed and reflected in firms' property inventories. (Newly arrived machines were commonly included in property inventories after they were installed and ready to be operated; prior to that, new arrivals would be reflected in balance sheets in the "expansion account" but not in property inventories, which are the primary source of our panel data on firm capacity.)

new owners, making sure that we updated the breakdown of machine capacity of acquiring firms each time they expanded their capacity through acquisitions.³³

The third source of capacity changes were the removal of aging machines and their destruction in accidents such as fires or earthquakes. Fortunately, fires and other destruction were rare events, and those that did happen are documented in company histories, including the details of which machines were lost. As for machine removal, those were also uncommon (Japanese firms used their machines while conducting necessary repairs along the way for many decades). Inasmuch as the machines were decommissioned, such events concentrated in larger firms and in later years of our sample, where inventories of properties in company reports are especially detailed, so we had no problem identifying the machines that were being removed and updating the capacity breakdown of the remaining machines accordingly.

While we were thus able to match all orders available in Lancashire archives to firm capacity changes, and then follow those machines through acquisitions and possible destruction and removal as above, the converse is not always true. That is, not all firm capacity changes identified in our panel data using Japanese archival sources could be matched to orders placed with the British manufacturers. There are, broadly speaking, two reasons for that. One reason is that originals of the orders may be missing from Lancashire archives (such as the small order Amagasaki Spinning placed with Brooks & Doxey mentioned above). Rather than drop all such observations, we imputed machine characteristics (high-end or low-end) to new machine arrivals reflected in firms' property inventories whenever we had unambiguous evidence allowing us to do that (once again, the case of 400 spindles of Brooks & Doxey machines added to Amagasaki Spinning capacity around 1897 mentioned above serves as an example). Such "imputed orders" comprise about 10 percent of the total orders we matched to firms' property inventories, both in terms of the number of cases and the total number of spindles (Table A1).

Table A1. Machine orders: actual and imputed

	Number of orders	Number of spindles ordered
From Lancashire archives' originals	430	2,439,414
Imputed	50	301,106
Fraction imputed	0.10	0.11

The second reason is second-hand market transactions where firms bought machines decommissioned by other firms. We do have several documented instances like that in company histories, and those are accounted for in our capacity breakdown. However, we do not have systematic records of purchases in the second-hand market, and in cases where we were also missing company reports, we had to drop observations where we could not know the origins of the machines, from the analysis that required machine characteristics. This affected 11 percent of all firms but only seven percent of all observations. Moreover, firms with unknown machine origins were generally small, so their fraction in the total number of capacity-weighted observations in our data is just about two percent (Table A2). Thus, dropping them from the analysis affects it only marginally.

Table A2. Total spindle-observations matched to underlying orders

	# of firms	# of observations	Capacity (# of spindles): Total	Of which: matched	Fraction matched
Matched sample	105	1,911	61,719,337	60,824,202	0.99
Unmatched sample	13	154	1,020,912	0	0.00
Fraction matched	0.89	0.93	0.98	1.00	0.97

³³ In addition to firm-by-firm acquisitions, there was one plant-by-firm acquisition, in which a firm bought just one of the plants from another firm. We know exactly which machines were involved in this acquisition too.

As Table A2 shows, for firms for which we do have the data on how their machine capacity evolved through orders, acquisitions and removal/destruction (the “matched sample”), the correspondence between machine capacity changes recorded in Japanese archival sources (based on property inventories) and our calculated changes combining the data sources above is nearly perfect. Specifically, we were able to match 60.8 million spindle-observations out of 61.7 million total in the matched sample of 105 firms, the match rate of more than 99 percent. This brings the fraction of total industry capacity (including also firms for which matching turned out to be impossible) that we could match to the characteristics of the machines comprising it to 97 percent (Table A2).

A.2 Historical trends

Figure A1. Dynamics of the number of firms and industry output



Table A3. Imports, production, and exports, 1880-1900 (units: *kori*=0.18 ton)

Year	Imports from:		Total	Dom. production	Export	Dom. supply
	England	India				
1880	84,367	10,908	95,324			95,324
1881	73,515	17,932	92,421			92,421
1882	61,708	21,794	84,324			84,324
1883	55,687	26,448	82,135	2,326		84,461
1884	48,790	21,797	70,622	5,687		76,309
1885	40,437	30,887	71,324	3,370		74,694
1886	45,251	36,850	82,101	16,217		98,318
1887	54,103	56,884	111,096	25,273		136,369
1888	77,583	80,546	158,281	33,142		191,423
1889	62,194	80,488	142,929	69,959		212,888
1890	59,703	46,566	106,588	108,374	31	214,931
1891	42,624	15,560	58,123	160,207	108	218,222
1892	53,494	27,527	81,534	213,489	109	294,914
1893	48,426	16,216	65,174	222,223	1,053	286,344
1894	45,353	7,778	53,555	304,584	11,796	346,343
1895	44,157	4,472	49,876	383,565	11,776	421,665
1896	63,859	2,854	67,373	428,864	41,916	454,321
1897	52,380	355	54,555	544,461	140,116	458,900
1898	52,697	353	54,563	670,067	229,445	495,185
1899	27,101	252	28,339	785,612	341,202	472,749
1900	30,035	101	30,941	647,484	208,732	469,693

Source: Takamura (1971), Vol. 1, pp. 146, 183; Association data (dom. production in 1883-85).

