

Productivity Adjustments and Learning-by-Doing as Human Capital

by Jim Bessen © 1998*
Research on Innovation
jbessen@netreach.net
610-565-7517

Working Paper

Original: 8/20/97

Revised: 3/98

Abstract: This paper measures plant-level productivity gains associated with learning curves across the entire manufacturing sector. We measure these gains at plant startups and also after major employment changes. We find: 1.) The gains are strongly associated with a variety of human capital measures implying that learning-by-doing is largely a firm-specific human capital investment. 2.) This implicit investment is large; many plants invest as much in learning-by-doing as they invest in physical capital and more than they invest in formal job training or R&D. 3.) This investment differs persistently over industries and is higher with greater R&D; hence measured returns to R&D may be overstated. 4.) Consistent with a search model, the human capital investment is much larger following employment decreases than increases. We conclude that learning-by-doing is a major factor in wage determination, the adoption of new technologies and asymmetric employment adjustment costs.

* Thanks to John McClelland whose wisdom and hard work compensated for many of my shortcomings. Thanks also to John Haltiwanger, Barry Hirsch and David Macpherson for data and to Martin Baily, Gale Boyd, Doug Dwyer, John Haltiwanger, Ron Jarmin, Boyan Jovanovic, Katsuya Takii and seminar participants at the CES and the University of Pennsylvania for helpful comments. The research in this paper was conducted while the author was a research associate at the Center for Economic Studies, U.S. Bureau of the Census. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census or the Center for Economic Studies.

Learning-by-doing is widely assumed to be a source of human capital. Becker [1965] initially included learning-by-doing as a source of human capital, and labor economists routinely treat it as an informal counterpart to job training. Macroeconomists sometimes treat learning-by-doing as an important adjustment cost associated with employment changes and growth theorists treat it as a costly source of skills needed to adopt new technologies (see reviews by Hamermesh and Pfann [1996] and Jovanovic [1997]).

Yet little empirical evidence firmly establishes learning-by-doing as a source of human capital. Little evidence demonstrates that the observed productivity gains attributed to learning are, in fact, the result of knowledge gained through experience. In most of the literature learning-by-doing is examined strictly as a *production* phenomenon—productivity increases over time after plant startup and this increase is correlated with accumulated production volume.¹ As such, this correlation alone does not provide the basis for an inference that experience *causes* the productivity gains or that this experience constitutes a form of human capital.

Also, little attempt has been made to gauge the magnitude and frequency of this implicit human capital investment. A substantial case study literature on learning curves associates learning-by-doing with large productivity gains, often ranging from two-fold to ten-fold (see [Argote and Epple, 1990], [Dutton and Thomas, 1984] for reviews of this literature). Yet no studies look at the entire manufacturing sector. The most comprehensive study to date, by Bahk and Gort [1993], examines 15 four digit industries (pooled data for 41) and their methodology does not estimate the human capital investment. Also, almost all the studies look only at learning-by-doing following plant startup, not after the introduction of new technologies, the reorganization of production, or other technology shocks.

The link between this production phenomenon and human capital is significant because if, in fact, the large observed productivity gains result from an investment in learning, then this investment must be particularly valuable. Moreover, if learning-by-doing occurs frequently throughout industry, then the total investment in learning could be quite large. Large "hidden" investment in what is not a spot commodity has strong implications for wage determination, for macroeconomic adjustments in general and the process of technological change in particular.

This study attempts: 1.) to measure the breadth and depth of the productivity adjustments that occur during plant ramp-ups and also following large changes in production worker employment across the entire manufacturing sector, and, 2.) to determine the extent to which these production phenomena appear to result from human capital investment. We measure the productivity adjustments using plant level data from the Longitudinal Research Database developed by the Center for Economic Studies of the U.S. Bureau of the Census and we associate these measures with a variety of indicators of human capital investment.

We find that on average the investment implicit in productivity adjustment, measured by foregone output, is large relative to investment in physical capital. On a per-worker basis it is much larger than investment in formal job training. Moreover, the productivity adjustments exhibit important differences between industries and these differences appear persistent—they are correlated between the two types of events—and related to technical change and technical complexity. We also find that the productivity adjustment made after large changes in employment is highly

¹ Exceptions include Argote, Beckman and Epple [1990] who examined the link between learning, worker turnover and plant disruptions. Also, Jarmin [1996] found a link between learning-by-doing and worker wages in the instrument industry.

asymmetrical, much larger for decreases in employment than for increases, consistent with a human capital interpretation and also consistent with recent research on adjustment costs.

To explore whether these productivity adjustments can be interpreted as a human capital investment we performed a number of regressions at the plant level and, using aggregated data, at the industry level. At the plant level we found significant relationships between wages and productivity gain for both production workers and non-production workers. Also consistent with human capital theory, we found a significant relationship between the *change* in wages over the first five years and productivity gain for both classes of workers.

Industry level regressions found significant positive associations between productivity adjustment and job training programs, full-time year-round employment as a percent of employment, union contract coverage, post-secondary and college education, the percent of scientists and engineers in the workforce, and possibly work experience. Negative associations were found between productivity adjustment and job turnover, part-time work, and low levels of education.

Thus these productivity adjustments appear to result substantially from firm-specific human capital investments. Since these investments cannot be accounted for by formal job training, the productivity adjustments must therefore result from learning-by-doing. Moreover, these are relatively large human capital investments, perhaps an order of magnitude larger than the investment in formal job training in manufacturing. Learning-by-doing is thus important for understanding labor markets and technical change, for explaining macroeconomic employment adjustments and for accounting for the economic growth of nations.

Background

General theoretical approach

In order to measure the breadth and depth of learning-by-doing across the manufacturing sector it is necessary to derive a more general methodology than that used in previous studies. Moreover, a general survey must empirically distinguish learning-by-doing from other causal explanations of productivity gain rather than simply assume that learning causes these gains.

Following Bahk and Gort, we define learning-by-doing as the accumulation of knowledge and skills as a by-product of the production of goods and services. Our focus also is on *firm-specific* learning-by-doing because, presumably, to the extent that learning-by-doing generates general human capital, that knowledge is readily available to a new plant and would not need to be generated anew.

For a large panel of plants Bahk and Gort estimate a production function that includes a term for accumulated production volume. The coefficient of this term provides a presumed measure of learning. This approach follows the case study literature which finds a correlation between productivity improvement after startup and accumulated production volume—log of productivity is roughly a linear function of log volume up to a "plateau" point.²

² Other studies taking this approach include Jarmin [1996], Rapping [1965], and Argote, Beckman and Epple [1990].

A general survey of learning-by-doing might follow this approach and attempt to measure a correlation between productivity and cumulative volume at a large number of plant startups. There are three important limitations to this approach, however.

First, one would expect learning-by-doing to occur following a variety of different "shocks" to the production environment. For example, learning-by-doing might follow a change of workplace organization or the introduction of a new technology. Ideally one would hope to measure such learning even though simple cumulative production volume might no longer be an appropriate index of experience in these cases.

Second, the production function approach does not directly permit estimation of the human capital investment. Following standard human capital theory, the investment in firm-specific learning-by-doing is measured as foregone output. Calculating foregone output requires the time path of inputs in addition to a production function.

Finally, these studies rely on the rather strong assumption that the observed productivity gains are *caused* by accumulated experience. A general survey must realistically contend with many other possible causal explanations and sort out the relative significance of experience.

For instance, productivity growth following startup could be caused by formal training of the workforce, or inherent delays in the production process (such as the aging of the first batch of wine at a winery) or a slow ramping of capacity utilization. However, the common functional form of the productivity gain—the “learning curve”—over a wide variety of case studies has led researchers beginning with Wright [1936] to infer that the productivity gain is caused by “experience” in some form and that accumulated volume is a practical index of experience.

But still this inference ignores a variety of other possible explanations for a consistent association between productivity improvement and production volume. Some of these explanations might be characterized as different types of learning-by-doing, some might not be learning-by-doing at all, and the potential causes are so inter-related that it is difficult at best to untangle exactly what portion of the productivity gain can be attributed to learning-by-doing. For instance,

- If managers experience convex adjustment costs in ramping up a new plant, that is, if managerial effectiveness declines with the number of new processes concurrently managed, then managers might limit the number of new processes handled at any one time. They could do this by ramping up different production processes over time, while outsourcing the remaining production steps. Productivity might then increase with production volume, but for reasons that are not clearly related to learning-by-doing.
- The learning that takes place could be along the lines of the job-matching models. Employers (and workers) could learn of the quality of the job match (worker productivity, job suitability) over time. As this knowledge is gained, less productive workers might leave or be fired. Productivity then increases with experience, but the learning does not result in human capital among production workers (although it could be considered human capital for managers).
- The productivity improvement may result from R&D. R&D is often described as a search process and part of this search might occur during actual production. To the extent that production experience allows scientists and engineers to make productivity improvements this can be described as learning-by-doing. One might further consider the knowledge learned to be part of the human capital of the scientists and engineers. However, it is then conceptually difficult to attribute the source of the productivity gain to learning-by-doing as opposed to R&D.

And certainly industry-wide productivity effects—such as cyclical changes in capacity utilization or diffusion of knowledge—are likely to affect plant productivity during startup.

Thus the inference that experience causes the productivity gain is not warranted without further examination of alternate explanations. Indeed, in one of the few case studies that looked at micro-data of individual worker performance, Lazonick and Brush [1985] found a more complicated causal story. David [1975] initially examined productivity and cumulative volume for a large textile mill and found a classic learning curve relationship. Re-examining the same mill with a much richer data set, Lazonick and Brush found that part of the gain could be attributed to worker experience, but that part was also attributed to greater worker effort elicited through a speed-up program and a change in the composition of the workforce.³

Of course, this sort of micro-data is available for few case studies, let alone industry or sector-wide studies. Hence if we admit the possibility that other factors might cause the gain in productivity occurring at startup, then a general survey must measure learning-by-doing by some means others than the correlation between productivity and cumulative volume.

