

WORKING PAPER · NO. 2020-03

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

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

JANUARY 2020

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

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

January 2020

JEL No. D2,L1,N6,N8,O3

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 set involves removing the "supply-side constraint" by investing in new types of machines and technologies. This process involves a high degree of uncertainty, however, so firms that take steps in this direction tend to first introduce innovative products on experimental basis. We show that conducting such experiments is a key to firm growth. It not only provides opportunities to capture the market in high-end vertically differentiated products when successful, but also facilitates horizontal differentiation of the firm's products within its previous technical capabilities. In long-term outcomes over 20 years, 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.

Serguey Braguinsky
University of Maryland
Robert H. Smith School of Business and
Department of Economics
4558 Van Munching Hall
College Park, MD 20742
and NBER
sbraguinsky@rsmith.umd.edu

Atsushi Ohyama
Hitotsubashi University
Institute of Innovation Research
Naka 2-1
Kunitachi, Tokyo 186-8603
Japan
ohyama@iir.hit-u.ac.jp

Tetsuji Okazaki
Faculty of Economics
The University of Tokyo
7-3-1 Hongo, Bunkyo-ku
Tokyo 113-0033, JAPAN
okazaki@e.u-tokyo.ac.jp

Chad Syverson
University of Chicago
Booth School of Business
5807 S. Woodlawn Ave.
Chicago, IL 60637
and NBER
chad.syverson@chicagobooth.edu

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 of approaches emphasizes supply-side innovation, where development of new, innovative products delivers the 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 fully summarized by the number of (symmetrically related) products the firm makes.

However, all product additions are unlikely to be the same in reality. Adding a highly innovative product may affect growth differently than adding a product that is more similar to those the firm has already been producing. First, the demand and cost effects may be different, given that customers' willingness to substitute and firms' economies of scope are likely to depend on the degree of 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 effect on long-term firm outcomes. The 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 that are 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 by conducting entrepreneurial experimentation (Kerr et al., 2014; Cusolito and Maloney, 2018, Ch. 4) and 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 in incremental steps 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), our findings imply that high-growth firms have more product flexibility and improve the quality of their products to better respond to demand changes. While product/firm appeal is typically treated as a primitive in those studies, our analysis suggests that product/firm appeal arises endogenously through experimenting with innovative products. Therefore 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 the 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 range of product varieties and infrequent product changeovers. Experimenting firms, on the other hand, changed their product portfolios frequently and expanded their overall scope. Our findings indicate the role played by vertical product upgrading experimentation in generating 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.

The rest of the paper is organized as follows. In the next section we briefly describe our data and historical context (more details are provided in the appendix). Section 3 documents basic trends in industry and firm growth through product varieties expansion. Sections 4 and 5 examines the mechanisms through which vertical product upgrading led to product expansion through diversification and translated into firm growth. Section 6 probes channels through which product innovation may have contributed to subsequent diversification as well as presents some robustness checks, while Section 7 concludes and contains further discussion of our findings.

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 reports, 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 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

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.

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 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 and continuously applied throughout the whole 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 distinguish between product varieties that belonged to the high-end tail of the product variety space and those at the low-end tail. Counts above 20 are high-end product varieties, while counts 20 and below are low-end. Among our 35 total product varieties, 16 are high-end according to this classification and 19 low-end. We chose this

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.

threshold because counts from the higher 20s and up 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 from high-end varieties 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.

To investigate firms' product expansion patterns, we also 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

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 of a variety in semi-annual period t that is both novel and produced at a modest quantity. More precisely, for a

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.

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. Note that this definition allows for multiple experiments in the same product line. If an experimental product was temporarily discontinued before it reached the threshold and is re-introduced again below the threshold after at least one period of non-production, this is considered 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 1,167 firm-product varieties ever produced by a firm in our sample, 620 (53.1%) were produced on a scale above three percent of their firms' contemporaneous total output from the outset. The remaining 547 product lines were initially “experimental,” and 324 of those never reached the threshold scale.⁶ Hence, almost 60 percent of product lines introduced on experimental basis completely “failed,” that is, never scaled above the threshold despite possibly multiple experiments. If we restrict the observations to only new product lines that were introduced at some point after the start of our observation period (those not in a firm's production set at the time of entry or at the time of first observation), 62 percent of product lines introduced first on experimental basis failed to ever reach the scaling threshold.

In Panel B we look at all experimental product launches, counting experiments in the same product line as separate episodes if the product was dropped, possibly temporarily, without being scaled. The fraction of experimental products that did not scale is even higher, at 73 percent. 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.

Table 1 also shows that 91 out of 547 all experimental product lines (16.6%) and 76 out of 439 new experimental product lines (11.1%) represented upgrades according to our definition above. Recall that an experiment is an upgrade if it involves a high-end product that is of a count higher than any count the firm had produced above the three percent threshold before. Of these upgrade experimental product lines, 45 percent (51 percent for those introduced after the start of our observation period) never reached the threshold. For individual upgrade experiments, this fraction is even higher at almost 64 percent.

