How Much Did The Liberty Shipbuilders Forget?

Peter Thompson

Florida International University

July 2006

This paper produces new estimates of the rate of organizational forgetting in the well-known case study of US wartime ship production. Estimates obtained using data constructed from primary sources at the National Archives yield rates of forgetting that are much smaller than previously reported, and may well be zero. The richness of the data make it possible to control for variations in the product mix, to explore alternative formulations for the learning curve, and to investigate the relationship between organizational forgetting and labor turnover.

1. Introduction

Seventy years have passed since Wright's (1936) pioneering study, yet in all that time the empirical study of organizational learning curves has experienced remarkably little technical change. Wright observed that the unit labor requirement in airframe manufacturing declined at a constant rate with each doubling of cumulative output. This phenomenon, consistent with a loglinear relationship between productivity and cumulative output, has subsequently been found to apply in numerous settings, although reported rates of learning vary widely across industries and across firms within industries (Dutton and Thomas 1994, Yelle 1979).

Among the few noteworthy innovations in the study of learning curves has been the introduction of the notion of organizational forgetting. Several researchers (Argote, Beckman and Epple 1990, Argote, at al. 1997, Benkard 2000, Darr, Argote and Epple 1995, Epple, Argote and Devadas 1991, Epple, Argote and Murphy 1995) have noted that observed costs may actually increase during certain periods of a product lifecycle. In a neat appeal to symmetry with the literal interpretation of the learning curve, these researchers have argued that reversals in productivity can be attributed to organizational forgetting.[1]

There is considerable evidence that interruptions to production may be associated with organizational knowledge loss (e.g., Hirsch 1956, Baloff 1970). But the more recent studies have made the rather stronger claim that organizational forgetting occurs even under conditions of continuous production.[2] The evidence, drawn from a diverse set of industries, suggests that knowledge depreciation can be economically significant, although it varies widely across cases. Among pizza franchises, for example, Darr, Argote and Epple found that knowledge depreciates at the astonishing rate of 17 percent per week, implying that "roughly one half of the stock of knowledge at the beginning of the month would remain at the end of the month." In wartime construction of Liberty cargo vessels, Argote, Beckman and Epple report that knowledge depreciated at the rate of 25

---

1. There is a larger, and older, literature on individual forgetting (e.g., Ebbinghaus 1885, Finkenbinder 1913, Thorndike 1914, McGeoch 1932). Some papers in this area (e.g., Nembhard 2001, Nembhard and Uzumeri 2000a, 2000b) have studied individual forgetting in commercial organizations.

2. A parallel claim does not exist for individual forgetting. In fact, McGeoch (1932:357) wondered "if there are cases of forgetting which occur during use." The only examples he could come up with were a vague allusion to forgetting "sexual perversions" and typists forgetting to make errors (i.e. learning).

percent per month. Benkard's (2000) study of aircraft manufacturing by Lockheed generated an annual rate of depreciation of about 40 percent, or about 3 percent per month.

Although numerous explanations have been given for why organizations appear to forget in the face of interruptions to production (e.g., Anderlohr 1969), few have been given for continuous depreciation of knowledge. One – that technological change makes past experience increasingly irrelevant – is perhaps a misnomer.[3] A second explanation is that organizations often fail to record experiences because of inadequately-designed organizational memory systems (Landry 1999). A third is that tacit knowledge embodied in employees is lost to labor turnover. Only the third explanation has been subject to direct testing, and the evidence is a little thin. In wartime shipbuilding, Argote, Beckman and Epple (1990) found that labor turnover rates averaging 10 percent per month did not appear to affect productivity. Argote, at al. (1997) find a u-shaped relationship between productivity and turnover in an American truck plant. However, Argote (1999) has noted that the rank order of knowledge depreciation rates in several studies matches the rank order of labor turnover rates.

Having accepted the evidence along with the few explanations on offer, many authors have drawn attention to the consequences of organizational forgetting for firm profitability and some have proposed strategies to help firms retain their hard-won knowledge (e.g., Belason 2000, Cross and Baird 2000, Kransdorrf 1997). Perhaps just to be contrary, Peters (1999) has gone so far as to argue that forgetting is more valuable than learning and has proposed methods by which firms can increase the amount they forget. More considered arguments along the same lines can be found in Huber (1991), and Walsh and Ungson (1991). Yet other research has shown how forgetting undermines attempts to institute flexible production schedules and just-in-time manufacturing processes (Smunt 1987). Interest is now turning toward the broader consequences of organizational forgetting. Benkard (2000), for example, has called for new theoretical efforts to explain its strategic implications.

In view of this broad interest, this paper offers another look at a familiar case study, the Liberty shipbuilding program of World War II. The episode has long been a classic case study of learning (Searle 1945, Rapping 1965, Lucas 1993, Thompson 2001, Thornton and Thompson 2001) and it is also a seminal case study of forgetting (Argote, Beckman and Epple 1990). The paper exploits recently discovered data on unit labor requirements (ULRs), collected by the author from primary

---

3. Productivity losses associated with changes in technology are usually interpreted as incomplete knowledge spillovers (e.g. Irwin and Klenow 1994).

sources at the National Archives, to produce new estimates of the rate of forgetting. The data were originally collected by auditors because cost-plus contracts paid out on individual ships demanded accurate and contemporaneous records of labor utilization at the individual product level. The quality of the data in this study is therefore comparable to those used in Benkard's (2000) analysis of aircraft manufacturing.

The unusually rich data make it possible to address two possible pitfalls in estimating forgetting rates that I think have not received sufficient attention to date. The first of these concerns problems of aggregation bias induced by unobserved changes in the product mix that can complicate inference about learning and forgetting from more aggregate data than those used here. The second concerns the relationship that can exist between assumptions about the nature of learning and inferences about the rate of forgetting. In particular, if learning is bounded, the dominant loglinear specification can produce high imputed rates of forgetting even when none exists. These issues are discussed in Sections 2 and 3.

The main empirical results are reported in Section 4. A loglinear specification for learning with no attempt made to control for the product mix returns a rate of forgetting of 8.4 percent per month, about one-third of the rate reported in Argote, Beckman and Epple (1990). Accounting for product-mix changes reduces the estimate by half, to 4.2 percent. Somewhat surprisingly, replacing the loglinear specification with either of two bounded learning models has little effect on the results. Absent controls for changes in the product mix, the bounded learning models return rates of forgetting of 8.7 and 5.8 percent per month. In both cases, accounting for the product mix reduces these estimates significantly, to 5.7 percent and 3.6 percent respectively, rates very close to those reported in Benkard (2000). The addition of controls for labor turnover also yields surprising results. When added in the usual way as a level effect, turnover rates are found to be *positively* correlated with productivity, while the estimated monthly rate of forgetting declines further to between -1.4 and $+2.0$ percent. An alternative specification, in which the rate of forgetting is a function of labor turnover rather than time also produces no evidence of organizational forgetting.