Rather than base our measurement of learning-by-doing on this specific correlation, we take a more general approach. First, we consider learning-by-doing to be one instance of a more general phenomenon of “productivity adjustment.” That is, following certain discrete changes in production—the installation of a new plant, a re-organization of production, or the introduction of a new technology—plant productivity adjusts slowly, increasing over time for a limited period. The cause of this adjustment may be learning-by-doing or it may be something else, but in either case, the magnitude of the adjustment can be gauged using standard multi-factor productivity measurement techniques.

Second, we determine whether some or all instances of this productivity adjustment can be associated with the attributes of firm-specific human capital. That is, we determine whether plants or industries with greater relative levels of productivity adjustment are more likely to exhibit patterns of wages, employment attachments, etc. indicative of firm-specific human capital. *If* this productivity adjustment is associated with firm-specific human capital *and if* this human capital cannot be accounted for by formal training, *then* we may infer that the productivity adjustment is substantially caused by firm-specific learning-by-doing. More generally, we know that the implicit investment acts like human capital even though we may not know all the specific mechanisms producing the productivity change.

This approach is elaborated as follows.

Measurement of productivity adjustment

a. Absolute productivity change

The simplest measure of the magnitude of a productivity adjustment is a Tornqvist-Divisia index of multifactor productivity assuming, in effect, that the relevant productivity change is Hicks-

³Even here the causality is tangled: managers may have learned through experience what level of effort could be extracted from the workforce. Also, worker learning through experience may itself have a significant dependence on unmonitored effort if, in fact, the learning occurs through a search process.

neutral technical change. The change in total factor productivity between periods t and T is measured

$$\Delta_{tT} = \ln Q_T - \ln Q_t - \sum_{i=1}^4 \left(\frac{s_t^i + s_T^i}{2} \right) \cdot (\ln X_T^i - \ln X_t^i) \quad (1.)$$

where s_t^i is the output share of input i at time t and X_t^i is the measure of the input of i at t , and the four inputs are production labor, non-production labor, intermediates and capital. The details of this calculation are described in Appendix A.

b. Relative productivity change

Inevitably, plant productivity changes arise from a variety of sources, some of which may be specific to the plant, others common to the entire industry or economy as a whole. For example, cyclical shocks may change capacity utilization, influencing productivity. Also, industry-wide diffusion of technical knowledge may boost an individual plant's productivity.

Learning-by-doing would, of course, only be associated with that portion of productivity change specific to an individual plant. In this sense, the absolute productivity change measure may overstate the appropriate magnitude we wish to associate with human capital measures. As an alternative, we may also consider a *relative* productivity change

$$\begin{aligned} \Delta_{tT}^R &\equiv \Delta_{tT} - \Delta_{tT}^I \\ &= \Delta_{tT} - \bar{\Delta}_{tT}^I - \delta_{tT}^I \end{aligned} \quad (2.)$$

where Δ_{tT}^I is the industry productivity change calculated in parallel manner. This industry change can be further decomposed into a long-term average productivity change (averaged over the entire sample period), $\bar{\Delta}_{tT}^I$, and a deviation or "shock," δ_{tT}^I , for the particular period being studied. This deviation term thus captures industry-wide cyclical changes.

For a number of reasons, learning-by-doing might be causally associated with a higher rate of industry productivity change. Highly innovative industries may have newer technologies that require greater learning-by-doing effort. Also, the diffusion of general industry knowledge may require learning-by-doing to apply this knowledge to a specific plant (Cohen and Levinthal [1989]). Thus this relative measure of productivity adjustment may understate the productivity change associated with learning-by-doing. This may be particularly true in industries such as computers which experience very high rates of industry productivity growth. The absolute and relative measures may be thought of as upper and lower bound measures to bracket the relevant productivity change.

c. Implicit adjustment investment

The impact of a productivity adjustment is determined not only by the size of the productivity change, but also by the duration of the adjustment. A long, slow adjustment can be quite costly to a plant. A straightforward means of measuring the *complete* magnitude of a productivity adjustment is the foregone output, the difference between what the plant could have produced at full productivity during the adjustment period and what the plant actually produced. Assuming that the adjustment process is complete by the fifth year (this is largely supported by the case

study evidence and by our own results below), the foregone output in absolute and relative terms is⁴

$$\Delta Q = \sum_{t=1}^4 Q_t \cdot (e^{\Delta t^5} - 1), \quad \Delta Q^R = \sum_{t=1}^4 Q_t \cdot (e^{\Delta t^R} - 1). \quad (3.)$$

These measures can be interpreted as working capital investments. Ignoring discounting of the foregone output, the value of a plant can be expressed as

$$V \equiv \frac{P}{d} - \Delta Q \quad (4.)$$

where π is the profit rate at full productivity, d is the discount rate, so that π/d can be viewed as the return on investment ΔQ . Under free entry (with zero additional fixed entry cost), V will be driven to zero so that these two terms will be equal.⁵

Note that this treatment mirrors the standard treatment in human capital theory. In Becker [1965] or Hashimoto [1981] the investment in firm-specific human capital is shared between the firm and the worker and the total investment is measured by the foregone output (the worker's investment is foregone wages). Thus these measures of absolute and relative foregone output both provide the most general metric of productivity adjustment and also provide a measure of investment should the productivity adjustment be substantially caused by human capital.

Figure 1 displays the monthly productivity of a startup of a steel process reported in Baloff [1966], in this case productivity measured as tons per hour. The two types of productivity measures are displayed, Δ as the gain in productivity over the startup, and the area ΔQ as the cumulative foregone output.

Associating productivity adjustment with human capital

If, in fact, a substantial portion of the productivity gain is caused by firm-specific human capital, then a variety of associations should be found between the productivity measures and labor characteristics. Recent research has modeled learning-by-doing as a search process ([Muth, 1986], Jovanovic and Nyarko [1995], Auerswald, et al [1996], Bessen [1998]). We generally refer to these models in forming hypotheses relating productivity adjustments to labor characteristics.

First, to the extent that the human capital is firm-specific, workers have opportunity for *ex post* bargaining (Shaked and Sutton, 1985). According to traditional human capital theory, this gives rise to a sharing of the investment and an upward sloping wage-tenure profile [Becker, 1965, Hashimoto, 1981]. Firms will pay above-market wages in later periods and this should induce lower turnover and more permanent employment attachments.

Also, this human capital investment may be costly to monitor. If learning is achieved through a search process, then it may be difficult for managers to monitor a worker's effort at learning: poor performance over a short period of time may be due to lack of effort or simply to unlucky search trials. Moreover, the employee who acquires this knowledge may shirk or may quit, thus

⁴For simplicity, we do not discount this quantity. Since the largest gains occur during the first two years, this is a minor difference.

⁵This assumes that the firm selects an optimal time path for this working capital investment.

obviating the investment. This gives rise to implicit labor contracts or other complex means of eliciting worker effort and insuring against a loss of shared human capital investment ([Lazear, 1979, Carmichael, 1983]). This also gives rise to an upward wage-tenure pattern.

These attributes of human capital investment imply that wage determination will not occur as a simple spot market for homogenous labor and also that the introduction of new capital, new staffing levels and new technology will incur substantial costs. Generally, plants with greater learning investment should be associated with a better quality (in some sense) workforce, both because firms will want to selectively hire workers who are more likely to learn and also because the workers' better skills result in greater measured learning. Second, plants with greater learning investments will want to preserve these investments.

Thus if the process of productivity adjustment arises substantially from a human capital investment then standard results give us the following hypotheses:

1. Plants with greater productivity adjustment should pay higher wages, either because they hire better quality workers or because workers are able *ex post* to bargain a share of the returns to learning.
2. Plants with greater productivity adjustment should exhibit a greater increase in wages with tenure, especially during the first years of employment. This is either because workers' productivity increases in later periods [Becker, 1965, Hashimoto, 1981], or because an implicit labor contract is used to minimize shirking in the learning process [Lazear, 1979, Carmichael, 1983].
3. Plants with greater productivity adjustment should, in general, demonstrate more permanent employment attachments, that is, less part-time and temporary work and less job turnover.
4. Plants with greater productivity adjustment might also be more likely to have unions, either because firms recognize that unions may increase workers' attachment to the firm or because of workers' stronger *ex post* bargaining.
5. Plants with greater productivity adjustment may hire a more educated workforce, either because the additional education provides necessary skills or because the greater education signals a better ability to learn.
6. Plants with greater productivity adjustment may hire more experienced workers, either because these workers have acquired general skills useful for learning or because their work records signal a greater facility at learning.

If these hypotheses appear to be true, then it is reasonable to infer that the productivity adjustments are generated by some form of human capital. Additional hypotheses may shed some light on the particular nature of that human capital, for instance on the role of formal training:

7. Plants with greater productivity adjustment might also exhibit a higher level of company-sponsored formal job training because the formal training may be part of the learning process and/or because formal training may complement the informal learning process.

Additional hypotheses may be generated from a view of learning-by-doing as a complex search process. If learning-by-doing is largely human capital, an important additional question is exactly *who* acquires the knowledge. Zevin [1975] has suggested that learning-by-doing may be strictly a managerial phenomenon. Bartel and Lichtenberg [1987] propose that learning-by-doing is dominated by non-production workers hence the proportion of non-production workers declines after an initial learning period.⁶ On the other hand, the search models of learning-by-doing imply that large numbers of ordinary production workers are involved in the process of learning-by-doing. Taking the latter view, we further propose:

8. The above associations should hold for production workers as well as for non-production workers.

Nevertheless, some non-production workers, namely scientists and engineers, might be particularly important for learning-by-doing. The contribution of production workers to learning-by-doing may be, in fact, a necessary *complement* to the activities of scientists and engineers. This is what one would expect if new technologies have to be "debugged" in production. Anecdotally, new technologies frequently have a large number of "manufacturability" problems to be resolved by the trial-and-error efforts of production workers, hence we might expect formal innovation to be accompanied by greater learning-by-doing.