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 Lancashire archives in Preston, U.K. contains books with orders that cotton spinning firms worldwide placed with British textile machine manufacturers (including but not limited to Platt Brothers). We collected all the available data on Japanese orders placed from the inception of the industry in the early 1880s and until 1914.⁷

All machines were custom made. Each order observation 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.).

⁷ These orders had been previously examined and processed by Gary Saxonhouse and archived at the ICSPP after he passed away (Wright, 2011). However, with no originals, it turned out to be impossible to match ICSPP 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.

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 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 of spindles 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 only seven percent of all observations and 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 the technical characteristics of machines to the product varieties produced by firms that owned those machines.

As mentioned, high-end products and low-end products required different types of machines. We employ the matched orders data to distinguish between machines designed predominantly for high-end products and those for low-end products. More precisely, machines are defined to be high-end if their design 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 in our sample for which we have breakdown of machine

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.

capacity, 42 had at least one high-end machine at some point in time during our sample period. Among these, however, the capacity distribution was highly skewed. The mean number of spindles in high-end machines was 36,763, while the median was 16,128.

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 cuts was not under the control of any given firm.⁹ Because high-end products were with one partial exception exempt from these mandated cuts, such low-end output controls gave firms extra incentives 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 in different periods (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 variable over a given semi-annual observation period, in the third column in Table A6 in the appendix, represents our measure of the degree of relative output cuts imposed on low-end products.¹¹ In subsequent analysis, 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.

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, high counts were also affected in 1910-12. Because we are interested in the *relative* stringency of output cuts imposed on low-end products compared to high-end products, we subtracted the fraction of high-end output cuts during that period from the fraction of low-end output cuts when constructing our measure.

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 start of our data until about 1899, but after that starts increasing first in high-end product space and then in low-end products as well. 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, with the expansion in the total number of products driving almost all industry growth. 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 is accompanied by an increase in the number of low-end product varieties too. 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

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

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.

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,” capturing 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 their numbers of products between t and $t+1$, holding their base period market share fixed. 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 the average number of products of entrants is greater (less) at the time of entry than in later periods of their operation. A positive (negative) number in the fifth term means that the average number of products of exiting firms is greater (less) at the time right before they exit than in earlier periods of their 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 summarized as follows. First, the decomposition in Panel A, where all the product categories in the data are used, shows that the total increase in the average number of products per firm between 1893 and 1914 was 8.98. Of this, the within-effect accounts for a greater share (5.63) than the composition effect (3.35). However, when we divide the sample into two subperiods before and after the spike in product varieties in 1907, there are stark differences. From 1893-1906, the overall increase of 2.53 products per firm is more modest, and the composition effect (1.63) contributes more to the total change in the number of product varieties than the within effect (0.90). 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 (1.73), the within effect (4.73) becomes a dominant contributor to the total growth in average products per firm of 6.45. Also, 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), on the other hand, while new entry still contributed to the growth in the number of product varieties at about the same magnitude as before, there was an explosion in the number of product varieties produced by continuing firms. This growth came to dominate the overall expansion of the number of product varieties produced by the industry. Decomposition separately by high- and low-end products in Panels B and C of Table 2 presents essentially the same picture, although 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? To the extent that they were the same firms, did they first expand low-end product varieties and then move on to high-end product varieties, as implied by the learning-by-doing theory of Stokey (1988), for example? Or, did they first introduce high-end product varieties and then expand also their low-end product offerings?

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 out 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-end and low-end product varieties that we saw in the decomposition analysis. Moreover, those were also the firms that invested in high-end machines and grew to become the top firms in the industry (see also Table A7 in the appendix). This leads to obvious questions about the relationships between product upgrading and diversification and how these

might translate into firm growth. This is what we explore next.

3.2 Upgrade experiments and product varieties expansion

As seen above, almost two-thirds of all new products introduced by firms were initially launched on a small scale (experimental production), and 60-70 percent of those experiments did not result in scaling. 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:

$$y_{it} = \alpha + \beta_1 \text{cuml_upgrade_exp}_{it-1} + \beta_2 \text{cuml_divers_exp}_{it-1} + \beta_3 X_{it-1} + \gamma_i + \delta_t + \zeta_\tau + \varepsilon_{it}, \quad (1)$$

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; $\text{cuml_upgrade_exp}_{it-1}$ and $\text{cuml_divers_exp}_{it-1}$, 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 low- and high-end specifically). Columns (1)-(3) include the number of products (all,

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

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 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

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.

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 not statistically related 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 that of the relation with high-end machine expansion.

