Paper to be presented at the

35th DRUID Celebration Conference 2013, Barcelona, Spain, June 17-19

Impact of experience heterogeneity on individual learning curves

Kannan Srikanth
University of Southern Denmark
Strategic Organization Design
ksr@sam.sdu.dk

Anand Nandkumar
Indian School of Business
Strategy
anand_nandkumar@isb.edu

Dutt Dev Harsha Tadikonda
Indian School of Business
Strategy
harsha_tadikonda@isb.edu

Abstract

Learning by doing is one of the key mechanisms that one of the key mechanisms that influences capability development in firms. Recent research has shown that the manner in which individuals within organizations accumulate and learn from experience significantly contributes to the differences in productivity between organizations. In this paper we explore whether individual learning is fungible across work environments. We empirically show that individual learning curves are significantly influenced by the context of work and the nature of experience of the individual. We discuss implications for our understanding of learning curves and how they may influence capability accumulation.

Jelcodes:M11,-

Impact of experience heterogeneity on individual learning curves

ABSTRACT

Learning by doing is one of the key mechanisms that one of the key mechanisms that influences capability development in firms. Recent research has shown that the manner in which individuals within organizations accumulate and learn from experience significantly contributes to the differences in productivity between organizations. In this paper we explore whether individual learning is fungible across work environments. We empirically show that individual learning curves are significantly influenced by the context of work and the nature of experience of the individual. We discuss implications for our understanding of learning curves and how they may influence capability accumulation.

Keywords: Learning by doing, learning curve, fungibility of learning

INTRODUCTION

A considerable literature documents the extent to which prior experience contributes to knowledge accumulation and consequently to improved productivity (c.f. Argote et al, 2003). Commonly referred to as "learning by doing," this idea has played an important role in the analysis of firm productivity. Recent research has shown that the manner in which individuals within organizations accumulate and learn from experience significantly contributes to the differences in productivity between organizations (Reagans, Argote and Brooks, 2005). Building on this idea, a growing body of research has focused on exploring the factors that influence the learning curve. However, largely absent from these studies, is an examination of how fungible individual learning is across the different work environments within an organization. Our study fills this gap, by examining the extent to which individual learning is fungible across work environments that differ in the level of stress in the form of excessive loads placed on the individual workers.

Numerous studies have shown evidence of learning by doing in a variety of industrial settings, (Besanko et al., 2010). Moreover, many studies also provide evidence that learning curves differ substantially organizations (Epple, Argote, Murphy, 1996; Pisano et al, 2001; Reagans et al, 2005). Heeding the call for the need for further research on the factors underlying learning curves (Argote, 1993, Dutton & Thomas 1984, Lieberman 1984), a number of studies have identified factors that moderate the relationship between cumulative experience and organizational productivity (Adler and Clark 1991, Haunschild and Sullivan 2002, Ittner et al. 2001, Lapre and Van Wassenhove 2001, Lieberman 1987, Pisano 1994, Terweisch and Bohn 2001).

The manner in which knowledge at the individual worker level is accumulated has been shown to be a key factor affecting organizational learning curves (Argote, Epple and Devadas, 1991; Regan et al. 2005; Huckman and Pisano, 2006). For instance, Huckman and Pisano (2006), analyze the performance of surgeons operating at multiple hospitals and find that while same hospital experience contributes to performance improvements, experience at other hospitals does not significantly improve surgical outcomes. In a similar vein, Groysberg et al. (2008), find that "star" securities analysts who switch employers experience a decline in their performance, especially among those analysts moving to lesser firms. Narayana et al. (2009) analyze the degree to which task specialization enhances learning, and finds that excessive exposure to task variety is an impediment to effective learning. Boh et al. (2007), find that while rates of individual learning are greatest for specialized tasks, diverse experience does play a positive role in productivity improvement at the group and organizational level. These studies broadly suggest that individual learning in influenced by the organizational context in which skills are acquired, such as the diversity of tasks performed and the organization of work groups.

Another important contingency that influences individual learning is the level of stress or workload in the learning environment. The literature shows that stressful work environments, have been shown in a number of settings to adversely affect production. Scott et al. (2006) show that personnel working during shifts with excessive loads are observed to cause more medical errors in diagnosis and treatment than medical personnel working in less tiresome environments. Similarly, Evans and Kim (2006) analyze patient outcomes on days during which a hospital experiences an unexpected demand surge at hospitals and finds evidence of decreases in productivity during these surges in the form of an increase in readmission and decrease in a patient's length of stay. Kc and Terwiesch's (2009) analysis of patient outcomes for those admitted during periods of elevated hospital workload finds an increase in the mortality rate of cardiac surgery patients, implying a decrease in hospital quality during these periods. While most studies examine how the nature of the environment influences individual learning, our focus is on the degree to which individual learning is transferrable across different work environments, such as in normal versus high stress environments within the same firm. To this end, we make use of a unique dataset that from an emergency service call center that allows us to distinguish between the nature of prior experience as well as the work environments that an individual works in. Specifically, we are able to follow each worker across each shift, and because our data covers the entire universe of experience for each worker as well as the full set of calls handled by each worker during each shift, we can develop a measures of different types of experience by segmenting worker experience based on the pressure faced by each worker during their shifts.

We also make an empirical contribution. Prior work on learning curves, especially the ones that estimate learning under high work-load conditions, has either used productivity or quality as a dependent variable but not both. Presumably, there is a trade-off between productivity and quality – the faster you perform an activity, the more likely quality suffers. There is considerable evidence for such a speed-accuracy trade-off among individuals (Svenson and Maule, 1993). It is unclear whether learning by doing, especially under different contexts, can be predominantly directed toward increases in productivity or increasing quality. Not controlling for quality in a productivity regression can potentially overstate learning, while not controlling for productivity when quality is the dependent variable can potentially understate learning. We are able to tease out these different learning effects to provide a more complete picture of individual learning under normal vs. heavy workload conditions.