9. Greater productivity adjustment should be associated with a greater proportional employment of scientists and engineers (and greater R&D activity generally) because these skills partially cause the productivity gain and/or because the innovations developed by scientists and engineers require greater learning-by-doing to "debug."

The discussion so far has considered learning-by-doing largely in the context of plant startup. If (and only if) a large portion of the production workforce acquires knowledge through learning-by-doing, then large changes in workforce staffing may also require substantial investments in learning, hence substantial adjustment costs.

Moreover, the magnitude of this investment may depend on the nature of the organizational change associated with the change in employment. Bessen [1998] has argued for a link between learning-by-doing and the division of labor—the search process varies with the complexity of each worker's tasks. A greater degree of specialization implies easier search hence a faster rate of learning; but the greater specialization also limits the latitude of search resulting in a lower final productivity gain. An employment change accompanied by a total re-organization of the division of labor will, according to this theory, require a greater investment in learning, all else equal.

This suggests that the productivity adjustment might be asymmetrical with regard to the direction of the employment change. An increase in employment can typically be accomplished without disrupting the workflow of the existing, fully trained workforce—additional shifts can be added or additional facilities built (see Bresnahan and Ramey [1994] for a discussion of these adjustments in the auto industry).

However, a decrease in employment may require a re-organization. Of course if a shift has been added, it can be removed without incurring re-organization. But further decreases will require an

⁶Our data, in fact, show the opposite pattern, namely, the proportion of non-production workers increases over the first five years of a new plant as learning is largely completed. Bartel and Lichtenberg look at a longer time period and further analysis suggest that the entire picture of relative employments is quite complex.

established facility to be run by fewer production workers, hence tasks must be re-defined, hence a large portion of these workers must learn new tasks. In fact, we find that most episodes of large negative employment changes occur in plants where employment is declining in the long run, so re-organization will be common. Thus we propose the additional hypothesis:

10. A significant productivity adjustment should follow large changes in production employment and this investment should be larger following decreases in employment.

Last, the above discussion suggests two reasons for learning to vary between industries. First, some industries have higher rates of innovation that may require greater investment in learning as a complement to R&D activity. Second, some industries may have more complex technologies also requiring a greater investment in learning. Thus,

11. Different industries should exhibit persistent differences in the level of productivity adjustment and its associated investment.

The empirical plan is first to measure productivity adjustment after plant startups and following large employment changes and then to test these hypotheses using these measures. Following that, the paper interprets these results looking at the specific role of learning-by-doing.

Empirical Results I: measurement of productivity adjustments

Samples

We begin by measuring productivity adjustments. The samples are drawn from the Longitudinal Research Database (LRD) of manufacturing plants developed by the Center for Economic Studies of the U.S. Bureau of the Census [McGuckin and Pascoe, 1988]. This database is derived from the Census of Manufactures, which surveys all manufacturing plants every five years, and the Annual Survey of Manufactures, conducted on a sub-sample during interim years. The database is constructed to be representative of the entire manufacturing sector.

Our samples were chosen to capture distinct events where learning might be expected to occur, namely, new plant start-up and large employment changes at long-term continuing plants. Such samples are necessarily selective and do not represent the entire population of manufacturing plants, however, we are able to demonstrate persistent industry characteristics captured in these samples.

The sample of startups is a balanced panel of new plants from all industries except tobacco that survived at least five years and for which we have annual data for each of the initial five years. The sample includes only plants which were reported as new their first reported year, beginning in years 1972 - 87. Plants with missing or imputed data were excluded as well as plants that had capital changes greater than $\ln 2$ between years three and five (such plants were either re-starting or dismantling) and plants where the change in the ratio of output to deflated materials changed

by more than **ln 2** between years one and two and between years four and five (probable reporting error).⁷

This resulted in a sample of 5,625 plants. Characteristics of these plants are shown in Table 1. The selection criteria used to derive this sample tend to slightly under-represent small plants and plants owned by single-plant firms. By comparison, Table 1 also presents the characteristics of new plants which first entered during years 1972, 1977, 1982 or 1987, were captured in the Census of Manufactures for those years and which also survived at least until the following Census.⁸ To make our analysis more representative of new entrants generally, weights were derived from this Census panel based on size class and single-plant ownership status and these weights were applied to the 5-year startup panel for all subsequent analysis. This weighting procedure did not affect results substantively.

Although comprehensive, the LRD has two important limitations for studying learning-by-doing among startups by comparison to the large case study literature. First, the frequency of observation is annual as opposed to more frequent, typically monthly, data found in case studies. Second, and more important, the LRD does not accurately capture the actual startup (partial) year.⁹ The case study literature demonstrates that the most dramatic productivity progress occurs during the first several months of production. Hence estimates of productivity adjustment based on our sample are necessarily *sharply* understated. This is illustrated in Table 2, which calculates the simple productivity measures for the data shown in Figure 1. Changing from monthly to annual data can sharply reduce the measured magnitude because of a dilution effect. Losing initial months further reduces the measure. Note that this startup was particularly rapid, taking about six months versus an average of 15.5 months for Baloff's entire sample.

The analysis of large employment changes was based on a sample of continuing plants from the LRD.¹⁰ The base file included all plants which were included in every Census of Manufactures and Annual Survey of Manufactures from 1972 - 93. This requirement insures the sample is almost exclusively large plants (over 250 total employees). After screens for reporting error, this base file contained 10,916 plants.

From this base file two sub-samples were extracted: an "up spike" file and a "down spike" file. The Up file included all episodes where a plant's production worker employment rose by 20% or more in a given year and the plant did not experience a change in production worker employment of absolute magnitude greater than 20% (up or down) during the subsequent four years.¹¹ This four-year requirement was necessary to allow consistent measurement of learning investment. The Down file included all episodes where production employment fell by 20% or more and was

⁷Additional screens were provided to check for large changes in rentals of plant and equipment where annual rentals are greater than capital stocks and also for decreases in inventories that exceeded 2/3 of shipments (indicating a possible change in inventory accounting method).

⁸This sample excludes plants with zero shipments or production workers or imputed values.

⁹The overwhelming majority of startup records show full production from the first quarter and many of the records showing zero production the first year appear to be "placeholder" records where data was not collected, but entered as zeroes. Initial year observations with zero shipments were discarded and the following year was treated as the first observation in the sample.

¹⁰We thank John Haltiwanger for making this sample available to us.

¹¹The 20% cutoff identifies relatively large changes, yet provides a reasonably large sample size. Studies of employment flows find that changes of similar magnitude or greater account for a substantial part of total flows [Davis, Haltiwanger and Schuh, 1996].

stable for four subsequent years. Note that a given plant could appear in either file more than once. Summary statistics are provided in Table 1 ("year 1" is the year of the employment spike).

The magnitude of productivity adjustment following new plant startup

Figure 2 displays the absolute and relative delta measures for each year averaged over the entire sample of new plants. Judging by the relative TFP measure, productivity adjustment relative to the industry appears to be complete by the third year (recall that the initial partial year is missing). This corresponds generally with the case study literature. For example, in Baloff's study of 26 capital-intensive startups, the average duration of the learning phase (prior to "plateau") was 15.5 months, although a number of the plants had not yet reached the plateau. Similar durations are found in the Liberty Ship data.

The magnitude of productivity gains differs sharply from the case study literature, however, as can be seen in Table 3. An average productivity gain of 5.7% is much smaller than the 100% + gains found in the case studies. As explained above, the annual frequency and loss of initial months will attenuate these measures. Rough calculations of a classic learning curve, dropping the first six months, calculating annual averages and assuming a plateau at 24 months, produces productivity gains of about 15%—somewhat higher than the best industries reported here, but not wildly so. Considering that the case study literature may over-represent plants with large productivity gains, our results are not inconsistent and perhaps more representative.

Nevertheless, the measures of foregone production are quite substantial. The more conservative measure based on the relative TFP change averages \$353,000 per plant. A correction for the missing initial months (based on case study data tabulated in [Bessen, 1998]) suggests the actual figure might be roughly triple this amount. The base amount represents 12% of the mean capital investment; divided by the average number of employees, this represents over \$6,000 per employee. Moreover, looking at some of the industries that exhibit greater productivity adjustment, clearly many plants may invest more in productivity adjustment than they invest in physical capital. And in many industries the implicit investment per employee is far greater than company expenditures on job training or expenditures on formal vocational training.¹²

Possible associations between the magnitude of productivity adjustment and plant characteristics can be explored through regression analysis. Table 4 shows results of regressing Δ against several plant and industry characteristics including size class, the portion of production workers in the workforce, capital-production labor ratios and whether the firm owns only one plant. We also include measures industry productivity change at the 4-digit level: the average industry TFP change over the sample period, $\bar{\Delta}^I$, and the industry deviation from average during each plant's startup period, δ^I .

A number of interesting associations emerge. Productivity adjustment appears to be larger at more capital intensive plants and at plants owned by multi-plant firms (where the quality of management may be better). Plants initially under 10 employees appear to perform better, and the portion of production workers has a positive, but not statistically significant effect. Industry

¹² Means from the NLSY (discussed below) suggest that about 7% of production workers receive company training (including informal training) during each year on average and that this training typically lasts about a week. Given mean annual pay of about \$20,000, employers' average cost of subsidizing formal training over five years is roughly $5 \times .07 \times .02 \times \$20,000 = \$140$ per total employee. Employers may also lose some production during training (although for a startup, training can be done before production commences) and non-production workers are trained more frequently. Nevertheless, the magnitudes of investment are clearly far less.

shocks have an effect that is not statistically different from 1.0 as one might expect. Average industry productivity change has a positive, but smaller effect, suggesting an association between plant productivity adjustment and industry productivity change, but also systematic differences between the two. As above, the direction of causality is unclear.