It is also worth noting that the relationships between high-end machine expansion and product growth holds even conditioning on past upgrade experiments. As we will see immediately below (Tables 5 and 6), 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 presents a puzzle. Why is it that firms which grow by expanding their product varieties are those that first introduce new innovative and technically challenging products on an experimental basis? As mentioned, most upgrading experiments fail to scale up and to stay in a firm's portfolio. Thus the direct effect of upgrade experiments on the expansion of product varieties is limited, and it has no direct effect whatsoever on the expansion of low-end product varieties. We seek to address this puzzle in the analyses 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 having such machines was indeed associated with more product upgrade experiments. The decision to introduce high-end machines is a costly investment decision, so firms select into doing so. We take installations as given for now, but later we employ an instrument based on industry-wide mandatory output cuts of low-end products as a source of exogenous variation in incentives to conduct upgrade experiments.

At the beginning of our sample, Japanese mills were trying to catch up to the world technology frontier. Delivering high-end products was a difficult task entailing much uncertainty. In addition to procuring high-end machines, clearing this hurdle required 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.

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 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 reflect that firms that had installed 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 high-end machines. 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 much higher 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.) Once again, because high-end machines were not directly relevant to low-end products, this is 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

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.

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.

--- Table 4 around here ---

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 (again as a proxy for market ties), 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. In column (2) we look at the roles of 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 despite the much smaller number of observations is still statistically significant.

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. The presence of a merchant on the board of executives led on average to an additional 1.5

upgrade experiments per semi-annual period. Thus our proxy for having market ties is related to twice the increment to experimentation levels as is having a university-educated engineer, and it has a magnitude similar to high-end machine capacity expansions. This regression uses across-firm variation, so these results are not directly comparable to the firm-fixed-effects estimations in Table 3 above. However, it is worth noting the contrast between the high correlation between having market ties and upgrade experimentation and its lack of a significant relationship with product varieties expansion.¹⁶ In other words, merchants on executive boards accompanied experimentation, but they played little explanatory role in subsequent product varieties expansion and, as we will see below, overall firm growth. This suggests that market ties were important in helping firms cope with initial market uncertainty when introducing unfamiliar products on experimental basis, but subsequent expansion was less dependent on such ties and instead relied more on the experience and technological knowledge embodied in engineers.¹⁷

--- 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 by interacting an indicator for contemporaneous installment of new high-end machines with an indicator for the imposition of mandatory output cuts for low-end products, as described above in Section 2.4. One might be concerned that high-end installations are endogenous to 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 unlikely that firms would—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. Furthermore, Boren imposed mandatory low-end product output

¹⁶ 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.

¹⁷ As discussed below and in appendix A.11, the presence of merchants on executive boards was also closely related to the experiments leading to the selection of firms into purchasing high-end machines in the first place.

cuts in the wake of big, unexpected demand shocks like the Boxer rebellion of 1900 or the stock market crash of 1907. 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.¹⁸ The interaction of the mandated output cuts with high-end machine installations gives us an instrument for upgrade experimentation that varies both intertemporally and across firms.

In columns (4) and (6) of Table 5, we include this interaction term. As can be seen, the instrument is significantly related to firms starting more upgrade experiments. The effect is especially large in the subsample of firms that already have high-end machines. Among those firms, the 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 term capturing the arrival of new high-end machines amidst mandatory output cut periods (which of course affects firms' incentives to start product upgrading but not product diversification experiments), the factors that explain the start of diversification experiments are remarkably similar to those for product upgrade experiments in Panel A. In particular, 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 results in Table 3, which showed low-end product expansion was related to the same factors (cumulative past upgrade experiments and high-end machine expansion) as high-end product expansion.

One notable difference between the estimation results presented in Panels A and B of Table 5 is a more prominent role played by engineers compared to merchants in explaining diversification experiments as opposed to upgrade experiments. This is again consistent with the interpretation of market ties playing a bigger role in the initial push to expand the technology frontier, followed by accumulation of technological knowledge playing a key role in firm

¹⁸ High-end machine expansion also happened through acquisitions. There too it took at least a year, often more, to consummate an acquisition.

expansion through both high- and low-end product additions. As we will see in the next section, the interaction between accumulated past upgrade experimentation with the expansion of the number of low-end product varieties is a key determinant of firm output scale growth.

In appendix A.6 we examine how firms select into procuring high-end machines. A key finding, echoing what we already saw after the investment decision had already been made, is that among the 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 merchants as board members than firms that never installed high-end machines (see Table A10, Panel A, in Appendix A.6). Future adopters were somewhat larger in terms of output (measured using our 20-count equivalent physical units), but they were neither more diversified nor more likely to expand their low-end machine capacity prior to adopting high-end machines. Thus once again we see that diversification and expansion of low-end machines happened *after* firms successfully expanded into the high-end product space, not before. More selection analysis and discussion is in Appendix A.6.