Our results indicate that learning is not fully transferrable across work environments. Specifically, our findings indicate that neither is worker experience acquired during particularly busy, high-pressure shifts is not fully transferrable to slower, less demanding work environments nor is experience acquired in high-pressure environments transferrable to less demanding environments. Our results are robust to the addition of a variety of worker and call-specific variables, and are invariant to alternative specifications and definitions of our key independent variable which we use to differentiate the types of shifts experienced by each worker in our data.

This paper proceeds as follows. In section 2, we discuss the theoretical and empirical framework for identifying the effects of learning across work settings. In section 3, we describe our data and the characteristics of the setting in which it was collected while in section 4, we present our results. Section 5 concludes, and discusses directions for possible extensions of our findings.

2. FRAMEWORK TO MOTIVATE THE EMPIRICAL ANALYSIS

We motivate our empirical analysis with a stylized model. In particular, our goal is to explain how differences in work environments influences productivity, defined as number of patient calls per hour and errors which denotes one of two scenarios: dispatching ambulance when not required (false positives) and not dispatching ambulance when required (false negatives). In our model the manager of an organization forecasts the number of calls in a shift based on prior call volume. The actual number of calls can be higher or lower than the expected number of calls. With an eye on empirics, "normal" time is one in which the actual call volume is less than or equal to expected volume. "High Load" is one in which the actual call volume exceeds expected volume.

2.1 Production function

The productivity for an average worker per unit of time differs by the time environment with her productivity being A under normal time and yA under high load with y>1. Thus in a time interval *[0,* τ] or in a shift of duration τ, the number of calls serviced by an average worker in a high load environment is $yA\tau$ and just $A\tau$ under normal time environment. Let k denote the probability that a given shift receives high call load and 1-k the probability that a shift is in a normal environment. Thus in a shift involving N workers, the expected productivity Φ is just

$$\Phi = NA\tau[k + (1-k)y] \tag{1}$$

2.2 Number of workers per shift

Suppose λ_n is the arrival rate in normal environment and $\lambda_h \equiv y\lambda_n$ the arrival rate in high load environment. The expected number of call received per shift is

$$\mathrm{M} = \lambda_n \tau[k + (1-k)y] \equiv \lambda\tau \tag{2}$$

A manager determines the optimal staffing level, the number of workers in a shift N^*, by equating expected productivity with the expected number of calls in a shift. N* is just the ratio of the expected arrival rate to the service rate and is given by $N^* = \frac{\lambda_n}{A}$. The productivity that is expected of any given worker is just $\frac{\lambda_n \tau}{N^*}$ in normal environment and $\frac{y\lambda_n \tau}{N^*}$ under heavy load respectively.

2.3 Actual productivity of an individual worker

We model the observed productivity of an individual worker as a function of expected productivity of a worker. In order to understand the effect of learning, we also incorporate a learning parameter that captures the extent to which prior experience influences worker productivity. Specifically, the observed productivity of an individual worker under normal environment is just $X_n = \delta_n \frac{\lambda\tau}{N^*}\exp(\eta)$where η is the sampling error. In our data, for every given shift we only observe $\frac{\lambda_n\tau}{N^*}$ or $\frac{\lambda_h\tau}{N^*}$. We hence, write X_n as

$$X_n = (k + (1-k)y)\delta \frac{\lambda_n\tau}{N^*}\exp(\eta_n) \tag{3}$$

The learning coefficient δ is influenced by experience under heavy work load conditions (α), experience under normal time conditions (β), individual ability (θ). More precisely we assume that $\delta_n = \exp(\theta)\,\beta^b\alpha^a$ where a and b are parameters to be estimated. With this assumption,

$$X_n = s_n\alpha^a\beta^b \frac{\lambda_n\tau}{N^*}\exp(\theta + \eta_n) \tag{4}$$

where s_n=k+(1-k)y. Likewise, the observed productivity of an individual worker under heavy load environment is just $X_h = \delta_h \frac{\lambda\tau}{N^*}\exp(\eta)$where η is the sampling error and $\delta_h = \exp(\theta)\,\beta^c\alpha^d$. Once again give that we do not observe $\frac{\lambda\tau}{N^*}$, we write X_t as

$$X_h = s_h\alpha^c\beta^d \frac{\lambda_h\tau}{N^*}\exp(\theta + \eta_h) \tag{5}$$

where $s_h = \frac{k}{y} + (1-k)$. If a worker's prior experience in both normal and heavy load environments influences worker productivity similarly, then a=b=c=d. Alternatively, if prior experience is non-fungible across work environments, then a > c and b < d.

2.4 Errors

We model errors as a proportion of total number of calls handled by a worker. If φ is the probability that a worker commits an error, then observed errors, Y_n in a normal environment is given by

$$Y_n = \psi_n\, X_n^{\mu}\exp(\epsilon_n) \tag{6}$$

where, ε is the sampling error and μ is the rate at which productivity affects error rates. As with productivity, we assume that the nature of prior experience affects error rates such that $\psi_n = \exp(\gamma)\,\beta^{b_1}\alpha^{a_1}$ where, γ denotes individual ability, while a_1 and b_1 are parameters to be estimated. Likewise the number of errors under heavy load Y_h is just

$$Y_h = \psi_h\, X_h^{\sigma} \exp(\epsilon_h) \tag{7}$$

where, ε is the sampling error and σ is the rate at which productivity affects error rates under time pressure. $\psi_h = \exp(\gamma)\,\beta^{c_1}\alpha^{d_1}$ where, γ denotes individual ability, while c_1 and d_1 are parameters to be estimated.