Table A4. Imports and domestic production by variety 1891-93 (units: *kori*=0.18 ton)

Imports from:	Variety (count)	1891	1892	1893
India	10's-18's	2,449	4,507	1,164
	20's	17,570	17,657	12,376
	Subtotal	20,019	22,164	13,540
England	16's-24's	13,079	13,130	9,385
	28's-32's	13,727	14,967	8,737
	38's-42's	1,431	2,345	1,759
	32's doubled	2,506	10,827	12,045
	42's doubled	5,882		
	60's-120's gassed	5,377	7,809	8,532
	Subtotal	42,003	48,444	40,458
Total		62,022	70,608	53,998
<i>(Total import, stat.)</i>		58,123	81,534	65,174
Domestic production	<14's		43,677	43,052
	14's-15's		38,064	39,579
	16's		36,347	44,125
	17's-19's		8,918	10,617
	20's		49,919	68,604
	21's-24's		1,390	3,894
	25's-27's		104	31
	28's-32's		4,329	8,121
	33's-37's		113	131
	38's-42's		51	0.2
	Other		8,377	6,730
	Total	160,207	189,206	225,091

Source: Takamura (1971), Vol. 1, p. 184. Per source, import data come from port statistics in Kobe and Tokyo, therefore, does not correspond exactly to customs statistics data in Table 1 but captures most of it.

Table A5. Imports and domestic production by variety, 1898 (units: *kori*=0.18 ton)

Variety	Estimated import		Dom. production	Export	Domestic supply	
	min	max			min	max
<16's			160,355	3,246	157,109	157,109
16's-24's	4,248	4,248	453,936	198,246	259,938	259,938
16's			224,982	120,039	104,943	104,943
20's			208,479	77,905	130,574	130,574
28's-32's	3,186	3,186	22,032	45	25,173	25,173
38's-42's	531	531	2,115		2,646	2,646
60's gassed	6,372	7,434	19		6,391	7,453
80's gassed	9,558	10,620	3,044		12,602	13,664
100's gassed	2,655	2,655			2,655	2,655
120's gassed	1,593	1,593			1,593	1,593
doubled <32's			517		517	517
32's doubled	7,965	9,558	1,779		9,744	11,337
42's doubled	14,337	15,930	7,022		21,359	22,952
Total	50,445	55,755	650,819	201,537	499,727	505,037

Source: Takamura (1971), Vol. 1, p. 322.

A.3 Industry Association-mandated output cuts

Table A6 presents the details of Boren-mandated output cuts and (in the right column) our measure of their relative impact on low-end products. For output cuts imposed as mandated holidays, we converted those into percentage terms by dividing the number of holidays by 30 (e.g., 4/30=0.133). The numbers in the column “Implied fraction of lower counts mandated cuts” is the difference between the numbers in previous two columns.

Table A6. Mandatory output cuts during our sample period

		Output of 20 count and lower	Output of 21 count and higher	Implied relative fraction of lower counts mandated cuts
year	month	Policy detail		
1899	1	4 mandated holidays	4 mandated holidays	0
1899	2-12	None	None	0
1900	1-4	None	None	0
1900	5-7	4 mandated holidays	None	0.133
1900	8-12	40% of spindles to be idled or suspension of night shift	None	0.400
1901	1-3	40% of spindles to be idled or suspension of night shift	None	0.400
1901	4-12	None	None	0
1902	1-6	None	None	0
1902	7-12	4 mandated holidays	4 mandated holidays	0
1903	1-12	None	None	0
1904	1-12	None	None	0
1905	1-12	None	None	0
1906	1-12	None	None	0
1907	1-12	None	None	0
1908	1-4	5 mandated holidays	None	0.167
1908	5-12	27.5% of spindles to be idled or suspension of night shift	None	0.275
1909	1-12	27.5% of spindles to be idled or suspension of night shift	None	0.275
1910	1-4	20% of spindles to be idled or suspension of night shift	None	0.200
1910	5-9	None	None	0
1910	10-12	27.5% of spindles to be idled; or 12.5% of spindles to be idled, plus 4 days and 2 hours per workday mandated holidays	20% of spindles to be idled or 5 days and 2 hours per workday mandated holidays	0.075
1911	1-9	27.5% of spindles to be idled; or 12.5% of spindles to be idled, plus 4 days and 2 hours per workday mandated holidays	20% of spindles to be idled or 5 days and 2 hours per workday mandated holidays	0.075
1911	10-12	27.5% of spindles to be idled; or 12.5% of spindles to be idled, plus 4 days and 2 hours per workday mandated holidays	10% of spindles to be idled or 5 days mandated holidays	0.175
1912	1-3	27.5% of spindles to be idled; or 12.5% of spindles to be idled, plus 4 days and 2 hours per workday mandated holidays	10% of spindles to be idled or 5 days mandated holidays	0.175
1912	6-9	4 mandated holidays	4 mandated holidays	0
1912	10-12	None	None	0
1913	1-12	None	None	0
1914	1-7	None	None	0