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 with vertical product upgrading, followed by application of the knowledge gained to diversifying product offerings at the low end of the product spectrum. In this section, we examine if these factors also translated into higher growth of total output. Specifically, we look at how the cumulative number of past upgrade experiments and their interaction with the fraction of low-end products in the firm’s product portfolio affect the firm’s subsequent output growth. The basic estimating equation is:

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

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 . The estimation is conducted using panel data estimation with firm

fixed effects γ_i so as to capture the impact of firm-specific knowledge accumulation through upgrade experiments and the transfer of that knowledge into expansion of its low-end product portfolio. The parameters of interest are β_1 and especially that on the interaction term, β_3 . A positive β_3 would be consistent with a complementarity between vertical 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 becomes economically and statistically indistinguishable from zero, the coefficient on the interaction is positive and significant, at 0.04. 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, this implies that 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) from the mean is associated with a 1.2 percentage point higher growth rate at the mean number of cumulative upgrade experiments. The link between upgrade experiments and growth does indeed appear 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 ---

Recall that some new high-end products (specifically, relatively easy-to-produce mid-20s-count yarns) could be introduced, at least on a trial basis, before firms had fully installed high-end machines. In particular, as shown in appendix A.6 (Table A.10), some firms apparently used such upgrade experiments to determine whether to make costly investments in high-end machines and university-educated engineers in the first place. This suggests that while they were useful for firms to decide on their investment strategies, their usefulness for learning the

technology and hence post-investment growth could be more limited. To test this conjecture, in columns (3) and (4) we split the cumulative number of upgrade experiments conducted by a firm into those conducted with and without high-end machines.

The estimation results show that the cumulative number of upgrade experiments conducted with high-end machines had a larger and more precisely estimated association with firm growth than the cumulative number of all upgrade experiments. The estimates in column (3) imply that each upgrade experiment conducted with high-end machines installed was related to a 3.2 percentage point increase in the firm's growth rate. Firms that already had high-end machines saw average output growth of 6.2 percent, so each upgrade experiment added just over half of the mean growth rate in this sample. The interacted specification in column (4) also implies a much larger relationship with firm growth than its counterpart using all upgrade experiments. The mean cumulative number of upgrade experiments among firms with high-end machines is 2.21, and the average fraction of low-end products in this sample is 0.67. Therefore the coefficient on the interaction term of 0.087 implies that one additional upgrade experiment, conditional on the mean fraction of low-end products, was associated with a 5.8 percentage point higher output growth. This is slightly over 30 percent of the interquartile difference in growth rates across firms. Similarly, a one-standard-deviation increase in the fraction of low-end products was associated with a 6.3 percent higher total output growth rate at the mean number of cumulative upgrade experiments.

The decision to conduct product upgrading experiments may be related to other factors correlated with firm growth. To have a better sense of the causal connection between upgrade experimentation and firm growth, we again use mandatory output cuts as a source of exogenous variation in upgrade experimentation. Here, we employ the continuous measure of the extent of relative output cuts imposed on low-end products at time t , described in Section 2.4 above and presented in Table A6 in the appendix. We interact this with the (also continuous) growth rate of new high-end machine capacity during period t to construct an instrument for upgrade experiments that 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, we estimate our IV specification in two stages. First, similar to the specification in 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 output restrictions and high-end machine installation as well as other exogenous variables (firm age, logged firm output, etc.). Next, we construct the firm's implied cumulative predicted upgrade experiments by summing these period-specific predicted values. We use this constructed predicted upgrade experiments as an 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 fraction of mandated output cuts in period t , the growth rate in the number of high-end machine capacity in period t , and their interaction. The coefficient on the interaction term is positive and significant, consistent with firms that expanded high-end machines being pushed toward conducting upgrade experiments especially during times of output cuts. As a placebo test, we conducted a similar regression using the growth rate in low-end capacity instead. See column (2) in Table 4. There is no significant relationship between upgrade experiments and the interaction in this case.

--- Table 7 around here ---

We then use the constructed number of upgrade experiments to estimate an equation similar to (2). 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 what we found in column (1) in Table 6. Column (4) uses the interaction of constructed cumulative upgrade experiments and the fraction of low-end products. The coefficient on this interaction is twice as large as the corresponding coefficient in Table 6 and is estimated much more precisely. In fact, despite very different estimation methodologies, the coefficient on instrumented variables (and the implied magnitudes of growth effects) from the second-stage regressions in column (4) of Table 7 is very similar to the coefficient in column (4) of Table 6, where we looked specifically at upgrade experiments conducted by firms that already had high-end machines.

6. Transmission channels

How exactly does technical knowledge acquired through experimentation with upgrade products lead to the expansion of low-end product varieties? 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.¹⁹ S-twist and Z-twist yarn also differ in 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 within 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.

¹⁹ 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.

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, including high-end ones.)