In regression (2), industry dummies for 64 CIC industries (several of the smaller industry groups have been rolled up) are added. This increases the predictive power of the regression (significant at the 1% level) suggesting that independent industry factors are involved in productivity adjustment beyond industry productivity change. Generally, depending on the measure examined, the industry groups with the greatest productivity adjustment include most of the "high tech" industries (SIC 35 is weakened in the relative measures by the high industry TFP growth of the computer industry), but also some lower tech industries such as paper. Some research [Bessen 1996] suggests that learning might be associated with the complexity of the technology and these results do not conflict with that explanation.¹³

Productivity adjustment following large changes in production employment

Figure 3 displays the relative year-to-year change in productivity following up spikes and down spikes. We track productivity changes only *after* the year during which the spike occurred. As in the case of plant startups, we would expect *most* of the productivity change to occur during the year of the spike, however, we cannot isolate that portion of change during this year which can be attributed to learning-by-doing as opposed to capacity utilization changes. Moreover, if lumpy employment changes result from non-convex adjustment costs, then the year previous to the spike cannot be taken as a neutral baseline. Thus we discard the first measurement and the measures of learning obtained subsequent to an employment spike are quite similar to the measures obtained above for plant startups.

As can be seen, relative productivity changes after down spikes follow a course quite similar to plant startups although here the learning process may extend a bit longer. As predicted, the learning following an up spike is substantially less, and, in fact, the mean relative productivity change is not statistically different from zero.

Table 5 provides summary measures of productivity adjustment following each type of event.¹⁴ The mean implicit investment relative to capital and per-employee following down spikes is quite large (at over \$18 thousand per employee), larger than for plant startups. Again, this is a lower-bound estimate and the actual value may be several times higher. The investment following up spikes is smaller and so the overall pattern is highly asymmetric. This asymmetry corresponds to a similar asymmetry found in macroeconomic studies of adjustment costs (see [Hamermesh and Pfann, 1996] for a summary).

Persistence of industry differences in productivity adjustment

One important issue is whether the productivity adjustments found in startups have anything to do with the productivity adjustments found after large employment changes, especially after large

¹³The only industry which appears to deviate from this description on first examination is SIC 29, petroleum and coal products. Further examination revealed, however, that in our sample this industry is comprised largely of asphalt plants, not oil refineries.

¹⁴Trimmed means (trimmed of the 1% tails in the delta measure) are presented for the investment measures to eliminate outliers.

down spikes. If, as proposed, industry differences in productivity adjustment relate to the complexity of technology and the rate of technical change, then industries with high levels of productivity adjustment after one type of shock should also show high levels of gain after the other type, and vice versa. Although these two events are drawn from very different populations of plants and are clearly responses to very different stimuli, there might well be a correlation between the levels of response to each type of event.

One way to estimate this correlation is to linearly regress mean industry measures of productivity adjustment from one panel against the corresponding industry means from the other panel (plus a constant term). This estimation is complicated, however, by the high level of within-industry variance due to idiosyncratic plant effects and the small numbers of plants in many 4-digit SIC industries. The resulting large sampling variance in both dependent and independent variables leads to both heteroscedasticity and to measurement error. This regression can be estimated, however, by maximum likelihood methods as presented in Appendix B.

There were 368 industries populated in both the startup panel and the down spike panel. The regression was performed on two measures of productivity adjustment: mean ΔQ^R divided by mean deflated output and mean ΔQ^R divided by total employees (year 5). Both measures were correlated between the two panels at a high level of significance.¹⁵ Thus industries respond similarly after both types of events, consistent with the view that industry differences in learning-by-doing affect these responses.

Summary of productivity adjustment measures

To summarize,

- The productivity adjustment at plant startup and following large changes in production employment was found to be quite large. In some industries plants made an implicit investment larger than their investment in physical capital. The levels of investment per employee were often much larger than investment in formal job training, in some cases exceeding a year's average wages.
- The levels of productivity adjustment exhibit strong differences between industries (and within industries) and these differences appear to be persistent. The cause of the productivity adjustment at plant startup must be quite similar to the cause of the productivity adjustment following large decreases in employment. Generally, large productivity adjustment seems to be associated with complex technology and a high rate of innovation.
- The productivity adjustment following large employment changes is highly asymmetric, much larger for employment decreases than for employment increases. This productivity adjustment therefore provides an explanation for the asymmetric pattern of adjustment costs found in Baily, Bartelsman, and Haltiwanger [1996] and other studies of cyclical adjustments.

¹⁵We use ratios of industry means rather than mean ratios to obtain a measure that is output-weighted. This is the better measure to use when comparing different samples of very different plant sizes.

Empirical Results II: Explaining productivity adjustment

Human capital and productivity adjustment at the plant level

If productivity adjustments are substantially caused by human capital, then we would expect new plants with greater productivity adjustment to hire better quality workers. Therefore, we should expect a positive association between productivity adjustment and wages. Moreover, if both production workers and non-production workers acquire useful knowledge through learning-by-doing, this association should hold for both. Jarmin [1996] found such a relationship in SIC 38. This relationship might be expressed

$$W_i = \alpha \cdot Z_i + \beta_1 \cdot \Delta^i - \beta_2 \cdot \delta^I - \beta_3 \cdot \bar{\Delta}^I + \varepsilon_i \quad (5.)$$

where W_i is the real wage at plant i in industry I , and Z_i is an array of plant characteristics.

Rather than positing a specific relationship between industry and plant productivity, we include the absolute plant productivity change and decomposed industry measures. Generally we might expect that $\beta_1 = \beta_2$ if industry shocks affect all plants uniformly—the industry shock term then just serves to “back out” the contribution of industry shock to plant productivity growth. Also we might expect that $\beta_1 > \beta_3$ because the association between average industry productivity growth and plant productivity growth is weaker and also because industries with higher average productivity growth might hire better quality workers at higher wages.

Table 6 shows estimations of equation (5) for the startup panel. Wages for production workers are measured as the value in the fifth year of average plant hourly wage relative to the mean wage for the entire manufacturing sector as obtained from the NBER Productivity Database (the difference between the log plant wage and the log manufacturing wage). For non-production workers relative mean salary was used. In addition to the other plant characteristics used above, we add a measure of the plant's productivity relative to its 4-digit industry:

$$R_i \equiv \ln Q_i - \ln Q_I - \sum_{j=1}^4 \left(\frac{s_i^j + s_I^j}{2} \right) \cdot (\ln X_i^j - \ln X_I^j) \quad (6.)$$

where the subscript i refers to the plant and I refers to the industry.¹⁶ The quantities are calculated as above using the NBER Productivity Database.

The results show a positive, significant relationship between relative wages and productivity growth. Note that the positive relationship holds for both production and non-production workers, with a higher coefficient for non-production workers. Wages also increase with capital intensity and plant relative productivity and decrease with single plant ownership and, for production workers, decrease with the proportion of production workers in the workforce. As expected, the coefficients on δ^I are not significantly different from the opposite of the coefficients on Δ . After controlling for industry fixed effects, greater plant size is also associated with higher wages.

¹⁶This measure is similar to the relative TFP measure used by Christensen, Cummings and Jorgenson [1981] and Baily, Hulten and Campbell [1992]. The difference is that this measure effectively uses arithmetic means (assuming approximately constant returns to scale) while the Christensen, Cummings, Jorgenson measure uses geometric means. Our concern is to capture the productivity advantage (disadvantage) of the plant's product compared to average industry product and hence the plant's "competitiveness." Thus a product-weighted measure is more appropriate. The components of this calculation are the same as for the productivity measures except that here constant returns are imposed on the capital share.

Another hypothesis about human capital that can be tested at the plant level is the association between productivity growth and the increase in wages. In Table 7 we estimate

$$\Delta W_i = \alpha \cdot Z_i + \beta_1 \cdot \Delta^i - \beta_2 \cdot \delta^I - \beta_3 \cdot \bar{\Delta}^I + \varepsilon_i, \quad \Delta W_i \equiv \ln \frac{w_5^i}{w_1^I} - \ln \frac{w_1^i}{w_1^I} \quad (7.)$$

where w_t^i is the current dollar wage at the i th plant in period t and w_t^I is the current dollar wage for the entire manufacturing sector at t .

Again, a significant association is found between productivity growth and wage increases, an association somewhat stronger among non-production workers. Also again, the coefficients on the cyclical shock parameter are close to the opposite of the coefficient on learning. For production workers this increase is larger at larger plants and at plants with more production workers proportionally. For non-production workers this increase is greatest among the smallest plants (less than 10 workers initially) and is less at more capital intensive plants and at plants with proportionally more production workers.

This analysis only measures changes in average wages and cannot distinguish changes in individual worker's wages from changes in the composition of workers employed at a plant. Our hypothesis suggests that individual workers' wages might increase with tenure, and because our panel is a startup, all workers are new hires, hence average wages at the plant should increase. On the other hand, it is possible that the change in productivity could result from a change in the quality of the workforce brought about by compositional changes and reflected by the higher pay.

To test this possibility, regression (7) was also run on the down spike panel. In this case, for both production and non-production workers, productivity growth was *not* significantly associated with the change in wages, and so the productivity growth could not result from a change in the quality composition of the workforce. This is consistent with our original hypothesis since the workforce now consists of experienced workers in the flat portion of the wage-tenure curve. Given that the industry pattern of productivity growth in the down spike panel correlates with the industry pattern in the startup panel, it does not seem likely that such compositional changes could explain much of the productivity growth in the startup panel either.

It is also possible that plants might change the composition of their workforce as a *consequence* (rather than a cause) of productivity growth. This might be true, for instance, if a higher level of output required a proportionally greater number of marketing and sales or clerical personnel. This might explain part of the change in non-production wages. Among production workers, however, it is not clear why a plant might want to higher better quality workers *after* it has achieved productivity improvements.

Thus these wage regressions suggest:

- Plants with greater productivity growth tend to have higher wages and a greater increase in wages over the first five years, consistent with a human capital interpretation.
- Both characteristics are true for production workers as well as non-production workers.