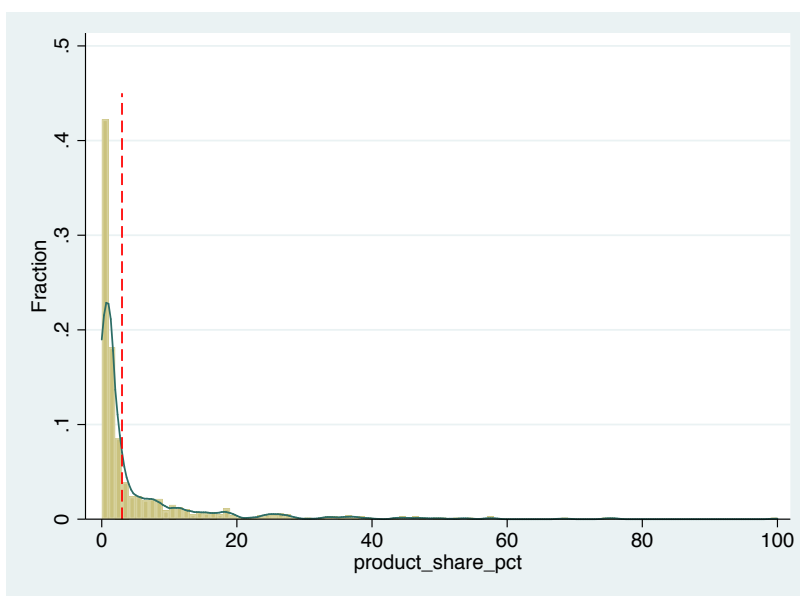
Source: compiled from Shoji, Otokichi, 1930. *Boseki Sogyo Tanshuku Shi* (History of Operational Curtailments in Cotton Spinning, in Japanese). Nihon Mengyo Kurabu, Osaka, Japan.

Since in the main text we use semi-annual data in our IV growth regression estimations, we converted the numbers from Table A6 into semi-annual data by multiplying the implied relative fraction of lower counts mandated output cuts by the number of months during which the restrictions were applied in any given semi-annual period, divided by 6. For example, 4 mandated holidays in the first half of 1900 were imposed for two months (May and June), hence our measure of mandatory output cuts impact on lower counts for the first half of 1900 is $0.133 \times 2/6 = 0.089$. For the second half of the same year it is $0.133 \times 1/6$ (July output cuts) + $0.4 \times 5/6$ (output cuts in the rest of the year) = 0.356, and so on.

A.4 Initial scale of new product introduction and definitions of product experimentation

Figure A2 presents the distribution of within-firm shares of newly introduced products. The mean and median within-firm shares are 5.27 percent and 1.39 percent, respectively. The majority of new products in our data are thus introduced as a small share of firm's total output, and new product introduction below three percent threshold are quite common. However, this initial small scale of new product introduction is much more pronounced among firms with high-end machines than for firms without such machines (the mean and median within-firm shares of newly introduced products for the first time for firms with high-end machines are 3.49 percent and 0.9 percent respectively, as opposed to 5.74 percent and 1.58 percent, respectively, for firms with no high-end machines).

Figure A2. Distribution of within-firm share of newly introduced products



mean	s.d.	10th	25th	50th	75th	90th	N
5.27	10.43	0.08	0.33	1.39	4.91	15.21	793

While in the main text we have used the three-percent cutoff of the within-firm share of newly introduced products as our empirical definition of an experimental product, we conducted sensitivity analysis and confirmed that the findings are robust to using reasonable alternative thresholds (such as two percent or four percent). It is also possible to define experiments in alternative ways. For example, we may consider a newly introduced product to be experimental if it was discontinued (at least for a certain period of time) after being introduced. Since firms may not know which product varieties are going to be continued and which are going to be discontinued at the time of their introduction, defining experimental product varieties in this way gives rise to a “selection on outcome” concern. We did try such alternative definitions anyway, and confirmed that our basic results still continue to hold. Details are available upon request.

A.5 Concentration of output, high-end machines, and product varieties

Table A7. Concentration of output, high-end machines, high-end and low-end products

	High-end machine capacity			
Output quintile	1893	1899	1907	1913
1	-	486	256	144
2	-	2,816	1,799	866
3	-	2,322	5,313	12,469
4	5,445	3,087	4,351	47,190
5	2,966	16,314	67,100	122,159
	Number of high-end product varieties			
Output quintile	1893	1899	1907	1913
1	0.0	0.5	0.4	0.3
2	0.1	1.1	0.8	0.3
3	0.3	0.6	0.9	1.9
4	1.1	0.1	1.2	3.0
5	1.1	1.3	3.3	6.5
	Number of low-end product varieties			
Output quintile	1893	1899	1907	1913
1	2.3	1.8	1.8	2.3
2	2.7	3.4	2.9	2.7
3	4.1	4.1	2.9	3.7
4	4.6	3.8	5.4	5.0
5	6.2	5.6	5.4	9.2

Source: Our calculations using the data described in the main text and in the appendix.

Table A7 presents high-end machine capacity (measured by the number of high-end machines spindles installed), the number of high-end product varieties and the number of low-end product varieties by quintiles of total output (firm size measured by output scale, as in the right panel of Figure 3) at four points in time during our sample. The top panel shows that as high-end machine capacity increases from virtually zero at the start of our sample, it spreads out across firms of different size, but remains heavily concentrated among the largest firms (even though the composition of the largest firms changes due to higher growth rates of firms with high-end machines, as we document below). Not surprisingly perhaps, we can see the same pattern in the number of high-end product varieties produced in the middle panel. What is most interesting, however, is that the number of low-end product varieties (which do not require high-end machines for their production) also becomes heavily concentrated over time, especially in the last period (which corresponds to the second subperiod in Table 2 in the main text when there was a big increase in within-firm number of product varieties). When the top decile firms (in terms of output) are compared with the second decile firms in 1913 (not shown), both categories are similar in the number of high-end product varieties and high-end machine capacity, but the number of low-end product varieties in the former category is twice as many as that in the latter category. This difference corresponds to the difference in total output, indicating that the expansion of low-end product varieties is also critical for firm growth.