Section 5 concludes. In this case, changes in the specification of the learning curve had little effect on inferences about forgetting, although there is no reason to expect this finding to hold in other applications. In contrast, adequate controls for the product mix appeared to matter, and using them yielded modest rates of organizational forgetting, at least relative to most previous findings. Attempts to relate forgetting to labor turnover were unsuccessful. To the contrary, the inclusion of labor turnover data eliminates entirely any evidence of organizational forgetting. Although

these findings are robust across three specifications for learning, the concluding section includes a caveat. A major motivation for incorporating organizational forgetting into our learning models is to improve our ability to track the data. While it is certainly true that the incorporation of forgetting achieves this aim, it turns out that a wide range of rates of forgetting does almost as well as any other in terms of model fit.

2. Aggregation and the Product Mix

A challenge in measuring learning and forgetting rates with firm- or plant-level data is that at this level of aggregation inferences from output levels can easily confound within–product variations in productivity with changes in a firm's product mix.[4] Even minor product changes, unobservable to the econometrician, can induce significant declines in measured productivity as a firm tries to incorporate the changes into its production system. It is also typical for line speed to decline as the firm tools up to accommodate the new design. It then follows that design changes induce a negative correlation between line speed and productivity, and this can lead to an inference of organizational forgetting in the standard formulation.

The consequences of changes in the product mix are well-illustrated by two episodes from the Liberty ship program. Both episodes are more easily understood with some background on historical data. The standard data source for the Liberty ship program is Fischer (1949), which turns out to have some important limitations. To produce yard output data from Fischer's tabulations, it is necessary to combine a table listing, for each yard, monthly output of Liberty ships per way (building berth) with a table summarizing the number of ways in the yard. Labor inputs are obtained as the product of another two tables, one giving the average number of employees at the yard in each month, the other the average hours worked per week.[5] These are the data used

---

4. Of course, this is not a new concern. For example, Smunt (1987), and McCreery and Krajewski (1999), have used simulation experiments to draw attention to the effects of changing product mixes on learning performance.

5. There are also some minor complications. For example, the average hours worked refers to wage employees and not salaried employees; the number of workers is from a survey taken on the 15[th] of each month for most of the war but on the 30[th] of each month for some observations; the number of ways refers to the number authorized, not the number fully equipped and in operation. For an extensive treatment of capital measurement problems in wartime shipbuilding, see Thompson (2001).

in Rapping (1965) and Argote, Beckman and Epple (1990), although Rapping aggregated to annual data.

Output per unit of labor inferred from these data mislead in two ways. First, the output data include modifications to the standard dry cargo vessel – tankers, colliers, aircraft transporters, hospital ships, passenger ships, and navy training vessels – which induced significant shocks to measured labor productivity. Each of these shocks associates a temporary period of low productivity with low line speed. Second, the input data include labor used in the production of some vessels that are excluded from the output data. Toward the end of the war, many of the yards gradually switched production to new vessel types, most commonly the new Victory cargo ship. Liberty ship production did not entirely stop in a yard until several months after it had begun to divert labor to the newer vessels. Although the labor input data record employees working on all types of vessels, including new vessels for which labor productivity is very low, output data derived from Fischer (1949) assign all this labor to Liberty ships. The net effect is to give the appearance of a dramatic decline in productivity as the rate of Liberty ship production declined.

Figure 1, which details production at the Delta Shipbuilding yard between January 1942 and September 1945, illustrates the productivity consequences of product design changes. The yard built 185 vessels during the war, all of them Liberty ships. But of these, 54 were heavily modified designs produced in two distinct clusters. In the spring of 1943, the yard won a contract to produce 32 oil tankers. To satisfy the contract, the production schedule for standard Liberty ships already under contract was suspended on 18 May 1943, and the yard tooled up to produce the tankers. Between 18 May and 6 July 1943, when the last standard Liberty that had already been under construction in May was delivered, the yard was simultaneously producing two types of vessels. Construction of the standard Liberty recommenced gradually after 25 November 1943, and until 10 February 1944, when the last tanker was delivered, the yard was again distracted by the construction of multiple designs. In October 1944, the yard won a contract for 24 colliers. This time, the yard completed its outstanding contract for standard designs, progressively turning the yard over to collier construction as capacity became available. The completion of the collier contract ended the yard's involvement in the shipbuilding program. However, between 15 September 1944 and 4 May 1945, the yard was again constructing two designs.

The two periods in which the yard engaged in production of the modified designs are associated with a sharp increase in unit labor requirements (upper panel of Figure 1) along with a sharp
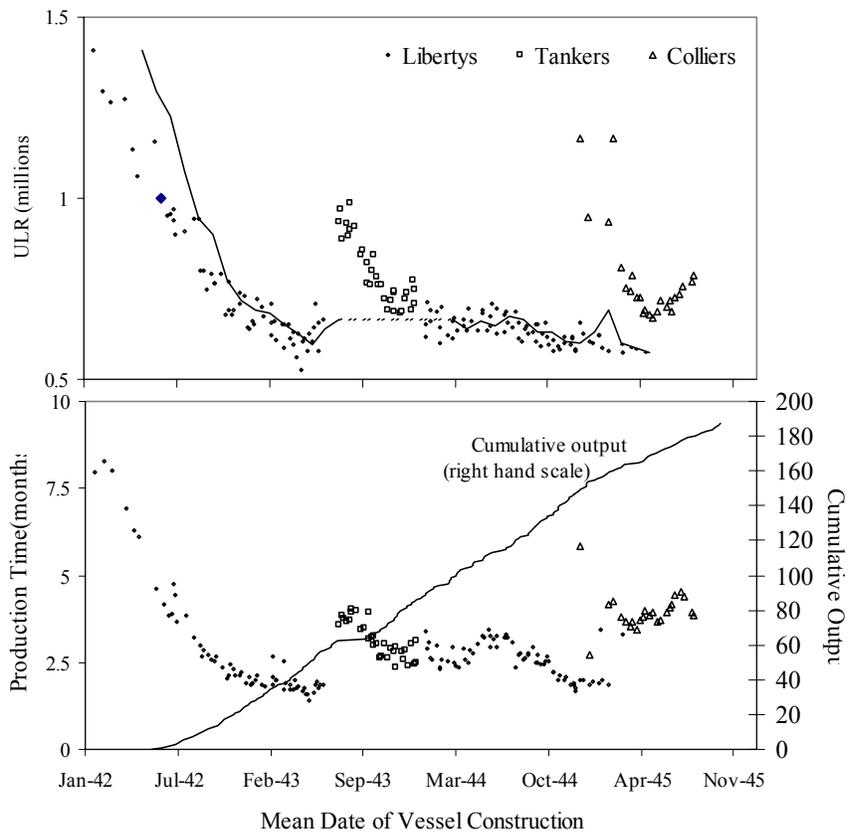
5

FIGURE 1. Delta Shipbuilding Corporation: unit labor requirements and production speed. Upper panel records unit labor requirements by vessel type and mean construction date (defined as the mid-point between dates of keel laying and delivery). Lower panel records production time by vessel type and mean construction rate. Upper panel also shows actual (solid line) and interpolated (dashed line) values for the average monthly ULR (by month of delivery) for standard Liberty ships alone. Lower panel also shows the cumulative number of vessels launched.

increase in unit production times (lower panel). Equivalently, design changes induce the same negative association between labor productivity and line speed that is predicted by forgetting models.[6] One can also see in the figure the impact on line speed in the standard Liberty of the complications caused by producing multiple designs simultaneously. However, it is evident that

---

6. The slope of the cumulative output curve in the lower panel approximates the line speed. Note the near-zero slope during the start up of the tanker production.