Human capital and productivity adjustment at the industry level

Ideally one would like to explore a much broader range of associations between productivity adjustment and human capital at the plant level; however, the LRD does not contain such measures.¹⁷ Instead, we can pursue an industry-level analysis. The above correlation between productivity adjustment measures in the startup panel and those measures in the down spike panel suggests that many productivity adjustments share a common industry-specific origin. For this reason, if the productivity adjustments are associated with human capital measures, these associations should also hold between aggregate industry measures, albeit with the possible introduction of considerable noise. Industry level measures, however, also allow us to test a wider variety of possible associations derived from other datasets.

Tables 8 and 9 list a variety of logit regressions for the probability that a given worker will have a specific characteristic:

$$P(\text{characteristic}) = \Lambda \left(\alpha \cdot \frac{\Delta Q^R}{EMP} + \beta \cdot Z \right) \quad \Lambda(x) = \frac{e^x}{1 + e^x}$$

where $\frac{\Delta Q^R}{EMP}$ is the ratio of mean relative foregone output for the industry to mean startup employment (year 1)¹⁸, and Z is a vector of other industry characteristics, namely, average four year industry productivity growth, net real capital per worker, and average annual industry employment growth.

Note that these other industry characteristics are all uncorrelated with $\frac{\Delta Q^R}{EMP}$, so although they have significant explanatory power themselves, the measure of productivity adjustment provides an independent source of variation. This is critical to our analysis.

The industry categories used are based on the Census Industry Classification (CIC), with some smaller industries rolled up, resulting in a maximum of 64 industries. This analysis is performed only on the startup panel because the down spike panel had substantially fewer industries with enough plants.

Table 8 is based on an unweighted pool of all manufacturing workers from the March Supplement to the Current Population Survey for 1979 and 1987 for privately-employed workers (excluding self-employed) who worked more than one week during the previous year.¹⁹ The regressions were estimated separately for production workers (occupation code 400-800 in 1979 and 500-900 in 1980) and non-production workers (everyone else).

The first set of regressions attempted to evaluate the permanence of the worker-employee match. Regression 11 estimated the probability that a worker was employed less than 50 weeks the previous year and the worker was not going to school, in the armed forces or retired for the

¹⁷ Troske [1994] has linked data between the LRD and the CPS, however, given the specialized nature of our samples, few matches would be found for our purposes.

¹⁸ This is thus an output-weighted mean for the industry which is appropriate for comparisons to the CPS and other representative databases.

¹⁹ Industry classifications were based on the 1980 CIC, but a "walkway" based on sex and industry was used to translate industry codes for 1979 workers to the 1980 CIC classification. Where 80% or more of a 1970 CIC class was not assigned to a single 1980 industry, the observation was dropped.

remainder of the year. This provides a measure of job turnover and also the extent to which an industry hires temporary workers. Regression 12 estimates the probability a worker worked part time last year, and regression 13, the probability a worker worked full-time all year. As can be seen, the coefficient on the implicit investment in productivity adjustment is significantly associated with more permanent matches for both production and non-production workers.

The next set of regressions measure the probability a worker has achieved a given level of education. For both production and non-production workers, the magnitude of productivity adjustment is negatively associated with the likelihood a worker has only 12 years of education or less, and is positively associated with the probability a worker has 16 years of education or more. For production workers only, productivity adjustment is positively associated with less than four years of post-secondary education.

Regressions 17 and 18 consider a worker's potential labor force experience ($\min[\text{age} - \text{highest grade} - 6, \text{age} - 17]$). For neither production workers nor non-production workers are the individual associations significant. However, the coefficient for workers with more than 25 years potential experience is significantly greater than the coefficient for workers with 10 or fewer years (at the 5% level of production workers and the 1% level for non-production workers). This suggests either a weak possible relationship between productivity adjustment and general experience or perhaps this commonly-used measure only poorly captures real experience.

The next set of regressions estimate the probability that non-production workers will fall into one of three occupational groups: managers, sales and marketing, and scientists or engineers. Only scientists and engineers are positively and significantly associated with productivity adjustment. Apparently either scientists and engineers are partially responsible for achieving the productivity gains or their activities complement the learning-by-doing activity, or (most likely) both.

This relationship is explored further in regressions 22 and 23 in Table 9. Regression 22 is the same regression, but run on the pool of all workers. Regression 23 is a similar OLS regression where the dependent variable is the average R&D intensity of the industry. The R&D intensity—R&D expenditures divided by revenues—was calculated from the NBER R&D Masterfile averaged over the years 1974 - 90. Again, both regressions show a significant relationship between R&D activity and productivity adjustment.

Regressions 24 and 25 explore the relationship between company training and productivity adjustment for production and non-production workers. Here the data are selected from the National Longitudinal Survey of Youth for years 1988 - 93 and we regress on the probability that a privately employed (excluding self-employed) manufacturing worker will have attended one or more company job training programs in any given year.²⁰ The ages of this cohort range from 23 to 35 during this period. As can be seen, the coefficient was significant only for non-production workers, although both were positive. For production workers, if the productivity adjustment results from learning-by-doing, this may be an activity that cannot always be aided by formal training.

²⁰The survey asks respondents for up to four training programs since the last interview. If any of these were company programs (formal training or seminars at work or not at work, not necessarily run by the employer), then the observation was counted as an incident of company training. We look only at the incidence of training, not the duration of each training episode; Bartels and Sicherman [1995] using the same database found most of the variation in training was over incidence.

Last, we consider productivity adjustment and the probability a worker (either production or non-production) is covered by a union contract. This analysis is based on data compiled from the Current Population Survey for 1983 - 91 by Hirsch and Macpherson [1993]. They report mean union coverage by CIC industry for each year using weighted data from the monthly CPS survey of outgoing rotation groups. We apply those means to the number of observations for each industry-year to perform a logit estimation on proportions data. The results show a positive association between productivity adjustment and union coverage.

In summary,

- The magnitude of the implicit investment in productivity adjustment at new plants is associated with a variety of indicators of firm-specific human capital, including the permanence of employment attachments, education, union coverage, job-training, and experience,
- These associations hold for both production and non-production workers, and,
- There is also an association between this investment and both R&D and the relative employment of scientists and engineers.

Analysis

Direction of causality

Thus there exist significant associations between the investment implicit in productivity adjustment and wages, wage growth and labor force characteristics. These simple associations do not, of course, indicate any specific path of causality. It is possible that causality could flow the other way, that is, changes in productivity, realized by means other than labor skill and effort, could subsequently give rise to higher wages and the given labor force characteristics.

For example, the increase in productivity might be a form of monopoly rents and workers might then be able to bargain a share of these. Or the increase in wages might be an unintended result of incentive pay that was instituted to reduce shirking.

The problem with such explanations is that it is very difficult to explain *all* of the associations found above. If wages rose as a result of bargaining or incentive pay, one might expect firms to employ *more* part-time labor rather than less. And these explanations provide no reason plants should hire more educated and experienced workers and more scientists and engineers.

Barring such a comprehensive explanation, some labor characteristic would seem to be a major cause of the productivity increases occurring at plant startups and after decreases in production employment.

Changes in composition and effort

It is also possible that the labor characteristic that produces much of the productivity gain is not an acquisition of new skills by a given body of workers. The causal factor could be a change in the composition of the workforce or a change in the effort elicited from the workforce.

For example, managers might learn workers' actual productivity over time and adjust the workforce to achieve greater productivity. Or managers could learn over time what level of effort to expect (to demand) from workers.

Although such explanations may provide a rationale for productivity growth during plant startups, these models do not explain the productivity growth following a large decrease in production employment. Presumably managers already know workers' productivity and output capability at a mature plant and any changes to the composition of the workforce or to effort benchmarks should take place *at the time of the decrease*, not after it. Indeed, average production wages rise several percent during the year of the down spike, but remain stable afterwards in our down panel; this suggests that worker quality may improve at the time of the decrease, but not afterwards.

Again, barring a model that can account for productivity gain following *both* plant initiation and employment decreases, the prime causal factor would appear to be the acquisition of skills or knowledge by the workforce. Clearly, the composition of the workforce and effort are an important part of the picture—as suggested above, learning-by-doing may require a higher quality workforce and greater worker effort—but they are not independently the main story.

Formal training and R&D

Finally, we must consider whether these skills are acquired by formal activity such as job training or R&D, or whether these skills are acquired informally via learning-by-doing. Clearly from the above results both job training and R&D are associated with productivity growth during an adjustment.

However, neither factor can explain the entire pattern of associations above. First, almost all of the above associations hold for production workers as well as non-production workers, suggesting that part of the skills are acquired by production workers. Moreover, the association between job training and the magnitude of productivity adjustment is particularly weak for production workers, suggesting that an important part of their skills are acquired informally. Second, thumbnail calculations suggest that job training would have to have an enormously higher return during these two adjustments than it apparently has in general.

Third, the timing is wrong. If formal training or R&D alone could produce the productivity gains, then why would a firm at a new plant wait until production is underway to perform these activities? A more suitable explanation is that job training and R&D account for part of the productivity gain, but that these activities are *complementary* to informal learning-by-doing. Also, perhaps, a portion of the R&D activity of scientists and engineers *is* learning-by-doing.

Thus by elimination, learning-by-doing emerges as a major causal factor (along with job training and R&D) of these large productivity adjustments that occur during plant startups and following employment decreases.

The significance of learning-by-doing human capital

Our estimates provide a basis for thumbnail calculations of the ongoing costs of adjusting to technology and employment shocks. If we assume that the total foregone output from a down spike, after correcting for the initial year observation as above, is three times the (relative) observed figure (as suggested by analysis of case study learning curves), then each down spike

costs about 44% of a year's output. This compares to an average cost of 73% of a year's output from an analysis of 40 case studies of plant startups [Bessen, 1998].²¹

Now about 4% of the plants in our base sample experienced large down spikes in a year on average. If we assume a roughly equivalent number of other shocks per year, then the total cost of adjusting to such shocks runs about 4% of total manufacturing output. Moreover, if we assume that employee turnover requires that new workers experience similar learning with similar foregone output, then the cost could be much higher.