Table A8 presents the average numbers of all product varieties and of low-end product varieties among those, produced by firms that did and did not have high-end machines at a given point in time. We present these statistics for the whole sample as well as for two subperiods in Table 2 in the main text. Firms with high-end machines produced on average almost twice as many product varieties as firms that did not have such machines over the whole period, with the gap larger in the second period than in the first (all the differences in means in Table A8 are statistically highly significant). Even more interestingly, firms that had high-end machines produced more *low-end* product varieties compared to firms that only had low-end machines, and the gap, once again, is larger in the second subperiod than it is in the first. When limited to firms that had high-end capacity above the mean of 36,000 spindles (14 out of 42 firms that had high-end machines), the gap with firms with no high-end machines is even much larger.

Table A8. Number of products and low-end products by firms with and without high-end machines

		Firms with high-end machines			Of which: firms above the mean # of spindles in high-end machines	Firms with no high-end machines		
		1893.2-1914.2	1893.2- 1907.1-	1907.1-	1893.2-1914.2	1893.2- 1907.1-	1907.1-	1907.1-
Number of products	Mean	7.75	6.77	9.52	11.01	3.98	3.89	4.29
	St. Err.	0.19	0.18	0.42	0.46	0.07	0.07	0.17
	# obs	725	465	260	205	1031	810	221
Number of low-end products	Mean	4.94	4.59	5.57	6.00	3.77	3.71	3.99
	St. Err.	0.13	0.14	0.27	0.34	0.06	0.06	0.16
	# obs	725	465	260	205	1031	810	221

Source: Our calculations using the data described in the main text and in the appendix. The difference in means between firms with and without high-end machines are highly statistically significant using double-sided *t*-test.

A.6 Selection into high-end machine adoption

We have taken the separation of firms into those that adopt high-end machines with the aim of conducting product innovation and subsequent product diversification and firms that stick to scaling their existing products (or do not grow at all) as given. Here, we use our data to try to open the “black box” of this separation.

We first conduct a simple cross-sectional comparison of founding team characteristics between firms that installed high-end machines at some point during our sample (including those that entered with high-end machines right away) and firms that never had high-end machines. Since university educated engineers played a crucial role at all stages of utilizing high-end machines and firm growth, we test whether firms that managed to secure the services of such engineers at the time of their founding (incorporation) were more likely to have high-end machines. We also test whether market ties, proxied by the presence of merchants among members of founding teams, played an independent role in high-end machine adoption. The results of estimating the effects of these two founding team characteristics are presented in columns (1) and (3) of Table A9. In column (1) the dependent variable is the firm-specific dummy equal to one if the firm had high-end machines at some point during our sample and zero otherwise, while in column (3), where the sample is limited to post-1893 entrants, it is the dummy equal to one if the firm entered with high-end machines and zero otherwise (pre-1893 entrants basically did not have a chance to install high-end machines at the time of entry).

The estimation results show that firms which installed high-end machines were indeed much more likely to employ a university educated engineer already at the time they were founded (note that no firm employed more than one at the time they were founded). The presence of merchants on the founding teams is also positively associated with high-end machines, but the magnitude of the coefficient is much smaller and it is statistically not significant at conventional levels. Historical evidence documented in Kinugawa (1964), company histories and other sources shows that it was both (i) forward-looking firms realizing that they needed to secure high-level engineering human capital if they wanted to install high-end machines and try to adopt new, unfamiliar production technologies required to operate at the high end of the product space; and (ii) hired educated engineers playing a key role in working with the British suppliers to order, customize, and install high-end machines. Thus, while we cannot claim “causality” whereby having university educated engineers at founding led to adoption of high-end machines, we do see strong selection (positive assortative matching) on both sides (founding teams and engineers), associated with the adoption of high-end machines.

To probe for this assortative matching further, we add firm size (measured by the number of spindles installed) at the time of entry as another regressor. Initial entry size has been shown to be associated with the entrepreneurial ability both theoretically (Lucas, 1978) and empirically (Klepper, 2010), so we use it as a proxy for entrepreneurial ability of the founding team in our data. Estimation

results presented in columns (2) and (4) of Table A9 show that higher entry size in our data is indeed positively and statistically highly significantly associated with the adoption of high-end machines. However, while whatever impact the presence of merchants had on high-end machines adoption is almost completely absorbed by entry size, the coefficients on university educated engineers employed at founding, although smaller in magnitude, remain economically and statistically significant. Indeed, comparing the magnitudes of the coefficients on engineers and entry size, we can see that having a university-educated engineer has an estimated association with high-end machines adoption which is equivalent to an increase in more than 10,000 spindles (more than one standard deviation from the mean) in entry size. Thus, hiring university-educated engineers' human capital can be seen to be a critically important independent factor in high-end machine adoption, alongside higher entry size.³⁴

Table A9. Founding Team Differences Among Firms with and without High-End Machines

DV:	Probability of ever having high-end machines		Probability of entry with high-end machines	
VARIABLES	(1)	(2)	(3)	(4)
University educated engineer employed at founding	0.333*** (0.111)	0.241** (0.109)	0.313** (0.132)	0.230* (0.125)
Merchant a member of founding team	0.071 (0.094)	0.018 (0.090)	0.064 (0.100)	0.024 (0.093)
Firm capacity (in thousands of spindles) at entry		0.019*** (0.005)		0.018*** (0.005)
Constant	0.287*** (0.069)	0.170** (0.074)	0.144* (0.073)	-0.003 (0.080)
Observations (firms)	105	105	69	69
R-squared	0.093	0.191	0.091	0.231

Linear probability estimation. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The sample in columns (3) and (4) is limited to post-1893 entrants. Probit estimations lead to very similar results.