6

excluding the modified designs from the productivity data tell a rather different story about the time path of productivity. As the upper panel shows, there is little, if any, decline in the productivity of labor employed in the production of standard Liberty ships, despite a sustained interruption to production.

The product design changes at the Delta shipyard were not an anomaly. Ten of the sixteen Liberty shipyards produced modified Liberty ships or different vessels altogether alongside the standard Liberty at some point during the war. Moreover, most of these yards wound down Liberty production by switching labor to entirely new vessels whose output is not recorded in Fischer's statistical summary. Monthly data on production and labor inputs will be very misleading unless one can fully account for variations in the output mix. To accomplish this, one needs data on individual production units, which this study uses.

3. The Learning Curve

Given the unit productivity data available, the following model is a reasonable formulation for a prototypical learning-forgetting model:

$$q_{ij} = F_j(E_{ij})K_{ij}^{\alpha}e^{Z_{ij}\beta} + u_{ij} , \tag{1}$$

where q_{ij} is labor productivity (the inverse of the unit labor requirement) on the ith unit produced at plant j, K_{ij} is physical capital, and Z_{ij} is a vector of potential additional regressors.[7] Experience, E_{ij}, is assumed to depreciate at a constant exponential rate δ and to accumulate as a result of production experience:

$$\begin{cases} E_{1j} = 0 \\ E_{ij} = 1 + e^{-\delta(t_{ij} - t_{i-1,j})}E_{i-1,j}, & i = 2, 3, \dots \end{cases} \tag{2}$$

where t_{ij} is the calendar time at which unit i is produced.

3.1 Loglinear learning

Theory provides little guidance as to the form in which knowledge should enter the ULR (c.f., Muth 1986). The standard approach is to assume a loglinear learning curve of the form

---

7. A minor variation to (1) assumes a multiplicative disturbance. In all regressions reported below, however, the additive disturbance outperforms the multiplicative error on goodness of fit criteria.

$F(E_{ij}) = A_j E_{ij}^{\lambda}$, where A_j is a yard-specific constant.[8] This specification assumes equiproportionate differences in yard productivity at any given level of experience, an assumption strongly at odds with the Liberty ship experience. Early differences in productivity across yards were largely driven by the state of readiness of the shipyard at the time production commenced, delays in the delivery of steel, and delays in labor recruitment (Lane 1951). These were start-up constraints that had no impact on yard productivity at the end of the war. To allow for deviations from equiproportionate yard effects, I therefore assume a loglinear specification of the form,

$$F_j(E_{ij}) = A_j E_{ij}^{\lambda_j} , \tag{3}$$

that allows for both different initial productivities and different learning rates.

A potential danger of the loglinear specification is that, if productivity is bounded, it is capable of generating the appearance of forgetting, where none exists. It would be valuable, therefore, to supplement analysis using (3) with alternative specifications that allow for bounded learning even in the absence of forgetting. Relatively little attention has been paid to the issue of functional form, so I think it useful to expand briefly on the possibility that spurious forgetting rates can be induced by unbounded learning. The introduction of forgetting has two major impacts on the learning curve. First, it generates a steady state for knowledge that bounds productivity. The obvious analogy is the neoclassical capital accumulation equation with constant investment, $k_{t+1} = i + (1 - \delta)k_t$. In the absence of depreciation, k grows without bound, while depreciation imposes a steady state $k^* = i / \delta$. Second, forgetting increases flexibility in fitting the curvature of the learning curve. Unbounded learning and deviations in curvature from the loglinear specification are precisely the two main weaknesses of loglinear learning that have been identified in the past (Auerswald et al. 2000, Muth 1986), and the appearance of forgetting can then simply be the result of a misspecified learning curve.

To illustrate just how much this could matter empirically, a bounded hyperbolic learning curve without forgetting, of the form $q_{i+1} = (i + p)/(i + p + r)$ was used to generate a time-series of 100 nonrandom observations for productivity. A learning curve of the form $\ln q_i = a \ln E_i + u_i$, where E_i evolves according to (2), was then fit to the data. The imputed knowledge depreciation rates so

---

8. Although alternative functional forms have sometimes been applied in different settings (Carlson 1961, 1973, Nembhard and Uzumeri 2000a), the loglinear model remains by far the most widely employed specification.
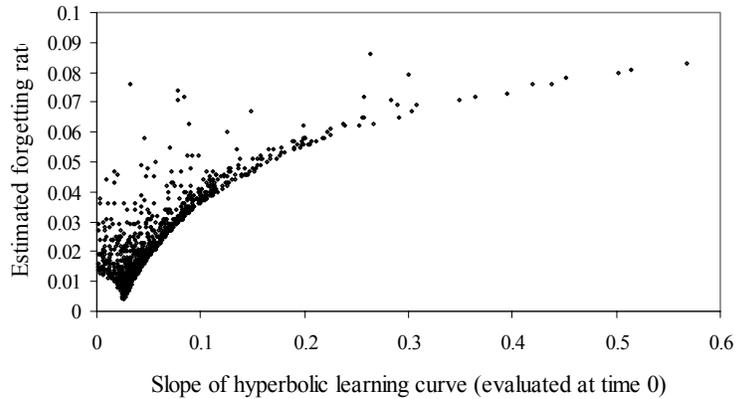
FIGURE 2. Imputed forgetting rates obtained from fitting a log-linear learning curve to artificial data obtained from a bounded hyperbolic learning function. 1,000 replications.

obtained from 1,000 replications generated by random parameter draws of $p, r \in [0,10]$ vary between 0.4 and 11.7 percent, even though no forgetting is present in the data generating process. Moreover, there is a systematic relationship between the imputed depreciation rate and various characteristics of the learning curve. For example, Figure 2 shows a strong positive relationship between the estimated depreciation rate and the slope of the hyperbolic functions evaluated at zero experience. More generally, easier technologies in which early learning is rapid and the terminal productivity is approached more quickly generate higher imputed depreciation rates. It is perhaps no surprise, then, that "[a]n interesting hypothesis that is consistent with our results is that more technologically sophisticated organizations exhibit less depreciation than less technologically sophisticated ones" (Argote 1999, pp. 60-61).