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 did more 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 0.22, while the mean number of within-count portfolio rebalancing events is 0.47. Hence an additional upgrade experiment was associated with an increase in portfolio rebalancing within counts by about 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 0.15, implying that an additional upgrade experiment is associated with an increase in portfolio rebalancing across counts by about 15 percent of its mean (equal to 1.0).

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.²⁰

--- Table 8A around here ---

²⁰ 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, remained 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 product 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 also at the low end (cf. Roberts, 2004, pp. 37-38).

Table 8B confirms that these transmission channels of upgrade experiments contributed to firm growth. A rise in portfolio rebalancing from t to $t+1$, regardless of whether it was within- or across-count, was economically and statistically significantly associated with accelerations in firm growth. One extra within- (across-) count portfolio rebalancing is associated with 1.1 (1.3) 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 2.22 (2.12). Hence, a moving from no within-count (across-count) portfolio rebalancing to its conditional mean is associated with a 2.4 (2.8) 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 experiments in Table 6 and those on the number of portfolio rebalancing events in Table 8B remain almost the same. The two relationships therefore operate independently of each other.²¹

--- 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 of upgrading experimentation affects the demand appeal of its low-end products. We utilize the method of estimating quality for horizontally differentiated products proposed by Khandelwal (2010), 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.²² The 20-count yarn is also just at the borderline between low-end and

²¹ 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 (of either type) is associated with about one percentage point increase in capacity utilization rates. Details are available upon request.

²² The 20 count accounted for 27.6 percent of industry output over the sample, but its importance goes far 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 toward the end of the 1880s, and in 1890, Osaka Spinning Company exported the first experimental batch of 20-count yarn to China. This marked the start of the transition to a 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. Being one of the only two yarn counts listed on the Osaka Three Articles Exchange (alongside 16 count), it totally dominated the trade volume (over 95

high-end products, which allows us to instrument for its price using a plausible cost shifter that would be difficult to obtain in our data for other product types.

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 output cuts imposed on low-end products. As noted, rebalancing the portfolio entailed adjustment costs, so its frequency can be used as a proxy for such costs. Lower adjustment costs allowed firms to substitute more easily from 20-count output to counts above 20; hence, such firms would feel less pressure to reduce the 20-count price during periods of slow demand than would firms facing higher adjustment costs.²³ This gives us a supply-side source of price variation that is plausibly uncorrelated with quality (demand appeal).

Before we proceed further, we note that while most firms (and all important firms) had 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.), indicating some degree of horizontal differentiation, the within-count substitution across brands appears to be very high, as indicated by a very small degree of variation in prices.²⁴ This makes it hard to obtain precise estimation of the effect of price on demand but we tried anyway.

Thus, to instrument for the price of 20-count yarn, we construct an across-count portfolio rebalancing measure for counts from 17 to 48 (that is, excluding very low and very high counts) and interact it with the time-varying degree of mandatory output cuts. The results of this first-stage estimation (which also includes semiannual time fixed effects and firm fixed effects which are included in the second stage) are presented in the first column of Table 9A. Lower portfolio adjustment costs (more frequent portfolio rebalancing by a firm) are strongly positively associated with the firm's 20-count price during mandatory output cuts periods, even controlling for both firm and time fixed effects. Again, because portfolio rebalancing distribution is highly skewed, it is illuminating to compute the effect of going from doing no portfolio rebalancing to its mean conditional on being positive and mandated output cuts being in place. This conditional

million yen gross trade volume in 1914 as opposed to just over 1,320 yen for the 16 count), and was traded every month, as opposed to only January and December for the 16 count (Tokei Nempo, 1915, pp. 166-169).

²³ 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.

²⁴ As can be seen from Table A11 in appendix A.8, the coefficient of variation of 20-count price across the sample was just 2.8 percent, while 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.

mean is 0.34, so going from doing no portfolio rebalancing to its conditional mean is associated with about one percent higher price, more than half of the interquartile dispersion during periods of output cuts (which at 1.8 percent is slightly higher than 1.6 percent in the whole sample).

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 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.

In the second stage, we regress the firm’s (logged) share of industry-wide output of 20-count yarn on the instrumented logged price as well as the measure of across-count portfolio rebalancing and time and firm dummies. The results, in Table 9B, indicate that the own-price elasticity is negative and large in magnitude, about -7.5. 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 (which is itself consistent with a high willingness of buyers to switch yarn suppliers). 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, and the same other controls. This “naïve” regression produces an own-price elasticity of -3, 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.²⁵

We then use this demand appeal measure as a dependent variable to 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

²⁵ In our setting, Khandelwal’s “outside good” is an amalgam of cotton yarns of all other counts.