We next utilize observations on firms which we observe while they did not yet have high-end machines they would have in the future. Among the 42 firms that had high-end machines at some point during our sample, 19 entered with high-end machines right from the beginning but 23 installed them at some point after entry. Four of these 23 added their first high-end machines prior to the start of product variety data, so we do not get a chance to observe them before that. For the remaining 19 firms we do have this opportunity, and we use this sample to look into the issue of “pre-treatment” selection by comparing those firms *before* they actually adopted high-end machines with the firms that never introduced such machines during our sample. If the “treatment” (adoption of high-end machines) was an outcome of a random draw, we should not observe systematic differences in the pre-treatment period.

In Panel A of Table A10 we present the results of the estimations, where the dependent variable is a dummy equal to one if the firm adopted high-end machines in the future and zero otherwise. We regress this on various actions and characteristics *before* they had high-end machines (if any). To err on the side of caution, we exclude not only observations after the high-end machines appear on firms' balance sheets but also all observations pertaining to the period after firms placed their first high-end machines order (if any), even though the machines are not reflected in the balance sheets yet. While this results in losing observations (including all observations on three of the 19 firms whose orders were placed early), it

³⁴ While, once again, we cannot claim causality based on our data, it is worth noting that in a recent paper, Choi et al. (2019) establish causal effect from founding team's human capital to startup growth.

makes sure that, especially when looking at experimental production, we do not inadvertently pick up experiments conducted with high-end machines that may have been partially installed.³⁵

The estimation results in column (1) in Panel A of Table A10 show that firms that ordered high-end machines in the future (hereafter, “future adopters,” for short) tended to introduce significantly more new products on experimental basis than firms that never adopted high-end machines, already prior to adoption. Strikingly, they were even more likely to conduct product upgrading experiments than diversification experiments compared to other firms during the pre-adoption stage. Although, as mentioned, producing most high-end products was impossible without high-end machines, stretching low-end machines enabled some small-scale production at the (lower end of the) high-end product space, and this is what we apparently observe in these data. Thus, one interpretation of these results is that potential entrepreneurial firms were conducting small pre-adoption experiments to determine whether they should make the costly decision of investing in high-end machines (Kerr et al., 2014).

According to the model in McCardle (1985), experiments with new technologies result in a separating equilibrium where some firms (receiving bad signals) decide not to adopt the new technology, while other firms (receiving good signals) decide to go ahead with investment. We test this in another regression, the results of estimating which are presented in Panel B of Table A10. In column (1) we calculate the fraction of successful experiments (products first introduced on experimental basis and subsequently scaled) among all experiments started by firms at time t and we use it as a key explanatory variable in a regression where the dependent variable is, once again, the dummy equal to one if the firm was a future adopter and zero otherwise. Because we are conditioning on starting at least one experiment at t , the number of observations is much smaller than in Panel A. Nevertheless, there is an economically important and statistically significant at the 10 percent level relationship between success in experimental production and future high-end machine adoption—the standard deviation from the mean in this sample is 0.44, so one standard deviation from the mean success rate is estimated to increase the probability of high-end machine adoption in the future by about 8 percentage points, or by about 40 percent, given that the mean fraction of future high-end machines adoption in the sample is 0.2. Thus, future adopters both conducted more experiments than non-adopters and were more successful in the experiments they did conduct, in line with the predictions in McCardle (1985).

In columns (2) and (3) of Panel B in Table A10 we look at how the duration of successful and failed experiments, respectively, was different between future adopters and non-adopters. As can be seen from the corresponding coefficients, in the case of successful experiments, those conducted by future adopters were significantly shorter than those conducted by non-adopters, but the opposite is true of eventually failed experiments. The first finding is consistent with future adopters having better success odds (higher exponential arrival rate of success in any given experiment). The latter finding renders itself to several possible interpretations, the most likely ones (not necessarily mutually exclusive) being (a) extra confidence future adopters had about their chances to eventually succeed which made them persist longer despite a string of bad signals, and (b) higher value placed by future adopters on experiments as a means of accumulating technological and marketing knowledge, even regardless of outcome.

In column (2) of Panel A in Table A10 we add more regressors capturing time-varying firm characteristics—employment of university educated engineers, market ties (proxied by merchants as board members), the situation in top management teams (TMT), expansion of (low-end) machine capacity, as well as output scale and scope. The coefficient on upgrading experiments is reduced in magnitude but remains economically and statistically highly significant in predicting which firms are going to order high-end machines in the future. Among other variables, market ties (as captured by the presence of merchants on boards) significantly positively affect the probability of adding high-end machines in the future in this sample, but the number of university educated engineers employed has no such impact. Thus, even though having a university-educated engineer at the time of founding is an important predictor of high-end machines at entry or at some point in the future (see Table A11 above),

³⁵ The results remain similar if we use all observations prior to high-end machines appearing on the balance sheets.

firms that did not adopt high-end machines at the time of entry did not on average employ more educated engineers than other firms either, until they decided to order such machines for the first time or even later.