Clearly, such adjustments are considerable. They are much larger than expenditures on formal job training. But these are more than one-time adjustment costs. The strong associations between productivity adjustment and labor characteristics indicate that firms treat this expenditure as an *investment* of firm specific capital. The knowledge and skill generated constitute a productivity-enhancing asset. And this investment is comparable in size to the net investment in physical capital.

Yet this is a “hidden” investment. It is difficult for firms to calculate in advance, it is ignored by accounting rules and never appears on a balance sheet. This may explain why firms frequently require hurdle rates of 25% or 30% on new investment projects even though their risk-adjusted cost of money is far less. Return-on-investment calculations do not usually include anticipated output foregone during ramp up, yet this is, in reality, a large part of their total investment.

This hidden investment may be particularly important for new technologies. Our calculations suggest that the costs of adopting new technologies may be much larger than their invention costs, given that basic and applied research only cost about 1% of output annually. The stock of knowledge associated with a new technology includes not only knowledge developed among scientists and engineers, but also knowledge developed by production workers.

Hence econometric estimates of the returns to R&D may be drastically overstated because they ignore a substantial quantity of learning-by-doing investment in the denominator. That is, our results show learning-by-doing investment is correlated with direct investment in R&D. The observed returns to new technology are the product of investment in R&D (capitalized R&D expenditures) and investment in learning-by-doing (capitalized foregone output) which is correlated with R&D, but omitted from the estimation. If one dollar of foregone output is needed for every dollar of R&D (and the ratio may be higher), then estimated returns to R&D are *twice* the true return on investment.²²

Thus learning-by-doing represents a large hidden investment that exerts a substantial, yet often unrecognized, influence on R&D and on firm investment and wage determination.

Conclusion

We have shown that plants implicitly make large intangible investments during plant startups and after reductions in production employment. Moreover, this investment appears to be predominately a form of firm-specific human capital, strongly associated with a variety of human

²¹ The case studies, of course, suffer from a possible upward sample selection bias.

²² This assumes they share a common depreciation rate. Also the effect will be offset somewhat because Schankerman's “double counting” bias will be even larger [1981].

capital covariates at both the plant and industry levels. The pattern of these associations suggests that this human capital results largely from learning-by-doing.

This human capital appears to differ significantly between industries, generally being greater in industries with complex technologies and/or high rates of technological change. Learning-by-doing is able to explain a number of important differences between industries.

Also, the human capital appears to be invested among production workers as well as non-production workers. This is significant for two reasons. First, it means that learning-by-doing imposes substantial adjustment costs on changes in production employment, especially negative changes. This has well-known macroeconomic implications, but a further implication is that the adjustment costs may vary with industry and that economies with more complex technologies and more rapid technological change may experience greater adjustment costs.

Second, it means that although R&D and job training for non-production workers are both positively associated with learning-by-doing, neither can fully “explain” the productivity gains achieved during learning. Put another way, this suggests that the learning activities of production workers *complement* both R&D and formal training. Thus learning-by-doing may be a substantial hidden factor in both wage determination and in estimations of the returns to R&D and to job training. Estimates of these returns may be highly overstated.

More generally, learning-by-doing may influence the rate of innovation and productivity growth. Large learning-by-doing investment implies that the cost of adopting new technologies is large. Jovanovic [1997] argues this may be critical to economic growth. To the extent that firms' managements, government policies and labor market institutions facilitate the shared investment in firm-specific learning-by-doing capital, different firms and nations will experience different rates of innovation.

Appendix

A. Measurement of total factor productivity quantities

Output was measured as the sum of shipments plus the change in inventories deflated by the appropriate 4-digit SIC deflator from the NBER Productivity database developed by Wayne Gray and Eric Bartelsman [1996]. The input factors used were production hours, non-production workers, the deflated sum of materials, purchased services and energy, and capital.

For the startup sample, net real capital was constructed using a perpetual inventory method where initial gross capital and investments in plant and equipment were deflated by 4-digit deflators developed by Bartelsman and Gray. Since capital is new, any depreciation series based on stochastic retirements is inappropriate. Additionally, we found relatively high levels of capital retirements during the initial years, tapering off to more normal levels by the fifth year—44% of the plants experienced retirements greater than 10% of their final capital stock! Although retirements are usually seen as the result of declining useful service efficiency, a high level of retirements is, in fact, consistent with learning-by-doing—initial capital purchases, made when managers have the least knowledge of operations, are frequently inappropriate or inefficient and thus often retired early. Accordingly, capital stocks were assumed to decay according to a beta function representing the efficiency of their services and retirements were accounted for

explicitly.²³ For productivity calculations, capital stocks also included deflated inventories and rental expenditures for equipment and structures divided by an industry rental rate (see below).

Output shares for production and non-production labor were calculated as 1.17 times the respective wage bill over current-dollar output to account for mandatory employer payroll taxes and benefits.²⁴ The output share for materials, purchased services and energy was simply the ratio of expenditures on these items to current-dollar output. Capital output shares were calculated by applying a capital rental rate to the current stocks. Rental rates were calculated using the 2-digit BLS series for current dollar capital cost divided by the total capital stocks for the corresponding industry from the NBER Productivity database. Note that this approach does not impose an assumption of constant returns to scale.

The productivity calculations for the Up and Down files were similar except that capital was constructed differently, as here stochastic retirements incorporated in the depreciation measure are appropriate. Capital was calculated on a perpetual inventory basis using 2-digit BEA deflators and 2-digit BEA depreciation series. The details used are the same as those described in Adams and Jaffe [1994]. Rentals were not included and retirements were not included explicitly, but are included as a stochastic term in the BEA depreciation figures. Other aspects of the productivity calculation remained the same.

B. Correlating industry productivity adjustment measures across panels

We wish to test for the correlation of the industry means of productivity adjustment measures between the startup sample and the down spike sample. Because of the differences between the samples, this must be performed at the 4-digit SIC level, otherwise differences in the composition of aggregate industries come into play. Many of these industries have very small populations, however, so sampling variance in both samples is large. The latent model is

$$y_i^* = \mathbf{b}_0 + \mathbf{b}_1 \cdot x_i^* + \mathbf{e}_i, \quad \mathbf{e} \sim N(0, \mathbf{S}^2) \quad (1)$$

where x^* and y^* are the true (unobserved) mean industry measures for the startup panel and the down spike panel respectively. Because of sampling variance, however, the actual observed means are

²³We used the BLS beta function where capital services at time t are $K(t) = K_0 \cdot (L - t) / (L - \mathbf{b} \cdot t)$ where K_0 is the initial value of the capital and L is the useful service life, and $\mathbf{b} = 0.5$ for equipment and 0.75 for structures. Industry service lives were taken from the BEA series of service lives by asset type and the asset type mix for each industry was derived from the 1977 Input-Output tables. Note that the LRD only reports retirement data for 1977 - 87. Prior to 1977 retirements were imputed as $\min(0, I - \Delta G)$ where I is investment and ΔG is the change in gross book value capital.

²⁴The figure 1.17 was calculated as the mean ratio of employment compensation to wages and salaries in the NIPA series for manufacturing workers 1972 - 87.

$$\begin{aligned}
 x_i &= x_i^* + u_i, & u_i &\sim N\left(0, \frac{\mathbf{s}_I^2}{n_i^x}\right) \\
 y_i &= y_i^* + v_i, & v_i &\sim N\left(0, \frac{\mathbf{s}_I^2}{n_i^y}\right)
 \end{aligned}
 \tag{2}$$

where σ_I^2 is the within-industry variance due to idiosyncratic plant effects (assumed to be uniform across industries) and n_i^x and n_i^y are the number of plant observations for the i th industry for the startup and down spike panels respectively (dropping observations where either are zero). Plugging (2) into (1) we get

$$y_i = \mathbf{b}_0 + \mathbf{b}_1 \cdot x_i + w_i, \quad w_i \sim N\left(0, \mathbf{s}^2 + \mathbf{s}_I^2 \left(\frac{1}{n_i^y} + \frac{\mathbf{b}_1^2}{n_i^x}\right)\right)
 \tag{3}$$

This cannot be estimated with WLS because the error is correlated with x (standard measurement error problem). Plugging in the normal distribution, the maximum likelihood function is

$$\ln L = -\frac{1}{2} \cdot \sum_i \left[\ln \left(\mathbf{s}^2 + \mathbf{s}_I^2 \left(\frac{1}{n_i^y} + \frac{\mathbf{b}_1^2}{n_i^x} \right) \right) + \frac{(y_i - \mathbf{b}_0 - \mathbf{b}_1 \cdot x_i)^2}{\mathbf{s}^2 + \mathbf{s}_I^2 \left(\frac{1}{n_i^y} + \frac{\mathbf{b}_1^2}{n_i^x} \right)} \right].
 \tag{4}$$

Using this function, (3) was estimated for two productivity adjustment measures: the mean industry foregone output (relative) divided by mean real output, and the mean industry foregone output (relative) divided by mean total employment. There were 368 industries with plants in both panels. Equation (4) was maximized and the likelihood value was compared to a similar likelihood function with the additional restriction that $\beta_1 = 0$. For both measures the improvement in the likelihood of the unrestricted model was highly significant using the likelihood ratio test (see table) rejecting the hypothesis of no correlation.