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:

$$y_{it} = \alpha + \beta_1 \text{cuml_upgrade_exp}_{it} + \beta_2 \text{low_frac}_{it} + \beta_3 \text{cuml_upgrade_exp}_{it} \times \text{low_frac}_{it} + \beta_4 X_{it} + \Delta_t + \varepsilon_{it}, \quad (3)$$

where y_{it} is the quality (demand appeal) measure of firm i 's 20 count at time t constructed as above; $\text{cuml_upgrade_exp}_{it}$ 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 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 product upgrading (the complementarity between experimentation with vertical product upgrading and product diversification) positively affects the demand appeal of 20 count, we expect β_1 (β_3) to be positive.

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

Table 9C presents the estimation results. A firm's upgrade experiments and their interaction with subsequent low-end product diversification are positively associated with the firm-specific demand appeal in 20-count yarn, especially when using the instrument for upgrade experimentation.²⁶ Interestingly, the coefficient on the indicator for a firm employing a university-educated engineer is also positive and statistically significant in all specifications. Thus, engineering talent has an independent effect on demand appeal of 20-count product. On the other hand, the number of across-count portfolio rebalancing conducted by a firm around 20 count in a given semi-annual period is negatively related to the demand measure. This suggests that during normal times, changing the portfolio balance across counts is associated with lower demand appeal of the firm's 20-count yarn, all other things equal.

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

²⁶ Similar to Table 7 above, we instrument the endogenous number of upgrade experiments conducted by firm i at time t by using the exogenous variation provided by the expansion of high-end machine capacity during period t , interacted with the degree of relative output cuts imposed on low-end products.

two quality indices are by and large consistent with each other, with a correlation between the two metrics of 0.63.

6.3 Robustness

6.3.1 Survival rates

Our results indicate that firms that installed high-end machines and conducted product upgrade experiments grew faster. While at first glance this might seem to imply that this reflects 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 experimenting firms having a lower chance of surviving. While such a variance increase might still be preferable for firms (limited liability companies are essentially an option, after all), this mechanism is qualitatively different than one where upgrade experimentation simply raises expected growth. To explore this possibility, we investigate survival patterns in our data.

In Panel A of Table 10, we present summary statistics on the status of firms at the end of our sample in 1914. Among the 105 firms for which we have machine data, 42 firms (40 percent) had high-end machines at some point. There were 33 firms that survived to the end of the sample, and 19 of these (58 percent) had high-end machines. Thus firms with high-end machines had a substantially higher probability of survival, not a lower one. This is not consistent with product upgrade experimentation being primarily a device that increases the variance of firms' growth rates.

Moreover, among the 72 exiting firms, we can distinguish those that exited by acquisition (53 firms) and by liquidating (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 increased the chances of exiting by 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 acquired firms that did not (a difference that is statistically significant at the 5 percent level).

Panel B of Table 10 reports results from hazard regressions using the complementary log-log model. Consistent with Panel A, the estimation result in column (1) in Panel B of Table 10 confirms that firms with high-end machines were more likely to survive than other firms, statistically significant at the 10 percent level.

--- Table 10 around here ---

To see more precisely what role the complementarities between upgrading and diversification may have played in survival, we limit observations to the time after high-end machines were first installed at each of the 42 firms that had high-end machines. We then calculate the average fraction of high-end products in the total number of all products produced by the firm over all observations. The idea is that firms that installed high-end machines but were unable to develop high-end products would suffer in terms of their survival probability, but so also would the firms that failed to convert the technology developed to produce high-end product varieties into longer-term expansion of the number of low-end product varieties. Hence, both too low and too high fractions of high-end products should be associated with higher risk of exit. Surviving firms should come predominantly from the middle part of this distribution, where they have a balance between the number of high-end and low-end products.

The estimation results in column (2) of Table 10, Panel B show this is exactly what happened. An indicator for a firm having an average fraction of high-end products between 30 and 70 percent is associated with more than 50 percent higher survival chances. This is statistically significant despite the small number of observations. In terms of raw data (not shown), 11 firms fall into this middle range of the fraction of high-end product varieties and nine of them survive; 31 firms fall into one of the two extremes (average fraction of high-end products less than 30 percent or higher than 70 percent) and only 10 of those survive.

To sum up, having high-end machines appears to (a) improve the firm's chances of survival, especially if it could develop high-end products successfully and then exploit complementarity between product upgrading and diversification to also expand low-end products and achieve a balanced product portfolio; (b) improve the chances of exiting by acquisition as opposed to shutting down, conditional on exit; and (c) improve the returns to shareholders conditional on exiting by acquisition. Thus the higher growth rates of firms that introduced high-end machines did not entail higher risks. Entrepreneurial action and the removal of supply constraint benefited both growth and survival chances.

6.3.2 Mergers and acquisitions

Firms can expand the range of their product varieties and scale up production through mergers and acquisitions. It could be that mergers and acquisitions generate in part the patterns of the complementarity between product upgrading and diversification that we documented above.