Table A10. “Selection” Among Firms with no High-End Machines

Panel A. Characteristics and actions associated with future high-end machine orders

	DV: Dummy equal to 1 if placed high-end machines orders in the future	
VARIABLES	(1)	(2)
Number of upgrade experiments started at t	0.418*** (0.147)	0.350*** (0.119)
Number of diversification experiments started at t	0.066** (0.031)	0.031 (0.038)
Firm age	-0.003 (0.008)	0.006 (0.009)
Number of university-educated engineers employed at t		-0.073 (0.057)
Dummy equal to one if merchant a board member at t		0.167** (0.079)
Shared TMT leadership at t		0.003 (0.096)
Dummy equal to one if a TMT leader ousted in the previous two years		-0.169*** (0.063)
Dummy equal to one if added low-end machines at t		-0.044 (0.042)
Logged firm total output at t		0.111** (0.047)
Logged number of products at t		-0.014 (0.058)
Constant	0.211** (0.080)	-0.734** (0.361)
Semiannual time dummies	Included	Included
Observations	964	921
R-squared	0.065	0.212

Pooled OLS. Robust standard errors clustered at the firm level in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Sample limited to firms before high-end machine orders were placed for the first time.

Panel B. Success rates and duration of successful and unsuccessful experimental production

	DV: Dummy equal to 1 if placed high-end machines orders in the future		
VARIABLES	(1)	(2)	(3)
Fraction of successful experiments among those started at t	0.208* (0.117)		
Average duration of successful experiments		-0.037** (0.015)	
Average duration of failed experiments			0.043*** (0.010)
Firm age	0.008 (0.006)	0.025 (0.016)	-0.000 (0.006)
Constant	0.201***	0.266**	0.203***

	(0.070)	(0.128)	(0.064)
Semiannual time dummies	Included	Included	Included
Observations	153	98	172
R-squared	0.294	0.430	0.335

Pooled OLS. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Looking at the top management teams (TMT), Agarwal et al. (2020) find that *stable* shared leadership (that is, having multiple TMT leaders who created division of labor while avoiding rifts leading to forced departures) contributed to firm growth and ability to attract engineering talent. We test this here by including also dummies capturing whether a firm had shared TMT leadership and whether there was a forced TMT leader departure in the two years preceding time t (where “forced” is defined as not due to death, illness, or resignation for personal circumstances unrelated to the firm). While shared TMT leadership by itself does not predict future adoption of high-end machines, we do find that within-TMT conflicts leading to forced departures are strongly *negatively* associated with high-end machines adoption. Since such departures represent disruptions and loss of managerial human capital, this finding is in line with the conclusions drawn in Choi, et al. (2019) using U.S. Census data. Finally, we also see that while future adopters were somewhat larger in terms of their output scale, they were neither more diversified nor more likely to expand their (low-end) machine capacity than other firms prior to adopting high-end machines. Thus, we have another strong piece of evidence showing that diversification and expansion of low-end machines happened *after* firms adopted high-end machines and successfully expanded into the high-end product space, not before.

To conclude, it appears that (a) conducting more experiments and being more successful in those experiments; (b) stronger market ties (merchants as members of boards); and (c) ability to avoid within-TMT conflicts leading to forced departures and loss of human capital were all salient features associated with “selection into treatment” (future high-end machine adoption). At the same time, most of the impact on growth, product varieties proliferation and hiring of educated engineers (as well as further increase in experimental production) happened already after high-end machines had been adopted and the firms had embarked on the process of growth using those machines and exploiting the complementarity channel from product upgrading to product diversification.

A.7 Flexible production system

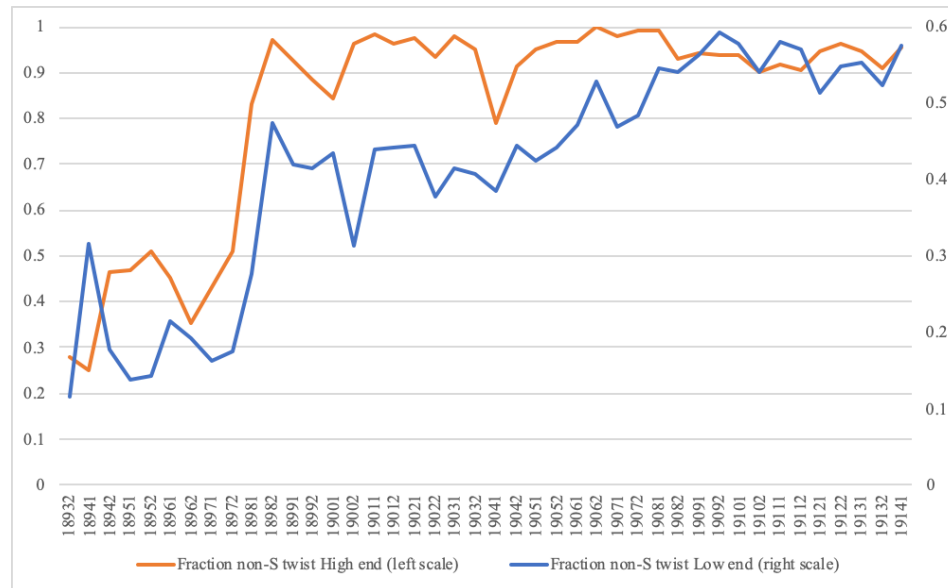
According to our interviews with Kanji Tamagawa (a prominent Japanese historian of cotton spinning technology), switching the direction of twist (from S to Z and vice versa) and/or changing the count to be spun between adjacent counts (say, from 16 to 18 or from 20 to 22) involved the following operations:

Operation	Required time
(1) Changing the draft (adjusting the draft change gear)	About 10 minutes
(2) Changing the twist number (adjusting the twist change gear)	About 10 minutes
(3) Changing the direction of the twist (changing the direction of the tin roller)	About 10 minutes
(4) Changing the spindles rotation speed (adjusting the gear)	About 10 minutes
(5) Changing the traveler (choosing a traveler of appropriate weight)	About 30 minutes

With appropriate skill, the first four operations could be completed during the time required to complete the fifth, so, under ideal conditions, the whole process would take about 30 minutes.