2.4 Estimating equations

Taking the natural log our estimating equation under normal time is

$$\ln(X_n) = \ln(s_h) + \theta + aln(\alpha) + bln(\beta) + \ln\left(\frac{\lambda_n \tau}{N^*}\right) + \eta_n \tag{8}$$

$$\ln(Y_n) = \gamma + b_1 \ln(\beta) + a_1 \ln(\alpha) + \mu \ln(X_n) + \epsilon_n \tag{9}$$

And under high load is

$$\ln(X_h) = \ln(s_n) + \theta + cln(\alpha) + dln(\beta) + \ln\left(\frac{\lambda_h \tau}{N^*}\right) + \eta_h \tag{10}$$

$$\ln(Y_h) = \gamma + c_1 \ln(\beta) + d_1 \ln(\alpha) + \sigma \ln(X_h) + \epsilon_n \tag{11}$$

If h denotes a high load environment, then equations (8) and (10) can be combined as the following system of equations.

$$ln(X) = \ln(s) + \theta + aln(\alpha) + h(c-a)ln(\alpha) + bln(\beta) + h(d-b)\ln(\beta) + \ln(Z) + \eta \tag{12}$$

$$ln(Y) = \gamma + \theta + a_1 ln(\alpha) + h(c_1 - a_1)ln(\alpha) + b_1 ln(\beta) + h(d_1 - b_1)\ln(\beta) + \eta \tag{13}$$

where $X \equiv (1-h)ln(X_n) + h\ln(X_h)$, $\ln(Z) \equiv (1-h)\ln\left(\frac{\lambda_n \tau}{N^*}\right) + h\ln\left(\frac{\lambda_h \tau}{N^*}\right)$, $\eta \equiv h\eta_h + (1-h)\eta_n$ and $\epsilon = \mathrm{h}\epsilon_{\mathrm{h}} + (1-h)\epsilon_n$

We use an instrumental variable regression using exclusion restrictions through ln(Z), the expected efficiency, which affect productivity but not error rate. Our use of expected productivity as an instrument is similar to Gowrisankaran et al (2003). It is worthwhile to reemphasize that productivity and error rates are simultaneously determined, and therefore estimating just one and not the other is likely to lead to biased inference.

EMPIRICAL SETTING, DATA AND ESTIMATION METHODS

We acquired data related to call center operators from India's largest emergency management firm. GVK EMRI (Emergency Management and Research Institute) is a leading emergency management services provider in India. It is a not for profit professional organization operating in the Public Private Partnership mode. EMRI handles medical, police and fire emergencies through the "1-0-8 Emergency service". Victims call the 108 number and are routed according to their requests for medical, police or fire emergencies. The EMRI service dispatches an ambulance to the emergency site to provide first-aid and transport victims to the nearest hospital. The services are offered free to victims, while every response is paid a fixed rate by the respective state governments with which EMRI has an agreement. Our data is restricted to medical emergencies.

The EMRI service is delivered through state-of-art call centers and has over 2600 ambulances across several Indian states. We received data from EMRI's operations in the state of Andhra Pradesh, the oldest and the largest of EMRI operations. When a victim calls the EMRI service a contact officer (CO) attends the call and enters basic information in a database. The CO also screens out prank calls and non-emergency calls. Once the CO has entered a case, the caller is routed by an automated switching network and is placed in queue for a Dispatch Officer (DO), who gathers further information and actually dispatches an ambulance to the patient.

It is important to note that it is not possible to control whether a specific DO receives specific calls. The Nortel switch automatically assigns the next call in the queue to the next free DO; therefore there is no confounding effect from selection that bedevils other learning curve studies at the individual level. In our case, more complex cases are not preferentially assigned to more competent agents.

The role of the DO is crucial in providing effective services to victims. The DO receives the basic information entered by the CO and takes the following further actions: (1) requests further information from the caller regarding the emergency; (2) identifies the caller location; (3) identifies the nearest available ambulance; (4) provides an estimated pick-up time by the ambulance and receives confirmation from the victim that this is acceptable; (5) in the case of acute emergencies patches an attendant physician to the caller and to the paramedic in the ambulance to provide first aid and any other specific instructions to care for the victims and (6) weeds out any prank calls not identified by the CO. The DO ends the call when (a) an ambulance is assigned, (b) receives confirmation from the ambulance driver and (c) provides an estimated arrival time to the patient.

The DO role requires considerable judgment – regarding the severity of the incident and the speed with which an ambulance needs to reach the scene, understanding the nearest available ambulance and working with the physician and paramedic as necessary. For example, the geographically closest ambulance may not be available or it may not be the ambulance that can most easily reach the victim, given traffic and infrastructure conditions. The DO has to take such considerations into account in assigning the ambulance. The DO needs to assess the severity of the case in deciding whether to patch the physician or provide special instructions to the paramedic in the ambulance. Such considerations make the DO's job amenable to significant learning based on experience.

Data and Variables

We ran regressions estimating the productivity and error rates of new DO's who joined the organization in 2010, for whom we have complete history of the calls they attended. Our first dependent variable is the actual productivity of the DO (X), measured as the number of calls the DO attended in a given shift. We estimate learning by assessing whether prior experience leads to an increase in DO productivity after controlling for the complexity of the call. The ambulance unit gets a variety of calls; some calls such as severe cardiac or poisoning cases are complex, where the DO has to patch in a physician with the paramedic in the ambulance to provide emergency first-aid before the patient is shifted. Other calls such as routine pregnancy cases simply involve assigning an ambulance to the patient. Our measure of DO productivity is weighted[1] by the complexity of cases serviced by the DO.