3.2 Bounded learning

Because the artificial data just generated contain no noise, Figure 2 likely exaggerates the practical consequences of assuming unbounded learning. Nonetheless, one would be well-served to explore simultaneously alternatives that do not *require* forgetting to bound productivity. Information theoretic models of learning provide such alternatives in a natural way. Prominent examples include the Prescott and Visscher (1980) task assignment model, the Jovanovic and Nyarko (1995) task learning model, and Jovanovic's (1979) labor turnover model. In the assignment model, workers are assumed to have different skills that make them comparatively more suitable for one of two tasks that must be undertaken by the firm. The firm does not know in advance

9

which worker should be assigned to which task, but must learn this over time by observing worker performance. Productivity attains a finite upper bound when all workers are assigned correctly. In the task learning model productivity is a decreasing function of the distance between a target and an action chosen by the firm. The target is unknown at first and must be learned over time. Productivity attains its upper bound when the firm has completely learned the target. In the labor turnover model, worker productivity depends on the quality of the match between worker and firm. Match quality is discovered slowly. Workers that turn out to make poor matches depart and are replaced. Firm productivity stops rising when all workers turn out to make matches good enough that they stay with the firm.

Unfortunately, these models are quite difficult to make empirically operational.[9] I therefore employ two specifications familiar to mathematical psychologists (e.g. Restle and Greeno 1970). Both generate bounded learning, and it turns out in each case to be simple to allow for adequate yard fixed effects. The first of these, known as the accumulation model, has the form

$$F_j(E_{ij}) = \frac{b_j + a_j \gamma E_{ij}}{1 + \gamma E_{ij}} \, , \tag{4}$$

where a_j, b_j, and λ are positive parameters. The second specification,

$$F_j(E_{ij}) = a_j - (a_j - b_j)(1 - \gamma)^{E_{ij}} \, . \tag{5}$$

is known as the replacement model. Both functions predict an initial productivity of b_j, and a terminal productivity of a_j, and so readily allow for yard effects as rich as those in (3). The parameter γ governs the rate of learning.[10] To interpret γ, consider a simplified model where expected productivity satisfies $\bar{q}_i = F(E_i)$ and $E_i = i$. In the replacement model, $\gamma = (\bar{q}_{i+1} - \bar{q}_i)/(a - \bar{q}_i)$, which is the change in expected productivity between adjacent units expressed as a fraction of the amount left to learn. The accumulation model is a little more compli-

---

9. Thornton (2000) has explored in some detail the finite-sample properties of the least squares estimator of the Jovanovic-Nyarko model. With four parameters only weakly identified by nonlinearity, it proved extremely challenging to obtain reliable parameter estimates for the model.

10. Equations (4) and (5) allow for yard-specific initial and terminal productivities but impose the same learning parameter on all yards. The loglinear model has no terminal productivity (absent forgetting). However, allowing for different learning rates in (3) implies all three specification have similar degrees of freedom and hence are comparable on a level playing field.

cated, as $\gamma = (\overline{q}_{i+1} - \overline{q}_i)/(a - \overline{q}_i - i(q_{i+1} - q_i))$ contains an extra term in the denominator. There is, of course, no constant relationship between γ and λ, but we may expect estimates of the latter to be an order of magnitude larger than the former.[11]

The names for the two specifications are derived from an intuitive appeal to two urn problems that generate these functions. Consider first the following replacement problem. There are two urns, A and B. Urn A contains a fixed number of marbles, of which a fraction b is red and a fraction $1-b$ is white. On each trial, one marble is drawn from A. If it is red a correct response is recorded, while a white marble indicates an incorrect response. Urn B contains an infinite number of marbles, a fraction $a>b$ of which are red. After each trial, a fraction λ of the marbles in A are replaced by marbles drawn from B. Let expected productivity be represented by the probability of a correct response. Then it can be shown (Restle and Greeno 1970, ch. 1) that expected productivity in the i^{th} trial is given by (4). The accumulation problem differs only in that after each trial a constant number of marbles from B are added to A without removing any from A. Let λ denote the number of marbles added after each trial, expressed as a fraction of the initial number of marbles in A. Then, (5) gives the probability of a correct response in the i^{th} trial.

This is a rather abstract way of thinking about learning. However, the urn problems capture the idea that learning can be thought of as either the replacement of incorrect ways of doing tasks with correct ways of doing them, or as the accumulation of new skills on top of existing skills. It is not obvious which of these is more appropriate for any given learning problem, individual or organizational. For example, learning to type can equally well be envisaged as a process of accumulating new skills or of replacing bad ways to type with good ways. Learning in wartime shipbuilding may equally be interpreted ambiguously. This is probably not important. What matters is that both models are computationally tractable, they are capable of producing non-zero initial, and finite terminal, productivity levels, and they offer alternative assumptions about how firms get from the former to the latter while still providing a compact measure of the learning rate. The next section estimates both bounded learning models along with the loglinear specification.

---

11. A back-of-the-envelope calculation illustrates. If the loglinear model is $\overline{q}_i = i^{\lambda}$, the replacement model implies a relationship, to first order, satisfying $\gamma \approx \lambda i^{\lambda-1}/(a - i^{\lambda})$. In the data the sample mean terminal productivity, a, is about 3.5 while the sample mean experience is $i=130$. We will produce estimates $\hat{\gamma} \approx 0.01$, which at the sample mean implies estimates for λ exceeding 0.2.

4. Estimation

The shipbuilding data used in this study have been fully documented in Thompson (2001), and are only briefly described here. Between 1941 and 1945, the US Maritime Commission procured 5,777 ocean-going vessels for war service. Prominent among these were 2,699 Liberty-type cargo vessels, an all-welded dry cargo vessel of some 7,000 displacement tons. Sixteen purpose-built yards were engaged in the production of Liberty vessels at one time or another. However, three of these yards produced 20 vessels or fewer and are excluded from the study, leaving a sample size of 2,662. Of these, 130 were modified designs, including tankers, colliers, tank and aircraft transporters. Yet another 31 vessels were constructed as standard designs, but delivered to the government incomplete for subsequent conversion to troop carriers or training ships. The data contain observations on the unit labor requirement, date of keel laying and date of delivery, for each of these 2,662 ships. For each yard, I use the number of ways in use as a proxy for capital, and I record monthly labor data, most notably average job separation and hiring rates. While I do not have complete comparable data for non-Liberty vessels, I do know the dates know when a Liberty shipyard was engaged in the production of non-Liberty vessels.[12]

In all the regressions that follow, vessels are ordered by date of keel laying. Ordering vessels by date of delivery produced smaller, and often negative, estimates of the rate of forgetting, but led to models that fit the data less well. I estimate the models using a computationally convenient two-step procedure. First, I conduct a grid search over δ, constructing new series for E at each step and estimating (1) by nonlinear least squares (NLLS) conditional on δ.[13] For the standard errors, I then approximate the gradient vector, $\mathbf{G}$, with respect to all parameters, including δ, using the delta method. The asymptotic variance of the coefficient vector is given by $\mathrm{var}(\hat{\beta}) = 2m(\mathbf{GG}')^{-1}$, where m is the mean square error of the NLLS estimator (Greene 2003, ch. 10).