However, if the task was to change to a count of yarn that was further apart (such as from 16 count to 20 or from 24 count to 32 count), such a change also required a change in roving to be fed into the machine, taking about two hours, so such adjustment was much costlier. Thus, firms with a more diversified product portfolio could change counts more easily.

Figure A4. Historical trend in the fraction of non-S twist products in the total number of product varieties



A.8 Low-end product quality

Table A11. Comparing 20-count Prices and Female Operative Wages Dispersion

	Coefficient of variation	Quartile coefficient of dispersion	90-10 percentile coefficient of dispersion
Prices of 20 count	0.028	0.016	0.031
Wages of female operatives	0.147	0.094	0.182

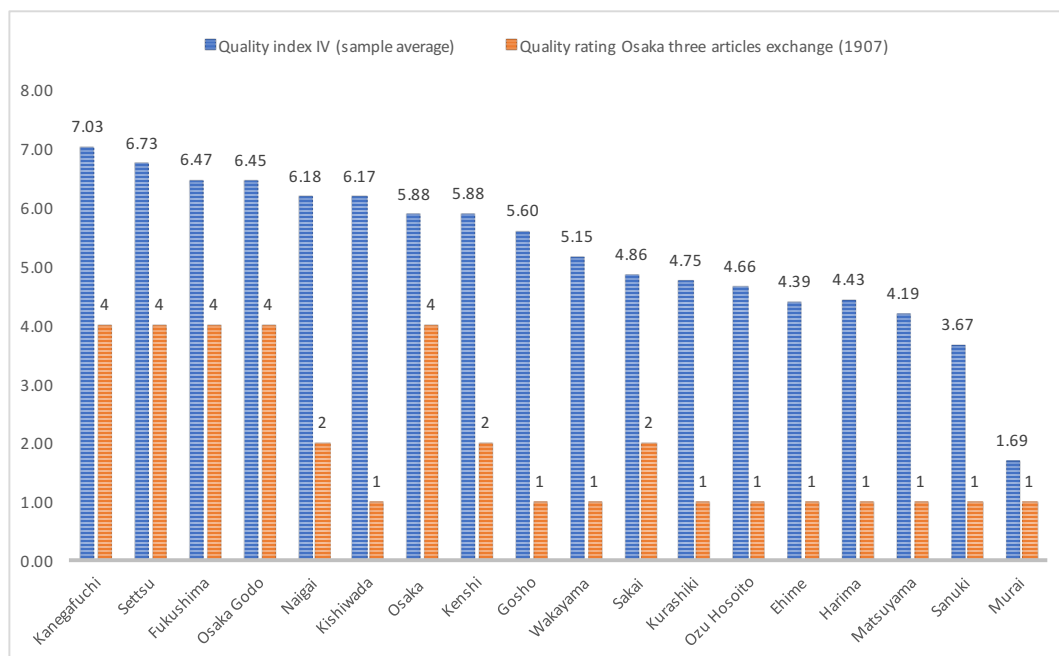
Note: The coefficient of variation is the ratio of the mean over the standard deviation. The quartile coefficient of dispersion is $(p75-p25)/(p75+p25)$. The 90-10 percentile coefficient of dispersion is $(p90-p10)/(p90+p10)$. Averages for the whole sample.

Table A12. OLS Regression of Logged Market Share of 20 Count on Logged 20-count Price

	DV: Logged market share of 20 count
Estimation	OLS
Logged 20-count price	-3.056** (1.191)
Constant	7.132 (5.432)
Semiannual time dummies and firm dummies	Included
Observations	743
R-squared	0.732

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

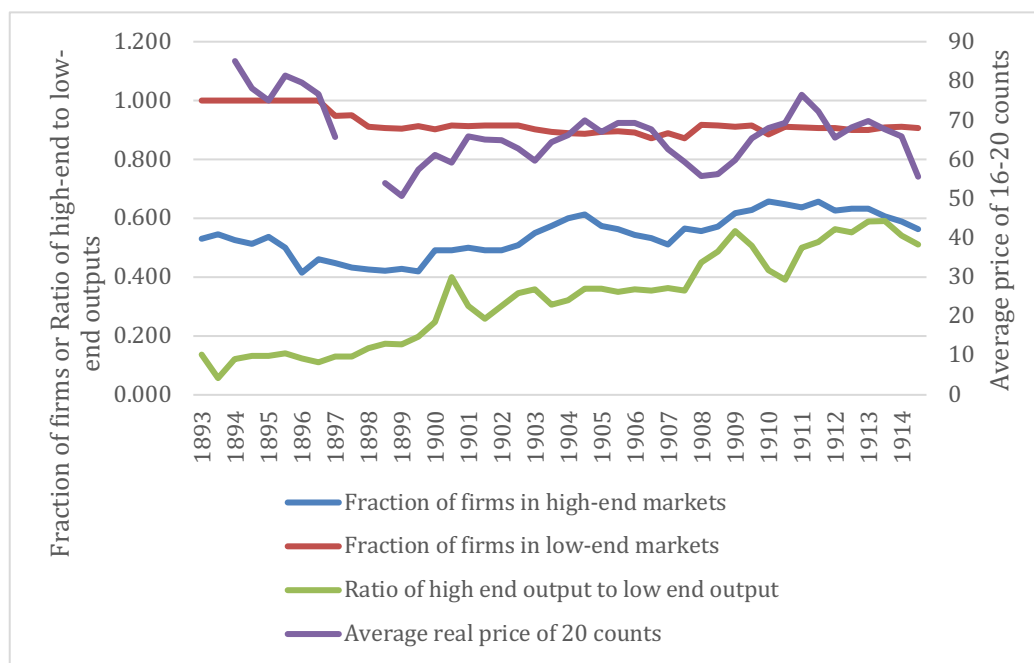
Figure A5. Estimated Quality Index and Historical Quality Ratings



Note: Quality rating is firm-level quality ratings of Z-twist 20 count in 1907 determined by the cotton yarn quality rating committee at the Osaka Three Articles Exchange (Arao, 1920, pp. 33-36). Rating 4 indicates the highest quality and rating 1 indicates the lowest quality. “Quality index IV” represents sample averages of our estimated quality measures for the same firms.