Our second dependent variable is the number of errors made by the DO in the current shift, weighted for call complexity. A senior experienced officer follows up on every call and codes whether the DO made the appropriate decision. We have data on both type I and type II errors – i.e., an ambulance should be dispatched but was not (false-negatives), and an ambulance need not be dispatched, but was (false-positives). The government carries out a periodic audit of its service providers including EMRI in order to decide on contract renewal and negotiate rates.

[1] Details regarding how we weighted the calls are available from the authors. Note that our results are robust to using un-weighted data as well.

Organizations with minimal error rates have a significant bargaining advantage with the government; therefore error rates are closely tracked by the firm.

Our primary independent variable is the agent's prior experience coded as the number of calls prior to the current shift that the DO has received in this job. In order to understand the impact of heavy workload conditions, we split prior job experience in two categories: number of calls operating under normal conditions, β (measured as ln(0.1+number of calls)) and number of calls working under heavy workload or time pressure conditions, α (measured as ln(0.1+number of calls)).

We coded heavy-workload shifts as those where the average calls per agent in the shift was higher than the estimated workload. To compute this variable, we first calculated the expected workload in any given shift based on data until the current shift. Presumably, the shift manager decides on staffing patterns based on estimated workload. If the actual number of calls received in the shift was higher than this estimated workload, we coded the shift as a heavy-workload shift.

Alternatively, for each shift we calculated the queue volumes of the number of calls waiting or agents to attend. Any shift that had a queue length above median (alternatively mean) was coded as a heavy workload or time pressure shift. This coding the advantage that it is endogenous to the workload faced by the peer group in any shift, automatically controls for different staffing patterns in any shift and is exogenous to agent characteristics.

Control Variables: To observe learning, we run within-subject regressions, using agent fixed effects to control for unobserved agent characteristics. We control 51 dummy variables for the week of the year, 6 dummy variables for the day of the week, and four dummy variables for shift (graveyard shift is the left out category). As instruments for the stage 2 regression, we also control for expected productivity by coding the prior history of the agents.

RESULTS

Table 1 shows the descriptive statistics for our data, and table 2 shows the correlation matrix. From table 2, we see that as expected, the number of errors is highly positively correlated with the productivity of the agents. The faster the agents service incoming calls, the more errors they commit. We also see that the experience under normal time and experience under heavy workload shifts are positively correlated, probably because the longer the agents have worked in the organization, the more likely they have served in both normal shifts and peak load shifts. Interestingly, we also see that both types of experience are positively correlated to both productivity and to errors. This suggests that experience mainly leads to productivity (speed) gains in our sample, which likely in turn drive increased error rates. In order to understand the contingent relationships between these variables, we now turn to regression analyses.

Table 3, specification 1 shows the main effect of the different types of experience on productivity and errors in an IV regression. Looking at the first stage in specification 1, we see that both types of experience lead to increased productivity among the agents. We also see that when the shift is a heavy work-load shift, agents work faster; the dummy variable is positive and significant. Our instrument, expected efficiency is also positive and significant. In the second stage, again as expected we see that higher the productivity (speed), the higher the error rates. Also, in heavy-workload shifts, the error rates increase. Finally, from the second stage specification, we also see that increasing experience, with either type of shift, leads to reduction in error rates. In fact, the coefficients are almost the same size, suggesting that the impact of experience in either kind of shift leads to about the same level reduction in error rates.

In specification 2 in table 3, we estimate how fungible the experience gathered from one operating environment is to working in another environment. For this purpose, we introduce two interaction effects – interacting the extent of experience in either type of regime with whether the focal shift is a heavy load shift. In the first-stage productivity regression, we see evidence of non-fungibility: whereas experience in heavy load environment allows individuals to be faster in the heavy load regime, increasing normal experience makes them slower in the heavy load regime. We also see that increasing normal experience leads the agent to committing more errors in the heavy workload regime, whereas increasing experience in that regime has not impact on errors.

Table 3a interprets the interaction effect in the 2x2 matrix (working in 2 types of regimes x increasing experience in these two types of regimes). We again see evidence for specialization. In the normal regime, only increasing normal experience has an effect on productivity; increasing experience in the heavy load regime has no effect. On the other hand, increasing normal experience has a more beneficial effect on subsequent work in the normal regime than in the heavy load regime. Similarly, increasing experience in the heavy load regime only impacts productivity in the heavy load regime; it has no impact in the normal regime. Turning to errors, we see that increasing experience in both kinds of regimes reduce errors committed in the normal regime. However, unit increase in normal experience has a much greater impact on subsequent errors in the normal regime than in the heavy load regime. On the other hand, experience gained in the heavy load regime has about the same impact on error rates in both these regimes. These results show that the experience gained in different types of work environments are not fungible.

Robustness Checks

We performed a number of robustness on our results. Before, we detail these checks it is worth reiterating two unique features of the sample. First, we have complete information on the population of new agents and the calls attended by these agents for one year in this organization. Therefore, our analyses are conducted on this population and not on a sub-sample. Second, incoming calls are assigned to agents by the technology employed on a FCFS basis. More complex calls are not assigned to more competent agents; therefore, we do not need to account for selection in our estimation.

Our sample consists of agents who are full time employees of the contact center that gave us the data as well as agents who are “trainees”. Trainees are either undergoing assessment for future employment or are members of other organizations who are getting trained in this center. As the first robustness check, though we estimate within-agent regressions, we wished to understand whether our results are driven by differences in the nature of agents. Table 4 replicates table 3 excluding these trainees from the sample. Our results are qualitatively identical.