---

12. Unit labor requirements were obtained from typescript records held in the National Archives [Boxes 35 and 37, Records of the Historian's Office, National Archives RG 178]. Dates of keel laying are from typescript records in Boxes 35 and 37, while dates of delivery were constructed by adding to the keel laying dates the construction times recorded in handwritten tabulations [various boxes]. The number of ways in use was also obtained from these data. Labor data were obtained from monthly reports of the Bureau of Labor Statistics forms BLS1761, Plant Operations, Box 36.

13. Because I assume an additive disturbance, the loglinear model must also be estimated by NLLS.

TABLE 1

NLLS estimates of learning curves

		-------------------ALL DATA--------------------			------CORRECTIONS FOR PRODUCT MIX------		
	Loglinear	Accumulation	Replacement	Loglinear	Accumulation	Replacement
	(1)	(2)	(3)	(4)	(5)	(6)
δ	0.084	0.087	0.058	0.042	0.057	0.036
	(.011)	(.010)	(.012)	(.010)	(.009)	(.012)
γ		0.008	0.011		0.007	0.009
		(.001)	(.001)		(.001)	(.001)
Average λ[a]	0.320			0.303		
α	−0.210	−0.109	−0.125	−0.103	−0.009	−0.033
	(.069)	(.072)	(.068)	(.067)	(.067)	(.066)
2 vessel types under construction				−0.084	−0.074	−0.083
				(.079)	(.020)	(.100)
3 vessel types under construction				−0.119	−0.113	−0.130
				(.115)	(.057)	(.156)
4 vessel types under construction				−0.239	−0.225	−0.217
				(.223)	(.039)	(.231)
N	2,649	2,649	2,649	2,488	2,488	2,488
$\bar{R}^2$	0.992	0.992	0.992	0.994	0.994	0.994

Dependent variable is the inverse of the unit labor requirement in millions of hours. The first vessel launched by each yard is dropped from the sample. Asymptotic standard errors in parentheses. Estimates of A_j, a_j and b_j not reported. [a] Average λ is the mean of thirteen yard-specific growth rates weighted by sample size: see Table A.1 for individual estimates and standard errors.

Table 1 reports the first results. Column 1 uses the loglinear specification from (3), and fails to account in any way for modified designs. The specification includes thirteen distinct learning rates and thirteen yard-specific constant terms. As usual, the estimates for the constant terms, which are precisely estimated and vary significantly across yards, are not reported. To avoid distracting clutter, the individual learning rates are relegated to Table A.1 in the appendix, while in Table 1 the sample-weighted mean is provided. The fit of the regression is remarkably good, as one would expect with strongly trending data. The regression returns an estimate for δ indicating rapid knowledge depreciation, at the rate of 8.4 percent per month. Nonetheless, this rapid imputed

rate of forgetting is about one third of the rate previously reported for this industry in Argote, Beckman and Epple (1990). I conjecture that the rate is lower because the ULR data do not misattribute to Liberty ship production labor used in the production of other types of vessels. Unsurprisingly, the average of the learning parameter, λ, is also about one third the rate estimated previously. However, one result suggestive of a specification problem is the point estimate for α, which at -0.2 indicates that larger yards have much lower productivity than small yards, and this is something we know not to be true.[14]

Columns 2 and 3 report the results for the accumulation and replacement models respectively. The surprising result is that the switch to bounded learning curves has only a modest impact on the estimated rate of forgetting. In fact, there is essentially no change in the estimate upon moving from the loglinear to the accumulation model, which continues to return a rate of about 8.5 percent. The replacement model does produce a lower estimate, however, which at 5.8 per cent per month represents a proportional decline of around 30 percent. One notable change is that the point estimates of α are much reduced and now statistically insignificant. The estimates of the learning parameter, γ, are also plausible. The gross learning rate, expressed as a fraction of the amount left to learn, is approximately one percent per vessel constructed.

Columns 4 through 6 make corrections for changes in the product mix. Doing so requires three adjustments. First, incomplete vessels of standard designs are included in the construction of the experience measure, but then excluded from the regressions. This is equivalent to assuming full knowledge spillovers from incomplete vessels to standard designs subsequently constructed, and has the effect of increasing the estimated rate of forgetting relative to assuming zero or partial spillovers. Second, non-standard Liberty designs were excluded from the data prior to constructing the experience measure. This is equivalent to assuming zero spillovers from non-standard designs to standard designs.[15] Third, to account for inefficiencies caused by disrupting the operations

---

14. The number of ways is more a proxy for the scale of operation than it is for the capital stock. Thompson (2001) used detailed data on capital authorizations for six yards, and showed that a significant amount of learning could be attributed to increases in the amount of capital per way.

15. This set of assumptions produced the best fit, although the difference from alternative treatments was modest. Benkard (2000) underline{estimated} the degree of spillovers from one aircraft design to another. Attempts to do so here failed to attain convergence, no doubt because of the relatively small number, and scattered appearance, of modified and incomplete designs.

14

of yards intended for production of a single vessel design, three indicator variables were added to record whether the shipyard was simultaneously producing other types of vessels.

The main effect of these changes is, as expected, to induce a marked reduction in the estimated rates of forgetting. For the loglinear model, $\hat{\delta}$ declines by exactly half to 4.2 per month. For the bounded learning models it declines by about 35 percent, yielding estimates of 5.7 percent per month for the accumulation model and 3.6 percent per month for the replacement model. The remaining parameters are only modestly affected by the adjustments for product mix. Despite the declines in the reported rate of forgetting, the estimates for the learning parameters, λ and γ, have also declined a little. All three estimates of α have also declined, and are now statistically insignificant. Thus, corrections for product mix plausibly yield estimates consistent with a constant returns to scale technology. Finally, the coefficients on the indicator variables for simultaneous production of multiple vessel types are negative, economically and statistically significant, and with the expected rank ordering.[16]

4.1 Labor Turnover

A frequent motivation for organizational forgetting is that knowledge embodied in workers is lost when they leave. If that is the case, labor turnover should be negatively correlated with productivity and, when measures of turnover are included in the model, the estimated rate of knowledge depreciation over time should decline. Columns 1 through 3 of Table 2 add as a level effect the recorded rates of labor hiring and separation in the month in which the vessel was delivered. In each case the estimate of δ does indeed decline, to 1.8 and 2.1 percent per month for the two bounded learning models. The estimated rate of forgetting is now *negative* but statistically insignificant for the loglinear case. However, the coefficients on the hiring and separation rates are consistently the wrong sign, although in five of six cases statistically insignificant, suggesting that increased labor turnover either has no effect on productivity, or raises it. These results are clearly inconsistent with conventional wisdom about the role of labor turnover.

One possible explanation for the implausible positive coefficients on the labor turnover measures is that the model may be fundamentally misspecified in this area. If organizational forgetting

---

16. Alternative specifications not reported also included controls for night shift work and Sunday shift work, intended as proxies for the intensity of production. The additional controls had no effect on the reported results.