Values of $\ln L$

	$\frac{\Delta Q^R}{Q}$	$\frac{\Delta Q^R}{EMP}$
Unrestricted model	-2146.95	-4374.29
Restricted model	-2192.44	-4419.07

Bibliography

- ADAMS, J. AND JAFFE, A., 1994. "The Span of the Effect of R&D in the Firm and Industry", *Center for Economic Studies Working Paper*, no. 94-7.
- ARGOTE, L. AND EPPLE, D., 1990. "Learning Curves in Manufacturing", *Science*, p. 920.
- ARGOTE, L., BECKMAN, S. AND EPPLE, D., 1990. "The Persistence and Transfer of Learning in Industrial Settings", *Management Science*, p. 140.
- ARROW, K., 1962. "The Economic Implications of Learning by Doing", *Review of Economic Studies*, p. 155.
- AUERSWALD, P., KAUFFMAN, S., LOBO, J., AND SHELL, K., 1996. "A microeconomic theory of learning-by-doing: an application of the nascent technology approach," unpublished.
- BAHK, B. AND GORT, M., 1993. "Decomposing Learning by Doing in New Plants", *Journal of Political Economy*, p. 561.
- BAILY, M., BARTELSMAN, E., AND HALTIWANGER, J., 1996. "Labor Productivity: structural change and cyclical dynamics", *NBER Working Paper Series*, no. 5503.
- BAILY, M., HULTEN, C., AND CAMPBELL, D., 1992. "Productivity Dynamics in Manufacturing Plants", *Brookings Papers: Microeconomics*, p. 187.
- BALOFF, N., 1966. "Startups in Machine-intensive Production Systems", *Journal of Industrial Engineering*, p. 25.
- BARTEL, A. AND LICHTENBERG, F., 1987. "The comparative advantage of educated workers in implementing new technology", *The Review of Economics and Statistics*, v. 69, no. 1 p. 1.
- BARTEL, A. AND SICHERMAN, N., 1995. "Technological Change and the Skill Acquisition of Young Workers", *NBER Working Paper Series*, no. 5107.
- BARTELSMAN, E. AND GRAY, W., 1996. "The NBER Manufacturing Productivity Database", *NBER Technical Working Paper*, no. 205.
- BECKER, G., 1965. "Human Capital", *University of Chicago Press*.
- BESSEN, J., 1998. "Discovery, Learning and Adoption of New Techniques: Choosing specialization to optimize technical progress", *Research on Innovation Working Paper*.
- BRESNAHAN, T. AND RAMEY, V., 1994. "Output Fluctuations at the Plant Level", *Quarterly Journal of Economics*, p. 593.
- CARMICHAEL, L., 1983. "Firm Specific Human Capital and Promotion Ladders", *Bell Journal*, v. 14 p. 251.
- CHRISTENSEN, L., CUMMINGS, D., AND JORGENSEN, D., 1981. "Relative Productivity Levels, 1947-1973, An international comparison", *European Economic Review*, v. 16 p. 61.
- COHEN, W. AND LEVINTHAL, D., 1989. "Innovation and Learning: the two faces of R&D", *Economic Journal*, v.99.
- DAVID, P., 1975. *Technical Choice Innovation and Economic Growth: Essays on American and British experience in the nineteenth century*. Cambridge University Press.
- DAVIS, S., HALTIWANGER, J., AND SCHUH, S., 1996. *Job Creation and Destruction*, MIT Press.
- DUTTON, J AND THOMAS, A., 1984. "Treating Progress Functions as a Managerial Opportunity", *Academy of Management Review*, v. 9, no. 2 p. 235.

- HAMERMESH, D. AND PFANN, G., 1996. "Adjustment Costs in Factor Demand", *Journal of Economic Literature*, v. **34** p. 1264.
- HASHIMOTO, M., 1981. "Firm-Specific Human Capital as a Shared Investment", *American Economic Review*, v. **71**, no. **3** p. 475.
- HIRSCH, B. AND MACPHERSON, D., 1993. "Union Membership and Coverage Files from the Current Population Surveys: Note", *Industrial and Labor Relations Review*, v. **46**, no. **3** p. 574.
- HUTCHENS, R., 1989. "Seniority, Wages and Productivity: A turbulent decade", *Journal of Economic Perspectives*, v. **3** p. 49.
- JARMIN, R., 1996. "Learning-by-doing and plant characteristics", *Center for Economic Studies Working Paper*, no. **96-5**.
- JOVANOVIC, B. AND NYARKO, Y., 1995. "A Bayesian learning model fitted to a variety of empirical learning curves," *Brookings Papers: Microeconomics*, p. 247.
- JOVANOVIC, B. 1997. "Learning and Growth," in Kreps and Wallis, ed., *Advances in Economics and Econometrics: theory and applications*, v. **II**, Econometric Society Monograph p. 318.
- LAZEAR, E. , 1979. "Why is There Mandatory Retirement?", *Journal of Political Economy*, v. **87** p. 1261.
- LAZONICK, W. AND BRUSH, T., 1985. "The 'Horndal Effect' in Early U.S. Manufacturing", *Explorations in Economic History*, v. **22** p. 53.
- MCGUCKIN, R. AND PASCOE, G., 1988. "The Longitudinal Research Database and Research Possibilities", *Survey of Current Business*, v. **68**, no. **1** p. 30.
- MUTH, J. 1986. "Search Theory and the Manufacturing Progress Function," *Management Science*, p. 948.
- RAPPING, L., 1965. "Learning and World War II Production Functions", *Review of Economics and Statistics*, v. **47** p. 81.
- SCHANKERMAN, M., 1981. "The Effects of Double Counting and Expensing on the Measured Returns to R&D", *Review of Economics and Statistics*, v. **63** p. 454.
- SHAKED, A. AND SUTTON, J., 1984. "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model", *Econometrica*, v. **52**, no. **6** p. 1351.
- TROSKE, K. 1994. "Evidence on the Employer Size-Wage Premium From Worker-Establishment Matched Data", *Center for Economic Studies Working Paper*, no. **94-10**.
- WEISS, A., 1995. "Human Capital vs. Signaling Explanations of Wages", *Journal of Economic Perspectives*, v. **9** p. 133.
- WRIGHT, T., 1936. *Journal of Aeronautical Science*, p. 122.
- ZEVIN, R., 1975. *The Growth of Manufacturing in Early Nineteenth Century New England*, New York: Arno.

Table 1. Sample characteristics

	5 Year Startup Panel	Census Startups	Up Spike File	Down Spike File
Number of plants	5,625	7,668	5,600	3,853
Mean employees year 1	63.1	54.9	623.1	571.8
Mean production workers year 1	47.7	42.6	444.4	377.4
Mean output year 1 (thousand 1987 \$)	6,180	4,777	90,114	72,579
Mean net capital stock year 1 (thousand 1987 \$)	3,441	2,531	30,062	39,332
Percent single-plant firms	39 %	49%	4%	3%

Table 2. Attenuation of productivity adjustment measures arising from data limitations

	Δ	$\Delta Q/Q$
Monthly data	123%	16%
Annual data, none missing	19%	17%
Annual data, 3 months missing	3%	3%
Annual data, 6 months missing	1%	1%
Annual data, 9 months missing	0%	0%

Calculations based on data from Baloff [1966] used in Figure 1.

Table 3. Productivity adjustment measures on startup plants

SIC	Industry	Δ_{15}	Δ_{15}^R	ΔQ	ΔQ^R	$\frac{\Delta Q}{CAP}$	$\frac{\Delta Q^R}{CAP}$	$\frac{\Delta Q^R}{EMP}$
	ALL MANUFACTURING (excluding tobacco)	.057 (.005)	.031 (.005)	1088	353	37%	12%	6.2
20	Food and kindred products	.023 (.013)	.003 (.013)	443	93	14%	3%	1.8
22	Textile mill products	.083 (.027)	.056 (.025)	1608	172	33%	4%	1.8
23	Apparel	.064 (.027)	.007 (.026)	155	-159	22%	-22%	-2.1
24	Lumber and wood products	.052 (.016)	.041 (.016)	224	85	12%	5%	1.8
25	Furniture and fixtures	.071 (.023)	.048 (.023)	39	-217	3%	-18%	-4.0
26	Paper and allied products	.095 (.032)	.088 (.031)	1332	1238	24%	22%	29.3
27	Printing and publishing	.109 (.020)	.111 (.020)	208	163	21%	16%	4.9
28	Chemicals and allied products	.066 (.026)	.027 (.025)	865	628	17%	12%	18.1
29	Petroleum and coal products	.007 (.036)	-.004 (.032)	-637	-983	-32%	-49%	-65.6
30	Rubber and plastic products	.138 (.018)	.103 (.018)	755	247	26%	8%	4.3
31	Leather and leather products	-.012 (.048)	-.026 (.041)	899	333	51%	19%	3.7
32	Stone, clay and glass products	.026 (.020)	.000 (.020)	448	160	18%	6%	4.7
33	Primary metal industries	.013 (.028)	.001 (.027)	2117	2577	31%	38%	38.0
34	Fabricated metal industries	.024 (.015)	.033 (.015)	70	95	3%	4%	2.0
35	Machinery except electrical	.072 (.019)	.020 (.017)	4170	207	120%	6%	2.9
36	Electric and electronic equipment	.066 (.029)	-.004 (.029)	3039	1369	66%	30%	13.3
37	Transportation equipment	.036 (.020)	.011 (.019)	2546	1617	59%	37%	21.8
38	Instruments and related products	.040 (.035)	-.039 (.035)	3002	1470	99%	48%	18.4
39	Miscellaneous manufacturing	.039 (.036)	.034 (.034)	397	30	26%	2%	0.5

Standard error of mean in parentheses.

ΔQ in thousands of 1987 dollars, ratios calculated as ratios of industry averages.

Table 4. Regressions on absolute TFP change, Δ_{15}

Regression no.	1	2	Variable means (std. dev.)
Intercept	0.104* (0.027)		
Production worker / total emp., year 1	0.042 (0.028)	0.020 (0.031)	.79 (.19)
10 < Employees <= 50, year 1	-0.099* (0.014)	-0.102* (0.014)	.45
50 < Employees <= 200, year 1	-0.136* (0.015)	-0.146* (0.016)	.28
200 < Employees, year 1	-0.127* (0.027)	-0.139* (0.028)	.05
Capital / prod. worker, year 1 (million 1987 dollars)	0.056* (0.019)	0.070* (0.019)	.148 (1.195)
Industry average TFP growth, $\bar{\Delta}_{15}^I$	0.769* (0.066)	0.620* (0.121)	.038 (.081)
Industry TFP growth deviation, δ_{15}^I	1.007* (0.062)	1.049* (0.064)	-.012 (.085)
Single plant firm	-0.031* (0.011)	-0.043* (0.012)	.49
CIC industry dummies		√	
Adjusted R-squared	.084	.106	

*Significant at 1% level.