Second, we changed the measure of our key independent variable – whether a given shift is under heavy load. We used an alternative measure and defined a heavy load shift as one that has more than the median or mean number of calls per shift. Table 5 shows these results; they are qualitatively similar to our baseline specifications. In the baseline we weight calls by their severity. For example, longer time taken to assign an ambulance or an error committed in a call involving a heart-attack victim is more serious than that involving a victim who has an ache in their leg. As a third robustness check, we estimated learning using un-weighted call times and error rates. Table 6 shows these results. Again, our main conclusions are robust to this check.

Finally, we seek to understand whether the non-fungibility of experience across different work environments is mainly pronounced in different types of errors committed by the agents. In table 7 shows the effect of different experience and work regimes on the number of false positive errors and false negative errors made by the agent. From table 7, we see that the non-fungibility is mostly visible in the false negatives. It appears that agents who have a lot of experience with the normal regime tend to deny an ambulance to genuinely needy cases when they are in the heavy workload condition. This may happen because these agents skip over some of their normal verification checks when they are working in the high load environment and this leads them to misclassify genuine, but potentially non-urgent cases as not worthy of service.

DISCUSSION

Recent research in the learning curve tradition attempts to understand why firms performing the same activities have very different learning curves (Argote et al, 2003). Prior research has suggested that learning curves may vary because of scope economies or differences in turnover between firms (Dutton and Thomas, 1984; Argote, 1999). If learning by doing is one of the most important means for firms to develop capabilities (Arrow, 1962; Nelson and Winter, 1982), and capability development is associated with competitive advantage (Teece, Pisano and Shuen, 1997; Teece, 2007; Felin et al, 2012), then understanding the factors that underlie the learning curve is important to understand the bases for competitive advantage. Recent work has therefore concentrated on understanding the factors that moderate the relationship between cumulative volume and productivity at the individual, group and organizational level of analysis (Hauschild and Sullivan, 2002; Stan and Vermeulen, 2012; Huckman and Pisano, 2006; Boh et al, 2007; Pisano et al, 2001; Huckman et al, 2009). It is from this perspective that we make novel contributions.

Reagans et al (2005) argue that the learning curve is influenced by two broad categories of factors: factors that influence individual learning and factors that improve coordination within the organization. We show that specialization at the individual level leads to faster learning; in our context learning in one work context – the normal work environment in our setting - is not fungible to execution in another environment – the heavy load setting, and vice-versa.

Our study has a few strengths in showing the non-fungibility of learning from experience at the individual level when compared to other studies in this domain. Huckman and Pisano (2006) showed that physician's experience in the same hospital is associated with better patient outcomes, but experience with other hospitals is not. They explain this by suggesting that physician's skills are specific to the assets available it he hospital. Potentially this specificity could include experience working with other staff in these hospitals, though that is not explicitly shown (also see Pisano et al, 2001; Edmondson et al, 2001). In contrast to these (mis)coordination based explanations, our study is more learning focused. We show non-fungibility in the same organization performing the same task but in different environments (normal vs. heavy load). Boh et al (2007) show that software engineers productivity increases with specialization, but the group's productivity increases with increasing task variety at the individual level (also see Narayana et al, 2009). Again, these explanations suggest that organizational learning places a premium of coordination ability, even at the expense of individual ability. Our results add further nuance; reduced specialization at the individual level can lead to poor outcomes for organizations when coordination is not an important consideration. The mechanism we identify, the non-fungibility of learning across different work contexts is also a novel explanation for the persistence of heterogeneity in learning curves between firms.

The principal implication of this result is that learning curves can be sensitive to employee allocation across different work environments. Our findings suggest that the rate of individual learning is influenced by matching individual experience to different work contexts. We show that individuals with more experience working in normal conditions reduce errors when working normal conditions as well as while working in heavy workload conditions. On the other hand, individuals with experience working in heavy workload conditions are more effective in heavy work load conditions when compared to normal conditions. This suggests that matching employees to work contexts according to their type of experience allows firms to walk down the learning curve faster than if they did not take these considerations into account.

However, we should note that employee experience is a function of which work contexts they worked in. A manager can influence experience as well as learning by assigning employees systematically to different kinds of work contexts in order to develop an experience and learning profile. For instance, our results suggest that managers can influence individual learning curves by assigning individuals to tasks differentially. To the extent that learning from experience is not completely fungible across different learning situations and application contexts, managers' ability to assign individuals across different situations can be important to optimally develop individual abilities. To the extent that such assignment significantly affects individual learning and is important for carrying out the firm's business, it impacts firms' capability development and competitive advantage.

One of the fundamental questions for strategy scholars is to explain the bases for firm heterogeneity (Barney, 1991; Peteraf, 1996). Our results suggest that managers' assignment of employees to work contexts is a mechanism by which managers contribute to firm heterogeneity that arises from differing patterns of capability development. Though such assignment patterns may be codified within the firm, they may not be immediately obvious to competitors, nor their impact on learning rates be identified within the firm. Insofar as the role of or the process of assignment is unknown to competitors, these could lead to causal ambiguity in capability development.

Another interesting insight from our paper is the importance of trade-offs in strategy. Porter (1980; 1990) argues that competitive advantage rests on tight integration of activities along the value chain, and to the extent that other organizations cannot replicate such tight integration. In this regard he suggests that firms have to choose between efficiency and flexibility. Efficient organizations have little slack and lower costs; but it comes at the cost of responsiveness, which implies the need t maintain higher slack levels. For example, a factory can be optimized to produce one type of product (Ford Model T), or it could be flexible to shift with changing demand conditions at the necessary cost of running the plant at less than full capacity.