TABLE 2

Learning, forgetting and labor turnover

	Loglinear	Accumulation	Replacement	Loglinear	Accumulation	Replacement
	(1)	(2)	(3)	(4)	(5)	(6)
δ	−0.014	0.016	0.021	−0.003	−0.002	−0.001
	(.011)	(.007)	(.012)	(.001)	(.001)	(.012)
$\tilde{\delta}$				−0.036	−0.017	−0.006
				(.012)	(.005)	(.010)
γ		0.006	0.008		0.007	0.008
		(.001)	(.001)		(.001)	(.001)
Average λ	0.293			0.244		
α	−0.052	0.011	−0.016	−0.039	−0.014	−0.032
	(.066)	(.066)	(.070)	(.062)	(.079)	(.065)
Hiring rate	0.646	0.549	0.498			
	(.088)	(.450)	(.639)			
Separation rate	0.011	0.075	0.062			
	(.076)	(.122)	(.096)			
2 vessel types under construction	−0.132	−0.110	−0.102	−0.154	−0.129	−0.115
	(.015)	(.020)	(.129)	(.165)	(.018)	(.089)
3 vessel types under construction	−0.191	−0.160	−0.148	−0.210	−0.183	−0.165
	(.040)	(.052)	(.184)	(.223)	(.047)	(.130)
N	2,233	2,233	2,233	2,233	2,233	2,233
$\bar{R}^2$	0.994	0.994	0.994	0.994	0.994	0.994

Some observations are lost due to missing data for labor turnover, and few observations remain with four vessel types under construction. All regressions include corrections and controls for product mix. See also notes to Table 1.

results from labor turnover, there should be forgetting only indirectly with respect to time and directly with respect to labor turnover. To explore this idea, consider the following variation on the assignment model of Prescott and Visscher (1980). A firm of age T employs workers indexed by a type, θ, where $\theta \sim N\left(0, \sigma_\theta^2\right)$, and must assign each worker to one of two tasks. Suppose productivity is given by

$$q_T = C + \sum_{t=0}^{T} \mu_t \left(\overline{\theta}_t^A - \overline{\theta}_t^B \right),$$ (6)

where $\bar{\theta}_t^j$ is the average type of workers with tenure t assigned to task $j=\{A,B\}$, and μ_t is the fraction of workers with tenure t. Clearly the firm would like to assign all workers with positive type to task A and those with negative type to task B. Let θ_i, denote the type of the i^{th} individual. The firm does not initially know the individual's type, but must learn it over time by observing signals satisfying $x_{i,t} = \theta_i + \varepsilon_{i,t}$, $\varepsilon \sim N\left(0,\sigma_\varepsilon^2\right)$. After t signals, the firm's posterior distribution of i's type, $F_t(\theta \mid \bar{x})$, is $N\left(t\theta_\theta^2 \bar{x}_{i,t} / \left(t\theta_\theta^2 + \theta_\varepsilon^2\right), \theta_\theta^2\theta_\varepsilon^2 / \left(t\theta_\theta^2 + \theta_\varepsilon^2\right)\right)$, where $\bar{x}_{i,t}$ is the mean of the t signals. The firm assigns worker i to task A if $\bar{x}_{i,t} > 0$ and task B otherwise (reassignment is possible whenever the sign of $\bar{x}_{i,t}$ changes). Suppose the firm is large, so that approximately half the labor force is assigned to each task. Then, expected productivity is

$$\bar{q}_T = C + \sum_{t=0}^{T} \mu_t \left(2\int_0^\infty \int_{-\infty}^\infty \theta dF_t(\theta \mid \bar{s}) dG_t(\bar{x}) - 2\int_{-\infty}^0 \int_{-\infty}^\infty \theta dF_t(\theta \mid \bar{x}) dG_t(\bar{s}) \right)$$

$$= C + 2\sigma_\theta \pi^{-\frac{1}{2}} \sum_{t=0}^{T} \mu_t \left(\frac{2t}{t\sigma_\theta^2 + \sigma_\varepsilon^2} \right)^{\frac{1}{2}} \tag{7}$$

where G_t, the unconditional distribution of $\bar{x}$, is $N\left(0, \sigma_\theta^2 + \sigma_\varepsilon^2 / t\right)$. In (7), the term in parentheses is strictly increasing in t, so the expected contribution of any worker rises with tenure. It is now easy to see how labor turnover affects productivity. Assume a constant separation rate, s, that workers of any tenure are equally likely to depart, and that departures are replaced immediately by new workers. A steady state distribution for μ (which requires $T \to \infty$) satisfies $\mu_0 = s$, and $\mu_{t+1} = (1-s)\mu_t$ for $t \geq 1$. This implies $\mu_t = s(1-s)^t$, $t = 0,1,2, \ldots$, and hence

$$\bar{q} = C + 2\sigma_\theta \pi^{-\frac{1}{2}} \sum_{t=0}^{\infty} s(1-s)^t \left(\frac{2t}{t\sigma_\theta^2 + \sigma_\varepsilon^2} \right)^{\frac{1}{2}}. \tag{8}$$

The term $s(1-s)^t$ is increasing [decreasing] in s for $t < [>](1-s)/s$, so that an increase in labor turnover increases the weight put on low values of the second term in parentheses, and decreases the rate put on high values. The net effect, of course, is to reduce $\bar{q}$.

It remains to make this idea operational in our data. Let L_t denote the number of employees at time t and let h_t and s_t denote the mean monthly hiring and separation rates over the period $(t,t+\Delta)$. Then, the fraction of inexperienced workers at time $t+\Delta$ can be written as

$$\rho_{t+\Delta} = 1 - \frac{(1 - s_t\Delta)L_t}{L_{t+\Delta}} = \frac{h_t\Delta}{1 + (h_t - s_t)\Delta}. \tag{9}$$

17

The specifications in the final columns of Table 2 assume that depreciation of knowledge in any month depends positively upon ρ. That is, knowledge evolves according to

$$\begin{cases} E_{1j} = 0 \\ E_{ij} = 1 + e^{-\tilde{\delta}\rho_{ij}} E_{i-1,j}, \quad i = 2, 3, \ldots \end{cases} , \tag{10}$$

where

$$\rho_{ij} = \frac{h_{ij}\Delta_{ij}}{1 + (h_{ij} - s_{ij})\Delta_{ij}} , \tag{11}$$

and s_{ij} and h_{ij} are the separation and hiring rates for the month of vessel delivery. I let the relevant period, Δ_{ij}, over which depreciation takes place for the i^{th} vessel at yard j be the difference in months between the date of keel laying and the date of delivery.[17] To interpret $\tilde{\delta}$, consider a steady-state employment level where $h_j = s_j$ and hence $\rho_j = h_j\Delta_j$. The depreciation rate of knowledge implied by (9) and (10) is then $\delta = \tilde{\delta}h_j \equiv \tilde{\delta}s_j$ per month. I estimate $\tilde{\delta}$ by the usual grid search, and then calculate $\delta = \tilde{\delta}\overline{s}$, where $\overline{s} = 0.104$ is the sample mean separation rate.[18] The results are reported in columns 4 through 6 of Table 2. As before, the fit of the model is very high, and the parameters other than δ are essentially unaffected by the changed specification. In contrast, the monthly rates of forgetting returned by the new specification are tiny and, although statistically distinguishable from zero in two cases, they are economically indistinguishable from zero in all three.