We examined this possibility by using industry M&A information (Braguinsky et al., 2015) and found that less than four percent of new product introductions coincide with an acquisition event. Conditional on new product introduction during a given semiannual period, the average number of newly introduced products is 1.8 in absence of a merger or acquisition and 2.3 when one occurs. Thus while there is some uptick in new product introductions at the time of acquisitions, acquisitions played a rather modest role in product variety expansion. 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.

6.3.3 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 temporal complementarity between product upgrading and subsequent diversification. 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 shares of firms operating in the low- and high-end product markets, the ratio of industry-level high-end output 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 a dependent variable and the number of firms in low-end product and high-end product markets and the average real price of 20 count yarns as explanatory variables, but 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.

6.3.4 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, particularly so for high-end varieties. Hence it does not appear that exporting was a driving factor in our results above.

This is explained by the nature of industry exports during our sample. 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). So 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.

7 Conclusions and Discussion

Using detailed historical panel data from the Japanese cotton spinning industry, we showed that 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. Subsequently, they increased the number of product varieties they produced inside their frontier. In other words, they first engaged in 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. The process of 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 in general lead to sustained growth in either output or the number of product varieties produced. 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 influenced by product upgrade experimentation.

The relationship between product variety expansion and firm growth uncovered in our study is not considered in standard models. 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. 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 contrast, in our case, introducing technologically more challenging products is tied much more strongly to growth than simple horizontal product proliferation. This highlights the importance of incorporating into analytical frameworks the

heterogeneity with respect to the particular ways and directions in which product variety expansion occurs.

References

- Bernard A B, Redding S J, and Schott P., 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 C J, Redding S J, and Weinstein D E., 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*. Oxford 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.

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)
All product lines ever produced	1,167	1.000	324	0.278
Of which: never experimental	620	0.531		
initially experimental	547	0.469	324	0.592
Of which: Upgrade lines	91	0.166	41	0.451
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

A. All product varieties								
Period	Total change (a)=(b)+(c)	Composition (b)	Within				Deviation from own average	
			Total (c)=(d) + (e)+(f)	Continuing firms			Entrants (e)	Exiting firms (f)
			(d)=(g)+(h)	Expansion (g)	Allocation (h)			
1893.2-1914.2	8.976	3.347	5.629	6.657	5.360	1.298	-0.234	-0.794
1893.2-1906.2	2.526	1.625	0.901	1.039	0.017	1.022	-0.254	0.116
1907.1-1914.2	6.450	1.722	4.728	5.619	5.342	0.276	0.019	-0.910

B. High-end product varieties								
Period	Total change (a)=(b)+(c)	Composition (b)	Within				Deviation from own average	
			Total (c)=(d) + (e)+(f)	Continuing firms			Entrants (e)	Exiting firms (f)
			(d)=(g)+(h)	Expansion (g)	Allocation (h)			
1893.2-1914.2	5.076	2.339	2.736	3.408	2.838	0.569	-0.124	-0.548
1893.2-1906.2	1.804	1.494	0.310	0.477	0.089	0.388	-0.144	-0.022
1907.1-1914.2	3.271	0.845	2.426	2.931	2.749	0.181	0.020	-0.525

C. Low-end product varieties								
Period	Total change (a)=(b)+(c)	Composition (b)	Within				Deviation from own average	
			Total (c)=(d) + (e)+(f)	Continuing firms			Entrants (e)	Exiting firms (f)
			(d)=(g)+(h)	Expansion (g)	Allocation (h)			
1893.2-1914.2	3.901	1.008	2.893	3.250	2.521	0.729	-0.111	-0.247
1893.2-1906.2	0.722	0.131	0.591	0.562	-0.072	0.634	-0.109	0.138
1907.1-1914.2	3.179	0.877	2.302	2.688	2.593	0.095	-0.001	-0.385