Our work adds a novel mechanism to understanding the trade-off between efficiency and flexibility. A flexible firm should cross-train its employees in a variety of tasks to retain the option of allocating them to different needs as and when they arise. However, this comes at the necessary cost of lower productivity of these workers in all these tasks. Our results suggest that learning from experience in one type of environment reduces effectiveness in another environment. This suggests that unless there are significant economies of scope from maintaining these different tasks within the same firm, specialist firms are more likely to out-compete these generalists because of their superior knowledge base rather than because of superior economies of scale.

Our third contribution is econometric. We argue and show that studies of learning at the individual level must control for the simultaneity of productivity and quality. Treating either variable by itself, even when controlling for the other, is likely to lead to erroneous results. More alarmingly, as we show, the normative implications arising from these results can be very different. Depending on the setting, this could have interesting public policy implications. For example, based on volume-outcome relationships in surgery, some states have laws for minimal volumes that a hospital must perform in order to offer certain surgical procedures; insurance providers cover certain procedures only in hospitals with certain levels of minimal volumes, which have competitive and anti-trust implications. Our specific investigation of the impact of level of workload on learning is also subject to regulation. For example, California enacted a law recommending a certain minimal nurse-patient ratio and New York enacted a law restricting the number of hours a resident may work in a week. Considering factors such as workload on surgeons, the trade-off between quality and efficiency can change these normative implications. This may be interesting future work in health economics.

Similar to all other papers, our work is also subject to limitations. Our major limitation is imposed by the type of data we have access to. Our data is from a single organization, though at a great level of detail at the individual level. Therefore, while we can measure learning at the individual level, we can only argue logically, rather than empirically, about organizational heterogeneity. However, our mechanism is novel, relating an individual level process to an organization-level outcome, and it is reasonable to suppose that the aggregation process, given our results at the individual level will lead to the proposed organizational outcome. Our data structure also has some interesting strengths, especially the random assignment of tasks to individuals that effectively controls for selection problem that bedevils studies of heterogeneity in individual learning curves. Second, our data is limited to one year, where we have complete data on individual work experience and outcomes. Therefore, the strongest identification comes from employees who newly joined this organization in this role in that year. We rely on individual fixed effects to control for any heterogeneity in agent characteristics.

In conclusion, we show that individual learning curves within an organization are significantly influenced by nature of experience and type of work context. This suggests that if a manager can effectively assign employees by matching their experience to tasks, then they could effectively improve learning in their organization. Since individual capability development results in improved productivity and quality, such matching could be a source of competitive advantage for the firm. This is one mechanism by which manager's can impact firm capability development that results in firm heterogeneity.

Table 1: Descriptive Statistics

Variable		Source of variation	N	Mean	Std. Dev.
Errors	# false positives and false negatives of a DO in a shift weighted by severity	DO, shift	15627	2.25	1.11
Productivity	Log actual number of calls attended by a DO in a shift weighted by severity	DO, shift	15627	2.83	1.23
Expected Efficiency	Log total number of calls received in a shift divided by number of DO in a shift	DO, shift	15627	3.32	0.25
Normal experience	Log number of cumulative calls of a DO in normal time	DO, shift	15627	5.26	2.00
High load experience	Log number of cumulative calls of a DO in time pressure	DO, shift	15627	5.29	2.09
High load	=1 if the current shift is under high load environment	Shift	15627	0.51	0.50
Department Transfers	=1 if DO was working in another department prior to becoming DO	DO	15627	0.69	0.46
Trainees	=1 if DO was a trainee	DO	15627	0.18	0.39
Observed Productivity(Raw)	Unweighted # of number of calls attended by a DO in a shift	DO, shift	15627	2.67	1.24
Errors (Raw)	Unweighted # false positives and false negatives of a DO in a shift	DO, shift	15627	2.09	0.99

Table 2: Correlation table

		1	2	3	4	5	6
Errors	1	1.00					
Productivity	2	0.93*	1.00				
Expected Efficiency	3	0.10*	0.14*	1.00			
Normal experience (β)	4	0.17*	0.27*	0.01	1.00		
High load experience (α)	5	0.14*	0.23*	0.08*	0.81*	1.00	
High load	6	0.08*	0.05*	0.32*	-0.02*	0.07*	1.00

Table 3: Estimating the main-effect of experience in normal vs. high load environment on productivity and errors (weighted)

	Specification 1		Specification 2	
	Errors	Productivity	Errors	Productivity
Productivity	0.81***		0.81***	
	(0.018)		(0.018)	
Normal experience (β)	-0.016***	0.071***	-0.021***	0.094***
	(0.004)	(0.012)	(0.005)	(0.014)
High load experience (α)	-0.016***	0.022**	-0.013***	-0.008
	(0.004)	(0.010)	(0.004)	(0.012)
Normal exp* High load			0.011**	-0.053***
			(0.005)	(0.016)
High load exp*High load			-0.006	0.072***
			(0.005)	(0.015)
High load	0.054***	0.044**	0.028	-0.062
	(0.008)	(0.022)	(0.019)	(0.056)
Expected efficiency		1.16***		1.165***
		(0.062)		(0.062)
Constant	-0.39	-3.88***	-0.365	-3.808***
	(0.359)	(1.085)	(0.359)	(1.086)
Week of Year Fixed Effects (51)	Y	Y	Y	Y
Day of Week Fixed Effects(6)	Y	Y	Y	Y
Shift Fixed Effects(3)	Y	Y	Y	Y
DO Fixed Effects	360	360	360	360
Observations	15,627	15,627	15,627	15,627
R-squared	0.898	0.293	0.898	0.294
Adjusted R-Squared	0.89	0.27	0.89	0.27
Wald Chi Squared	39787		39877	
F		14.83		14.84

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

Table 3a:

Errors

		Experience	
		Normal	High Load
Regime	Normal	-0.021***	-0.013***
		(0.005)	(0.004)
	High Load	-0.009*	-0.02***
		(0.005)	(0.005)

Productivity

		Experience	
		Normal	High Load
Regime	Normal	0.094***	-0.008
		(0.014)	(0.012)
	High Load	0.042***	0.064***
		(0.015)	(0.014)

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

Table 4: Estimating the spillover-effect of experience in normal vs. high load environment on productivity and errors (sub-sample after removing trainees)

	Specification 1		Specification 2	
	Errors	Productivity	Errors	Productivity
Productivity	0.809***		0.808***	
	(0.018)		(0.018)	
Normal experience (β)	-0.012**	0.070***	-0.016***	0.088***
	(0.005)	(0.014)	(0.006)	(0.016)
High load experience (α)	-0.020***	0.005	-0.018***	-0.019
	(0.004)	(0.012)	(0.005)	(0.013)
Normal exp* High Load.			0.011*	-0.045**
			(0.006)	(0.018)
High load exp*High Load			-0.005	0.063***
			(0.006)	(0.017)
High Load	0.056***	0.035	0.022	-0.065
	(0.008)	(0.024)	(0.021)	(0.061)
Expected efficiency		1.257***		1.260***
		(0.068)		(0.068)
Constant	0.58	-5.44***	0.596	-5.358***
	(0.365)	(1.098)	(0.365)	(1.099)
Week of Year Fixed Effects (51)	Y	Y	Y	Y
Day of Week Fixed Effects(6)	Y	Y	Y	Y
Shift Fixed Effects(3)	Y	Y	Y	Y
DO Fixed Effects	274	274	274	274
Observations	12,743	12,743	12,743	12,743
R-squared	0.897	0.308	0.897	0.309
Adjusted R-Squared	0.89	0.29	0.89	0.29
Wald Chi Squared	32100		32172	
F		16.34		16.30

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

Table 5: Estimating the effect of experience in normal time vs. time pressure on productivity and errors (weighted): Alternative measure of time pressure

	Specification 1		Specification 2	
	Errors	Productivity	Errors	Productivity
Productivity	0.865***		0.865***	
	(0.016)		(0.016)	
Normal experience (β)	-0.024***	0.033***	-0.027***	0.041***
	(0.004)	(0.011)	(0.004)	(0.011)
High load experience (α)	-0.019***	0.080***	-0.020***	0.061***
	(0.004)	(0.012)	(0.004)	(0.012)
Normal exp* High Load.			0.019***	-0.052***
			(0.005)	(0.015)
High load exp*High Load			-0.002	0.077***
			(0.006)	(0.016)
High Load	0.021**	0.091***	-0.071***	0.042
	(0.010)	(0.029)	(0.022)	(0.066)
Expected efficiency		1.182***		1.177***
		(0.057)		(0.057)
Constant	-0.399	-3.922***	-0.316	-3.871***
	(0.357)	(1.081)	(0.357)	(1.082)
Week of the Year Fixed Effects(51)	Y	Y	Y	Y
Day of Week Fixed Effects(6)	Y	Y	Y	Y
Shift Fixed Effects(3)	Y	Y	Y	Y
DO Fixed Effects	360	360	360	360
Observations	15,627	15,627	15,627	15,627
R-squared	0.899	0.294	0.899	0.295
Adjusted R-Squared	0.90	0.27	0.90	0.28
Wald Chi Squared	40301.64		40525.99	
F		14.94		14.94

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

Table 6: Estimating the effect of experience in normal time vs. time pressure on productivity and errors (un-weighted)

	Specification 1		Specification 2	
	Errors	Productivity	Errors	Productivity
Productivity	0.833***		0.832***	
	(0.014)		(0.014)	
Normal experience (β)	-0.011***	0.110***	-0.015***	0.139***
	(0.004)	(0.011)	(0.005)	(0.013)
High load experience (α)	-0.012***	0.051***	-0.009***	0.015
	(0.003)	(0.010)	(0.004)	(0.011)
Normal exp* High Load.			0.010**	-0.063***
			(0.005)	(0.015)
High load exp*High Load			-0.005	0.089***
			(0.005)	(0.014)
High Load	0.045***	-0.001	0.015	-0.145***
	(0.006)	(0.020)	(0.016)	(0.052)
Expected efficiency		1.288***		1.300***
		(0.058)		(0.058)
Constant	-0.260	-3.539***	-0.229	-3.436***
	(0.316)	(1.010)	(0.316)	(1.010)
Week of the Year Fixed Effects(51)	Y	Y	Y	Y
Day of Week Fixed Effects(6)	Y	Y	Y	Y
Shift Fixed Effects(3)	Y	Y	Y	Y
DO Fixed Effects	360	360	360	360
Observations	15,627	15,627	15,627	15,627
R-squared	0.900	0.393	0.900	0.394
Adjusted R-Squared	0.90	0.38	0.90	0.38
Wald Chi Squared	55004.97		55397.28	
F		23.17		23.22

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

Table 7: Estimating the effect of experience in normal time vs. time pressure on productivity and errors (weighted): false positives and false negatives