5. Conclusions

This paper has presented new estimates of the rate of organizational forgetting in wartime shipbuilding using unusually detailed and precise data on unit labor requirements. The richness of the data facilitated a focus on two potentially important confounding issues. First, unobserved changes in a firm's product mix can produce spurious evidence for organizational forgetting. Second, the estimated rate of forgetting may be sensitive to assumptions made about the learning

---

17. Other regressions, in which depreciation was assumed to take place between successive keel laying dates, produced large *negative* rates of forgetting, but fit the data as well as the specification reported here.

18. This is virtually identical to the sample mean hiring rate of 0.103, thereby justifying the steady-state employment assumption.
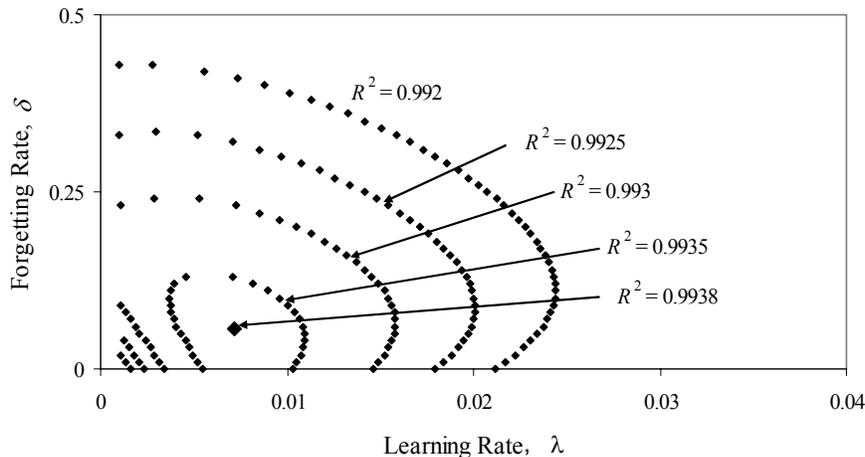
Figure 3. Contour lines for coefficient of determination, from accumulation model (Table 1, column 5). The contours are drawn by fixing d at different levels and then estimating the remaining parameters by NLLS.

process. It was found that for this data set, controlling adequately for changes in the product mix has significant effects on the estimated rate of forgetting, but the results were only moderately sensitive to the specification of the learning curve. The rate of organizational forgetting estimated in this paper is markedly lower than has previously been reported. In a seminal study on organizational forgetting, Argote, Beckman and Epple (1990) had reported an estimated 25 percent monthly rate of knowledge depreciation in the industry. The estimates reported in columns 4 through 6 of Table 1 of the present paper range from 3.6 percent to 5.7 percent.

Nonetheless, these are still significant rates of knowledge depreciation. Because much previous literature has suggested that organizational forgetting may be driven by labor turnover, I explore the relationship between productivity changes and labor turnover through two models, one which incorporates labor hiring and separation rates as simple level effects, and another which assumes organizational forgetting operates directly through turnover and only indirectly as a function of time. Labor turnover appeared to be largely unrelated to productivity changes. However, incorporating them essentially eliminated all evidence of organizational forgetting.

It is prudent to close with a note of caution. The fact that we are necessarily dealing with strongly trending data makes the estimated rate of forgetting quite sensitive to model specification, while allowing wide variations in the forgetting rate to account well for variations in the data (see Figure 3). Moreover, the relatively modest rates estimated in this paper compound the importance of uncertainty about the point estimate. For example, in Table 2, column 2, the point

19

estimate is a rate of forgetting of 1.6 percent per month, with a 95 percent confidence interval of [0.2 percent, 3.0 percent]. The strategic implications of organizational forgetting are vastly different at the two ends of the interval.

Yard-specific estimates of learning rates, loglinear specifications

SHIPYARD [a]	NO. OF LIBERTY SHIPS BUILT [b]	TABLE 1 COLUMN 1	TABLE 1 COLUMN 4	TABLE 2 COLUMN 1	TABLE 2 COLUMN 4
Bethlehem-Fairfield Shipyards, Inc. (Baltimore, MD)	384	0.172 (.015)	0.182 (.013)	0.205 (.012)	0.156 (.013)
California Shipbuilding Corp. (Los Angeles, CA)	306	.278 (.014)	0.325 (.014)	0.308 (.013)	0.258 (.018)
Delta Shipbuilding Co. (New Orleans, LA)	188	0.136 (.024)	0.234 (.025)	0.216 (.019)	0.202 (.018)
Todd-Houston Construction Co. (Houston, TX)	208	0.408 (.029)	0.348 (.024)	0.298 (.021)	0.243 (.020)
J.A. Jones Construction Co. (Brunswick, GA)	85	0.484 (.040)	0.443 (.034)	0.429 (.029)	0.371 (.028)
J.A. Jones Construction Co. (Panama City, FL)	102	0.545 (.044)	0.459 (.038)	0.449 (.032)	0.420 (.032)
North Carolina Shipbuilding Co. (Wilmington, NC)	126	0.211 (.016)	0.200 (.014)	0.224 (.014)	0.179 (.012)
New England Shipbuilding Co. (South Portlasnd, ME)	244	0.484 (.031)	0.419 (.025)	0.351 (.020)	0.331 (.015)
Oregon Ship Building Corp. (Portland, OR)	330	0.345 (.013)	0.310 (.012)	0.298 (.012)	0.219 (.021)
Permanente Metals Corp., No. 1 Yard (Richmond, CA)	138	0.186 (.020)	0.118 (.017)	0.165 (.019)	0.143 (.018)
Permanente Metals Corp., No. 2 Yard (Richmond, CA)	351	0.380 (.014)	0.349 (.013)	0.353 (.017)	0.278 (.020)
Southeastern Shipbuilding Co. (Savannah, GA)	88	0.326 (.044)	0.287 (.035)	0.275 (.029)	0.232 (.026)
St. John's River Shipbuilding Co. (Jacksonville, FL)	82	0.471 (.051)	0.415 (.041)	0.372 (.033)	0.335 (.031)

Asymptotic standard errors in parentheses. [a] Yards omitted because of small volumes are: Alabama Dry Dock Co., Mobile AL (20 delivered), Kaiser Co., Vancouver WA (2 delivered, 8 transferred incomplete to another yard), Marinship Corp., Sausolito CA (15 delivered), Walsh-Kaiser Co., Providence RI (10 delivered), and Rheem Manufacturing Co., Providence RI (1 delivered). [b] Includes modified designs and vessels delivered incomplete.

References

Anderlohr, G. (1969): "Determining the Cost of Production Breaks." *Management Review*, **58**(12):16-19

Argote, Linda (1999): *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Boston, MA: Kluwer Academic Publishers.