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)	(3)	(4)
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)		
Cumulative number of upgrade experiments with high-end machines at t			0.032*** (0.012)	-0.015 (0.019)
Cumulative number of upgrade experiments without high-end machines at t			-0.008 (0.024)	-0.027 (0.034)
Fraction of low-end products in total number of products at t	-0.018 (0.066)	-0.104 (0.080)	-0.023 (0.064)	-0.147* (0.083)
Cumulative number of upgrade experiments high-end machines x fraction of low-end Cumulative number of upgrade experiments high-end machines x fraction of low-end				0.087*** (0.031) 0.022 (0.026)
Dummy = 1 if university-educated engineer at t	0.096** (0.042)	0.101** (0.042)	0.102** (0.042)	0.104** (0.043)
Dummy = 1 if merchant board member at t	0.017 (0.025)	0.017 (0.026)	0.018 (0.025)	0.021 (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)	0.015** (0.007)	0.016** (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)	0.008 (0.014)	0.007 (0.014)
Ln(output) at t	-0.308*** (0.048)	-0.312*** (0.048)	-0.312*** (0.047)	-0.321*** (0.047)
Constant	2.497*** (0.356)	2.605*** (0.367)	2.528*** (0.350)	2.700*** (0.369)
Semiannual time dummies	Included	Included	Included	Included
Firm FE	Included	Included	Included	Included
Observations	1,608	1,608	1,608	1,608
R-squared	0.325	0.326	0.327	0.330
Number of firms	99	99	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 (1)	"Placebo test" (2)	Second stage (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	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		-0.089 (0.568)		
Logged installed high-end spindles in $t+1$, minus Logged installed high-end spindles	-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	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 $t-1$	0.221*** (0.075)	0.145* (0.083)
Cumulative number of diversification experiments at $t-1$	0.017 (0.014)	0.003 (0.026)
Dummy = 1 if university-educated engineer at t	0.275* (0.141)	-0.076 (0.240)
Dummy = 1 if merchant board member at t	0.049 (0.093)	0.119 (0.155)
Dummy equal to 1 if high-end machine expansion during period t	0.122 (0.129)	0.187 (0.120)
Dummy equal to 1 if low-end machine expansion during period t	0.139 (0.105)	-0.014 (0.082)
Total output (thousands of tons, adjusted to 20-count) during period t	-0.025 (0.028)	-0.096*** (0.034)
Constant	0.313 (0.229)	0.874*** (0.264)
Observations	1,605	1,605
Semiannual time dummies and firm dummies	Included	Included
Number of firms	99	99
R-squared	0.088	0.063

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 $t+1$	0.011** (0.005)
Change in the number of across-count portfolio rebalancing from t to $t+1$		0.013*** (0.005)
Dummy = 1 if university-educated engineer at t	0.103** (0.041)	0.103** (0.041)
Dummy = 1 if merchant board member at t	0.010 (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 $t-1$	0.009 (0.014)	0.009 (0.014)
Logged total output at t	-0.298*** (0.048)	-0.297*** (0.047)
Constant	2.409*** (0.337)	2.404*** (0.332)
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*** (0.008)	
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		0.003 (0.007)
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	-7.410 (7.059)
Number of across-count portfolio rebalancing (between 17-48 counts)	-0.097*** (0.028)
Constant	26.678 (34.562)
Semiannual time dummies and firm dummies	Included
Observations	743
R-squared	0.731

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Instrument: 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				
	Estimation	(1)	(2)	(3)	(4)
		OLS		IV	
Cumulative number of upgrade experiments at t	0.099*** (0.025)	0.012 (0.082)			
Cumulative number of upgrade experiments x fraction of low-end products at t		0.131 (0.113)			
Cumulative number of upgrade experiments at t (instrumented)			0.530*** (0.061)		-0.056 (0.223)
Cumulative number of upgrade experiments x fraction of low-end products at t (instrumented)					0.900*** (0.331)
Fraction of low-end products in total number of products at t	0.295 (0.301)	0.233 (0.304)	0.848*** (0.306)		0.362 (0.378)
Number of across-count portfolio rebalancing (between 17-48 counts)	-0.161*** (0.049)	-0.159*** (0.049)	-0.170*** (0.045)		-0.167*** (0.045)
Dummy = 1 if university-educated engineer at t	1.004*** (0.115)	1.013*** (0.116)	0.811*** (0.114)		0.812*** (0.114)
Dummy = 1 if merchant board member at t	0.254** (0.110)	0.236** (0.111)	0.095 (0.107)		0.009 (0.111)
Logged installed high-end spindles in $t+1$, minus Logged installed high-end spindles in t	0.051 (0.038)	0.048 (0.039)	0.048 (0.040)		0.045 (0.039)
Logged installed low-end spindles in $t+1$, minus Logged installed low-end spindles in t	0.348 (0.279)	0.340 (0.279)	0.278 (0.271)		0.274 (0.267)
Firm age	-0.018* (0.010)	-0.019* (0.010)	-0.051*** (0.011)		-0.059*** (0.012)
Constant	2.851** (1.423)	2.934** (1.425)	2.491* (1.412)		3.041** (1.432)
Semiannual time dummies	Included	Included	Included		Included
Observations	721	721	721		721
R-squared	0.618	0.619	0.641		0.645

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

Panel A

		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

Panel B

	Hazard regression	
	(1)	(2)
Dummy equal to one if had high-end machines	-0.466* (0.281)	
Dummy equal to one if had high-end machines and mean fraction of high-end products between 30 and 70 percent		-1.648** (0.745)
Educated engineer dummy	-0.412 (0.314)	-0.714 (0.444)
Merchant board member dummy	-1.815*** (0.326)	-1.814*** (0.497)
Firm age	0.249* (0.133)	0.237 (0.262)
Constant	-2.716*** (0.313)	-2.478*** (0.674)
Number of firms	103	42

Estimated coefficients are reported. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.