A.9 Market competition

Figure A6. Fraction of Firms, Output ratio, and Price



References cited in the appendix, not included in the references list in the main text

- Arao, Yoshimichi, 1920. *Menshi Boseki oyobi Menshi Torihiki* (Cotton spinning and cotton yarn transactions, in Japanese). Kogyo Kyoikukai Shuppanbu, Osaka.
- Agarwal, Rajshree, Serguey Braguinsky, and Atsushi Ohyama, 2020. “Centers of Gravity: The Effect of Stable Shared Leadership in Top Management Teams on Firm Growth and Industry Evolution.” *Strategic Management Journal*, 41(3): 467-498.
- Choi, Joonkyu, Nathan Godtsclag, John Haltiwanger, and J. Daniel Kim, 2019. “The Post Entry Growth and Survival of Business Startups: The Role of Founding Teams.” Unpublished.
- Fuji Boseki. 1998. *Fuji Boseki Hyakunenshi* (100 years of Fuji Boseki, in Japanese), Fuji Boseki, Tokyo.
- Gakushikai, 1886-1914. *Gakushikai Kaiinroku* (The Alumni List of University Graduates' Society, in Japanese). Tokyo (restricted use, only for members).
- Jinmeiroku, 1892, 1898, 1907. *Nihon Zenkoku Shoko Jinmeiroku* (All-Japan Registry of Traders and Craftsmen, in Japanese). Available at <http://dl.ndl.go.jp/search/searchResult?featureCode=all&searchWord=日本全国商工人名録&viewRestricted=0>.
- Kanebo. 1988. *Kanebo Hakunen-shi* [100 Years of Kanebo, in Japanese]. Kanebo Kabushiki Gaisha, Osaka.
- Klepper, Steven, 2010. “The Origin and Growth of Industry Clusters: The Making of Silicon Valley and Detroit.” *Journal of Urban Economics*, 67: 15-32.
- Kokajo (Company reports) (1883-1914), Osaka University Library: Osaka, Japan, in Japanese.
- Kumamoto Koto Kogyo Gakko, 1901-1914. *Kumamoto Koto Kogyo Gakko Ichiran* (Catalog of Kumamoto Technical college, in Japanese).
- Kurabo, 1953. *Kaiko Rokujugonen* (Looking Back at 65 years, in Japanese). Osaka, Kurashiki Boseki Kabushiki Gaisha.
- Kyoto Koto Kogei Gakko, 1905-1914. *Kyoto Koto Kogei Gakko Ichiran* (Catalog of Kyoto High School of Technology and Craftsmanship, in Japanese).
- Lucas, Robert E. Jr., 1978. “On the Size Distribution of Firms.” *Bell Journal of Economics*
- McCardle, Kevin F., 1985. “Information Acquisition and the Adoption of New Technology.” *Management Science*, 31(11): 1372-1389.
- Nagoya Koto Kogyo Gakko, 1908-1914. *Nagoya Koto Kogyo Gakkou Ichiran* (Catalog of Hagoya Technical college, in Japanese).
- Okamoto, Yukio. *Chiho Boseki Kigyo-no Seiritsu-to Tenkai* (Formation and Development of Cotton Spinning Firms in the Province, in Japanese), 1993. Kyushu Daigaku Shuppankai, Fukuoka.
- Osaka Koto Kogyo Gakko, 1900-1914. *Osaka Koto Kogyo Gakko Ichiran* (Catalog of Osaka Technical college, in Japanese).
- Osaka Sampin Torihikijo Seisan Jimusho, 1943. *Sampin Shoshi* (Brief History of Osaka Three Articles Exchange, in Japanese), Osaka, Japan.
- Sendai Koto Kogyo Gakko, 1910-1914. *Sendai Koto Kogyo Gakko Ichiran* (Catalog of Sendai Technical college, in Japanese).
- Shikibo, 1968. *Shikishima Boseki Nanaju-gonenshi* (Seventy five years of Shikishima Spinning, in Japanese), Osaka, Shikishima Boseki Kabushiki Gaisha.
- Tokyo Koto Kogyo Gakko, 1886-1914. *Tokyo Koto Kogyo Gakko Ichiran* (Catalog of Tokyo Technical college, in Japanese).
- Toyo Bōseki. 1986. *Toyoboseki Hyakunenshi* (One Hundred Years of Toyo Boseki). Toyobo (in Japanese): Osaka, Japan.
- Unitika. 1989. *Unitika Hyakunenshi* (100 years of Unitika, in Japanese). Available at <http://www.unitika.co.jp/company/archive/history/>
- Yakuinroku, 1893-1914. *Zenkoku Shogaisha Yakuinroku* (All-Japan Registry of Firms Executives, in Japanese). Available at <http://dl.ndl.go.jp/search/searchResult?featureCode=all&searchWord=日本全国諸会社役員録&viewRestricted=0>.