Argote, Linda, Sara L. Beckman and Dennis Epple (1990): "The Persistence and Transfer of Learning in Industrial Settings." *Management Science,* **36**(2):140-154.

Argote, Linda, Dennis Epple, Rukmini Devadas Rao, and Kenneth Murphy (1997): "The Acquisition and Depreciation of Knowledge in a Manufacturing Organization: Turnover and Plant Productivity." Working Paper, Carnegie Mellon University.

Auerswald, Phillip, Stuart Kauffman, José Lobo, and Karl Shell (2000): "The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing." *Journal of Economic Dynamics and Control*, **24**(3):389-450.

Baloff, N. (1970): "Startup Management." *IEEE Transactions on Engineering Management*, **17**:132-141.

Belason, Allan T. (2000): *Leading the Learning Organization*. Albany, NY: State University of New York Press.

Benkard, C. Lanier (2000): "Learning and Forgetting: The Dynamics of Aircraft Production." *American Economic Review*, **90**(4):1034-1054.

Carlson, J.G. (1961): "How Management Can Use the Improvement Phenomenon." *California Management Review*, **3**(2):83-94.

Carlson, J.G. (1973): "Cubic Learning Curves: Precision Tool for Labor Estimating." *Manufacturing Engineering and Management*, **71**(5):22-25.

Cross, Rob, and Lloyd Baird (2000): "Technology is not Enough: Improving Performance by Building Organization Memory." *Sloan Management Review*, **41**(3):69-78.

Darr, Eric D., Linda Argote and Dennis Epple (1995): "The Acquisition, Transfer, and Depreciation of Knowledge in Service Organizations: Productivity in Franchises." *Management Science*, **41**(11):1750-1762.

Dutton, John M., and Annie Thomas (1984): "Treating Progress Functions as a Managerial Opportunity." *Academy of Management Review*, **9**(2):235-247.

Ebbinghaus, Herman (1885): *Ueber das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Leipzig: Duncker and Humblot.

Epple, Dennis, Linda Argote and Rukmini Devadas (1991): "Organizational Learning Curves: A Method of Investigating Intra-Plant Transfer of Knowledge Acquired Through Learning by Doing." *Organization Science,* **2**(1):58-70.

Epple, Dennis, Linda Argote and Kenneth Murphy (1995): "An Empirical Investigation of the Microstructure of Knowledge Acquisition and Transfer Through Learning by Doing." *Operations Research,* **44**(1):77-86.

Finkenbinder, E.O. (1913) "The Curve of Forgetting." *American Journal of Psychology*, **24**:8-32.

Fischer, George J. (1949): "A Statistical Summary of Shipbuilding under the US Maritime Commission." *Historical Reports of the War Administration*, US Maritime Commission.

Greene, William H. (2003). *Econometric Analysis.* 5[th] edition. New York: Prentice Hall.

Hirsch, Werner Z. (1952): ""Manufacturing Progress Functions." *Review of Economics and Statistics*, **34**:143-155.

Huber, G.P. (1991): "Organizational Learning: The Contributing Processes and the Literatures." *Organizational Science*, **2**(1):88-115.

Irwin, Douglas, and Peter Klenow (1994): "Learning-by-Doing Spillovers in the Semiconductor Industry." Journal of Political Economy, **102**(6):1200-1227.

Jovanovic, Boyan (1979): "Job Matching and the Theory of Turnover." *Journal of Political Economy*, **87**(5):972-990.

Jovanovic, Boyan and Yaw Nyarko (1995): "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves." *Brookings Papers: Microeconomics 1995*, pp. 247-305.

Kransdorrf, Arnold (1997): "Fight Organizational Memory Lapse." *Workforce*, **76**(6):34-39.

Landry, John R. (1999): "Forgetful or Bad Memory?" Proceedings of the 32ⁿ Hawaii International Conference on System, Sciences – 1999. Paper 1069. IEEE Computer Society. http://computer.org/proceedings/hicss/0001/00011/0001toc.htm. Accessed 24 August, 2002.

Lane, Frederic C. (1951): *Ships for Victory*. Baltimore, MD: Johns Hopkins University Press.

Lucas, Robert E. Jr. (1993): "Making a Miracle." *Econometrica,* **61**(2):251-272.

McCreery, J.K., and L.J. Krajewski (1999): "Improving Performance Using Workforce Flexibility in an Assembly Environment and Learning and Forgetting Effects." *International Journal of Production Research*, **37**(9):2031-2058.

McGeoch, John A. (1932): "Forgetting and the Law of Disuse." *Psychological Review*, **39**:352-370.

Muth, John F. (1986): "Search Theory and the Manufacturing Progress Function." *Management Science*, **32**(8):948-962.

Nembhard, David A. (2001): "An Empirical Comparison of Forgetting Models." *IEEE Transactions on Engineering Management*, **48**(3):283-291.

Nembhard, David A., and Mustafa V. Uzumeri (2000a): "An Individual-Based Description of Learning within an Organization." *IEEE Transactions on Engineering Management*, **47**(3):370-378.

Nembhard, David A. and Mustafa V. Uzumeri (2000b): "Experiential Learning and Forgetting for Manual and Cognitive Tasks." *International Journal of Industrial Ergonomics*, **25**(2):315-326.

Peters, Thomas J. (1999): *The Circle of Innovation*. New York: Vintage.

Prescott, Edward C., and Michael Visscher (1980): "Organizational Capital." *Journal of Political Economy*, **80**(1):446-461.

Rapping, Leonard (1965): "Learning and World War II Production Functions." *Review of Economics and Statistics*, **47**(1):81-86.

Restle, Frank, and James G. Greeno (1970): *Introduction to Mathematical Psychology*. Reading, MA: Addison-Wesley.

Searle, Allan D. (1945): "Productivity Changes in Selected Wartime Shipbuilding Programs." *Monthly Labor Review,* (December):1132-1147.

Smunt, Timothy L. (1987): "The Impact of Worker Forgetting on Production Scheduling." *International Journal of Production Research,* **25**:689-701.

Thompson, Peter (2001): "How Much Did the Liberty Shipbuilders Learn? New Evidence for an Old Case Study." *Journal of Political Economy*, **109**(1):103-137.

Thorndike, E.L. (1914): *The Psychology of Learning.* NY: Teacher's College.

Thornton, Rebecca Achée (2000): *Essays on Learning from Experience and Knowledge Spillovers.* Ph.D. thesis, University of Houston.

Thornton, Rebecca Achée, and Peter Thompson (2001): "Learning from Experience and Learning From Others. An Exploration of Learning and Spillovers in Wartime Shipbuilding." *American Economic Review*, **91**(5):1350-1368.

Walsh, J.P., and G.R. Ungson (1991): "Organizational Memory." *Academy of Management Review*, **16**(1):57-91.

Wright, T.P. (1936): "Factors Affecting the Cost of Airplanes." *Journal of the Aeronautical Sciences*, **3**(4):122-128.

Yelle, Louis E. (1979): "The Learning Curve: Historical Review and Comprehensive Survey." *Decision Sciences*, **10**:302-328.