	FALSE POSITIVES Specification 1		FALSE NEGATIVES Specification 1	
	Errors	Productivity	Errors	Productivity
Productivity	0.732***		0.785***	
	(0.028)		(0.020)	
Normal experience (β)	-0.026***	0.094***	-0.026***	0.094***
	(0.008)	(0.014)	(0.006)	(0.014)
High load experience (α)	-0.007	-0.008	-0.014***	-0.008
	(0.006)	(0.012)	(0.005)	(0.012)
Normal exp* High Load.	0.011	-0.053***	0.015**	-0.053***
	(0.008)	(0.016)	(0.006)	(0.016)
High load exp*High Load	-0.012	0.072***	-0.007	0.072***
	(0.008)	(0.015)	(0.006)	(0.015)
High Load	0.061**	-0.062	0.028	-0.062
	(0.029)	(0.056)	(0.021)	(0.056)
Expected efficiency		1.165***		1.165***
		(0.062)		(0.062)
Constant	-0.933*	-3.808***	-0.500	-3.808***
	(0.554)	(1.086)	(0.399)	(1.086)
Week Fixed Effects (51)	Y	Y	Y	Y
Day Fixed Effects(6)	Y	Y	Y	Y
Shift Fixed Effects(3)	Y	Y	Y	Y
DO Fixed Effects	360	360	360	360
Observations	15,627	15,627	15,627	15,627
R-squared	0.753	0.294	0.872	0.294
Adjusted R-Squared	0.75	0.27	0.87	0.27
Wald Chi Squared	13861		32310	
F		14.84		14.84

Notes: Standard errors in parenthesis. * Significant at 90% level. ** Significant at 95% level. *** Significance at 99%.

REFERENCES

Adler, P. S., K. B. Clark. 1991. Behind the learning curve: A sketch of the learning process. Management Sci. 37 267–281.

Argote, Linda. "Group and organizational learning curves: Individual, system and environmental components." British journal of social psychology 32.1 (1993): 31-51.

Argote, Linda. Organizational learning: Creating, retaining and transferring knowledge. Springer, 1999.

Argote, Linda, Bill McEvily, and Ray Reagans. (2003) "Managing knowledge in organizations: An integrative framework and review of emerging themes." Management science 49.4: 571-582.

Besanko, D., Doraszelski, U., Kryukov, Y., & Satterthwaite, M. (2010). Learning-by-Doing, Organizational Forgetting, and Industry Dynamics. Econometrica, 78(2), 453-508.

Dutton, John M., and Annie Thomas (1984). "Treating progress functions as a managerial opportunity." Academy of Management Review: 235-247.

Epple, Dennis, Linda Argote, and Rukmini Devadas. (1991) "Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing." Organization Science 2.1 58-70.

Epple, Dennis, Linda Argote, and Kenneth Murphy. (1996) "An empirical investigation of the microstructure of knowledge acquisition and transfer through learning by doing." Operations Research 44.1: 77-86.

Evans, William N., and Beomsoo Kim. (2006) "Patient outcomes when hospitals experience a surge in admissions." Journal of Health Economics 25.2: 365-388.

Fillmore, Ian, and Devin G. Pope. (2012) The Impact of Time Between Cognitive Tasks on Performance: Evidence from Advanced Placement Exams. No. w18436. National Bureau of Economic Research,.

Fong Boh, Wai, Sandra A. Slaughter, and J. Alberto Espinosa. (2007) "Learning from experience in software development: A multilevel analysis." Management Science 53.8: 1315-1331.

Groysberg, Boris, Linda-Eling Lee, and Ashish Nanda. (2008): "Can they take it with them? The portability of star knowledge workers' performance." Management Science 54.7 1213-1230.

Haunschild, Pamela R., and Bilian Ni Sullivan. (2002) "Learning from complexity: Effects of prior accidents and incidents on airlines' learning." Administrative Science Quarterly 47.4 609-643.

Huckman, Robert S., and Gary P. Pisano. (2006) "The firm specificity of individual performance: Evidence from cardiac surgery." Management Science 52.4 473-488.

Ittner, Christopher D., Venky Nagar, and Madhav V. Rajan. (2001) "An empirical examination of dynamic quality-based learning models." Management Science 47.4 563-578.

Kc, Diwas S., and Christian Terwiesch. (2009) "Impact of workload on service time and patient safety: An econometric analysis of hospital operations." Management Science 55.9 1486-1498.

Kim, Youngsoo, Ramayya Krishnan, and Linda Argote. (2012) "The Learning Curve of IT Knowledge Workers in a Computing Call Center." Information Systems Research 23.3-Part-2 887-902.

Lapré, Michael A., and Luk N. Van Wassenhove. (2001) "Creating and transferring knowledge for productivity improvement in factories." Management Science 47.10 1311-1325.

Lieberman, Marvin B. (1984) "The learning curve and pricing in the chemical processing industries." The RAND Journal of Economics 15.2: 213-228.Lieberman, Marvin B. "The learning curve, diffusion, and competitive strategy." Strategic management journal 8.5 (1987): 441-452.

Narayanan, Sriram, Sridhar Balasubramanian, and Jayashankar M. Swaminathan. (2009) "A matter of balance: Specialization, task variety, and individual learning in a software maintenance environment." Management Science 55.11: 1861-1876.

Pisano, G. P. (1994) Knowledge, integration, and the locus of learning: An empirical analysis of process development. Strategic Management Journal. 15 85–100.

Poulose, Benjamin K., Wayne A. Ray, Patrick G. Arbogast, Jack Needleman, Peter I. Buerhaus, Marie R. Griffin, Naji N. Abumrad, R. Daniel Beauchamp, and Michael D. Holzman. (2005) "Resident work hour limits and patient safety." Annals of surgery 241, no. 6: 847.

Reagans, Ray, Linda Argote, and Daria Brooks. (2005) "Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together." Management Science 51.6: 869-881.

Scott, L. D., Rogers, A. E., Hwang, W. T., & Zhang, Y. (2006). Effects of critical care nurses' work hours on vigilance and patients' safety. American Journal of Critical Care, 15(1), 30-37.

Terwiesch, Christian, and Roger E Bohn. (2001) "Learning and process improvement during production ramp-up." International Journal of Production Economics 70(1), 1-19.