Standard errors in parentheses.

Table 5. Measures of productivity adjustment following different events.

Industry	Δ_{15}	Δ_{15}^R	ΔQ	ΔQ^R	$\frac{\Delta Q}{CAP}$	$\frac{\Delta Q^R}{CAP}$	$\frac{\Delta Q^R}{EMP}$
Startup	.057 (.005)	.031 (.005)	1088	353	37%	12%	6.2
Up employment spike	.030 (.005)	-.006 (.005)	8917	1816	35%	7%	3.3
Down employment spike	.117 (.006)	.065 (.006)	17969	9490	58%	30%	18.5

Standard errors in parentheses.

ΔQ in thousands of 1987 dollars, means for ΔQ for up and down files are shown trimmed of 1% tails, ratios calculated as ratios of industry averages.

Table 6. Regressions on Relative Wages

Regression no.	3	4	5	6	Variable means (std. dev.)
	Production Workers		Non-production Workers		
Intercept	.053 (.030)		-.326* (.050)		
Absolute productivity change, Δ_{15}	.162* (.015)	.162* (.014)	.282* (.026)	.287* (.026)	.060 (.404)
Relative TFP, year 1	.218* (.018)	.184* (.018)	.312* (.032)	.314* (.032)	-.083 (.322)
Industry average TFP growth, $\bar{\Delta}_{15}^I$	-.366* (.066)	-.192 (.116)	-.245 (.114)	-.115 (.205)	.038 (.081)
Industry TFP growth deviation, δ_{15}^I	-.245* (.065)	-.158 (.063)	-.265 (.110)	-.165 (.112)	-.012 (.085)
Production worker / total emp., year 1	-.229* (.029)	-.112* (.029)	.001 (.050)	.019 (.053)	.79 (.19)
10 < Employees <= 50, year 1	-.004 (.014)	.044* (.014)	.218* (.025)	.229* (.025)	.45
50 < Employees <= 200, year 1	-.014 (.016)	.080* (.016)	.237* (.027)	.275* (.029)	.28
200 < Employees, year 1	.054 (.026)	.138* (.026)	.294* (.047)	.318* (.048)	.05
Capital / prod. worker, year 1 (million 1987 dollars)	.173* (.019)	.152* (.018)	.108* (.032)	.084 (.033)	.148 (1.195)
Single plant firm	-.095* (.012)	-.085* (.012)	-.043 (.020)	-.014 (.020)	.49
Region dummies	√	√	√	√	
CIC industry dummies		√		√	
Adjusted R-squared	.116	.424	.054	.156	

*Significant at 1% level.

Standard errors in parentheses.

Table 7. Regression on change in wages

Regression no.	7	8	9	10	Variable means (std. dev.)
	Production Workers		Non-production Workers		
Intercept	-.148* (.031)		.974* (.059)		
Δ_{15}	.286* (.015)	.285* (.015)	.492* (.031)	.480* (.032)	.060 (.404)
Industry average TFP growth, $\bar{\Delta}_{15}^I$	-.008 (.076)	.121 (.138)	-1.047* (.143)	-.877* (.267)	.038 (.081)
Industry TFP growth deviation, δ_{15}^I	-.289* (.072)	-.273* (.074)	-.546* (.141)	-.482* (.146)	-.012 (.085)
Production worker / total emp., year 1	.075 (.032)	.102* (.035)	-1.171* (.064)	-1.350* (.071)	.79 (.19)
10 < Employees <= 50, year 1	.060* (.016)	.069* (.016)	-.169* (.034)	-.188* (.034)	.45
50 < Employees <= 200, year 1	.086* (.018)	.097* (.019)	-.177* (.037)	-.212* (.038)	.28
200 < Employees, year 1	.076 (.031)	.081 (.032)	-.058 (.060)	-.089 (.062)	.05
Capital / prod. worker, year 1 (million 1987 dollars)	.039 (.022)	.039 (.022)	-.152* (.041)	-.140* (.042)	.148 (1.195)
Single plant firm	-.010 (.013)	-.007 (.013)	.038 (.025)	.027 (.026)	.49
Region dummies		√		√	
CIC industry dummies		√		√	
Adjusted R-squared	.063	.074	.119	.131	

*Significant at 1% level.

Standard errors in parentheses.

Table 8. Logistic regressions on industry characteristics

Regression no.	11	12	13	14	15	16	17	18	19	20	21
Characteristics of Industry	Permanence of employment			Education			General experience		Occupation (non-production)		
	Changed jobs / temp. work	Part time worker	Full time year round worker	High school or less	Some post-secondary schooling	College or more	Years in labor force <= 10	Years in labor force > 25	Manager	Marketing / sales occupation	Scientist / engineer
Production workers ($N = 18,178$)											
ΔQ^R per employee <i>marginal effect</i> [†]	-.0020* (.0007) -1.0 %	-.0050* (.0016) -0.7 %	.0025* (.0006) 1.5 %	-.0033* (.0007) -1.2 %	.0028* (.0008) 0.9 %	.0042* (.0014) 0.3 %	-.0008 (.0006) -0.5 %	.0012 (.0006) 0.7 %			
TFP 4-year growth rate, $\bar{\Delta}_{15}^I$.7392 (.3073)	-2.5086* (.7623)	-.2058 (.2860)	-1.0745* (.3135)	.8528 (.3453)	1.2910 (.5612)	-.0118 (.2822)	-.5616 (.3039)			
Capital per employee	-.0038* (.0004)	-.0071* (.0010)	.0044* (.0004)	-.0042* (.0003)	.0039* (.0004)	.0037* (.0006)	-.0016* (.0003)	.0007 (.0003)			
Annual employment growth rate	-.0541* (.0083)	.0605* (.0169)	.0301* (.0077)	-.0990* (.0099)	.0881* (.0108)	.1089* (.0202)	.0514* (.0078)	-.0394* (.0076)			
Constant	-.9811* (.0310)	-2.4337* (.0666)	.4835* (.0283)	1.8690* (.0336)	-2.1167* (.0368)	-3.5919* (.0681)	-.6298* (.0278)	-.7784* (.0273)			
Non-production workers ($N = 9,905$)											
ΔQ^R per employee <i>marginal effect</i> [†]	-.0035* (.0012) -1.0 %	-.0082* (.0015) -1.9 %	.0043* (.0009) 2.1 %	-.0047* (.0007) -3.2 %	.0002 (.0008) 0.1 %	.0050* (.0008) 3.0 %	-.0019 (.0008) -1.2 %	.0009 (.0008) 0.5 %	-.0002 (.0009) -0.1 %	-.0026 (.0012) -0.7 %	.0056* (.0010) 1.8 %
TFP 4-year growth rate, $\bar{\Delta}_{15}^I$.0023 (.3350)	-7.3239* (.7668)	2.1824* (.2899)	-1.6907* (.2320)	.7027* (.2323)	1.0015* (.2121)	-.3921 (.2163)	-1.0066* (.2514)	.5062 (.2576)	-3.3040* (.4588)	1.9310* (.2544)
Capital per employee	-.0032* (.0007)	-.0064* (.0010)	.0045* (.0005)	-.0034* (.0004)	-.0008 (.0004)	.0043* (.0004)	-.0014* (.0004)	.0002 (.0004)	.0005 (.0004)	-.0018* (.0007)	.0037* (.0005)
Annual employment growth rate	-.0276 (.0154)	.1375* (.0183)	-.0397* (.0118)	-.0769* (.0098)	-.0147 (.0112)	.1051* (.0108)	.0669* (.0104)	-.0446* (.0103)	-.0312* (.0120)	.0965* (.0163)	.0779* (.0152)
Constant	-1.8790* (.0558)	-1.7103* (.0680)	.8282* (.0426)	.0711 (.0350)	-1.1177* (.0402)	-1.1420* (.0382)	-.5212* (.0366)	-.8068* (.0371)	-1.4658* (.0429)	-1.8801* (.0571)	-2.3227* (.0541)

*Significant at 1% level. Asymptotic standard errors in parentheses.

[†]Marginal effect shown is the mean effect of a standard deviation change in the independent variable evaluated at all data points.

Table 9. Additional regressions on industry characteristics


Regression no.	22	23	24	25	26
	Scientist / engineer	OLS: R&D as % of sales	Received company training	Received company training	Covered by union contract
Group:	All		Production worker	Non-production worker	All
ΔQ^R per employee <i>marginal effect</i> [†]	.0082* (.0010) 0.9 %	.0148* (.0055) 0.5 %	.0035 (.0020) 0.6 %	.0054* (.0018) 2.2 %	.0023* (.0001) 1.3 %
Capital per employee	.0057* (.0005)	.0028 (.0028)	.0055* (.0009)	.0037* (.0008)	.0028* (.0001)
TFP 4-year growth rate, $\bar{\Delta}_{15}^I$	2.6199* (.2488)	10.1879* (3.0221)	2.0689* (.6567)	1.5096* (.4302)	-3.8289* (.0819)
Annual employment growth rate	.1986* (.0151)	.1117 (.0871)	.0957* (.0285)	.0464 (.0253)	-.0636* (.0019)
Constant	-3.6244* (.0518)	1.0696* (.3167)	-3.0306* (.0961)	-1.9694* (.0885)	
Year dummies					√
No. of observations	28,083	58 $\bar{R}^2 = .25$	5,040	2,864	--

*Significant at 1% level.

Standard errors in parentheses.

[†]Marginal effect shown is the mean effect of a standard deviation change in the independent variable evaluated at all data points.

Figure 1. Monthly productivity of steel process startup with productivity adjustment measures



Source: Baloff [1966]

Figure 2. Mean Year-to-year Productivity Changes after Startup




Figure 3. Mean year-to-year changes in relative productivity after up and down employment spikes

