Managing Software Implementers in the Information Services Industry:
An Example of the Impact of Market Growth on Knowledge Worker Productivity and Quality

Edward G. Anderson Jr.
University of Texas Department of Management
Austin, Texas

May 1, 1998

Abstract

This paper uses control theory techniques to solve the staffing problem for knowledge workers under continuous demand growth. The novelty of this particular formulation of the staffing problem lies in the use of time-discounting in combination with a special cost structure. The resulting formulation was solved for an optimal policy which reflects the influence of both shortfall and salary penalties. Under real-world conditions, this policy quickly evolves to a steady-state solution which is stationary in a proportional sense to evolving capacity requirements. An illustration of the policy in the setting of software implementers of enterprise requirement planning systems is developed.

Additionally, the paper will examine some of the implications of this model for a typical implementation firm. In particular, the paper will show that under many real-world operating conditions, the interaction of knowledge workers and market growth can drive firms to deliver lower levels of service at a higher cost. Finally, the paper puts this research into a framework with the experience-curve and technology supply-chain literature to outline possible directions for future research.

1 MOTIVATION

It is estimated that currently over half of U.S. firms' capital spending is related to information technology. Much of this has been driven by the explosive growth of all-encompassing enterprise resource planning (ERP) software packages such as SAP, PeopleSoft, and Baan. The majority of these implementations seem to be the price paid for fixing the "millenium bug" in which software systems mistake the year 2000—recorded in most older systems as "00"—for 1900 (Gross 1997). However, while the millenium bug might be a disaster, the cure has inflicted its own pains on customer firms. Due to the highly complex nature of these systems, ERP system installation requires a veritable small army of experts—known as implementers—to customize the software to the customer's business processes. Very often ERP implementations go awry. One senior information technology consultant commented that installing SAP was much like a war: "It costs a lot; it tends to escalate; and a lot of people get shipped home in body-bags."[1] While there are many reasons for the difficulties involved in these implementations, one contributor has been the inexperience of the implementation workforce. Unlike traditional software, which is either "shrink-wrapped" or custom-built by the firm in question, installation of ERP systems requires expert knowledge that a customer firm would not normally possess. These outside expert require, in addition to some general business process experience, a deep knowledge of a particular software package's quirks. Developing this expertise is a lengthy process, sometimes requiring over a year of course-work and on-the-job training. In a market growth situation, the ranks of experienced implementers will necessarily be thin. Hence, training these knowledge workers will become a major expense for the implementation company. Because of the time lags involved, it has not made a great deal of sense for the customer firms to develop their own implementers. Particularly, as these employees, once trained, will find themselves more valuable to outside consulting firms than to the customer firm that trained them (Parker 1997). Until recently the software vendors themselves have also

[1] This comment as well as much of the other ERP industry information in this paper comes primarily from interviews with Mr. Robert Keating, a senior business analyst with Computer Sciences Corporation (CSC) Consulting Group of Cambridge, Massachusetts. Additional information was provided by interviews with Mr. Steven Baggette, a senior programmer-analyst with International Business Machines of Southfield, Michigan. Any errors in this paper are the fault of the author.

remained out of this business, leaving installation to third party implementation firms such as consultancies like Computer Sciences Corporation (CSC), Price-Waterhouse, or any of a number of small specialist firms. Managing the hiring and training of implementers in advance of double- and sometimes triple-digit annual growth rates is a complicated dynamic problem, mismanagement of which can be fatal to the implementation firm.

Accordingly, this paper will use dynamic programming techniques to solve analytically the staffing problem for knowledge workers under continuous demand growth. The novelty of this particular formulation of the staffing problem lies in the use of time-discounting in combination with a special cost structure. This allows the development of a simple, steady-state solution which is stationary in a proportional sense to evolving requirements. Additionally, the paper will examine some of the implications of this model for a typical implementation firm. In particular, we will show that under many real-world operating conditions, the interaction of knowledge workers and market growth can drive implementing firms to deliver lower levels of service at a higher cost. Finally, we will put this research into a framework with the more traditional experience-curve literature to outline some possible directions for future research.

2 The Problem

In this paper, we build upon many previous works in the field of staffing problems. From a simulation perspective, Sterman's (1988) analysis of service quality at People Express Airlines studies issues quite similar to those investigated in this paper. From the analytic side, there have historically been two approaches to optimizing staffing problems that are differentiated by their objectives. One approach typified by Bartholemew (1973) is to guide an organization to a desired steady-state structure of groups of individuals by pay-grades and experience-levels. The other approach traces back to Holt et al. (1960), which optimizes personal planning with respect to some set of production requirements and costs. However, the Holt et al. or "HMSS" (so-name for the book's four authors) model used a single constant employee productivity in their groundbreaking study and, hence, could not represent individual or group learning. Ebert (1976) modified the "HMMS" model to include an experience curve increasing productivity with cumulative output. This approach, while appropriate for learning embodied in capital or organizational processes, does not capture the effects of learning by *individual* employees. Others, such as Grinold (1977), used Markovian flow models to differentiate the effectiveness of workers by experience and grade-level. Gaimon and Thompson (1984) extended the art significantly by using non-linear contributions by various employee groups to the production function. The bulk of these formulations have, however, been designed to optimize under conditions of either finite horizons or limited growth. In contrast this model will sacrifice some complexity relative to previous works in order to grapple effectively with the special demands of managing long-term growth.

Typically, implementers of enterprise requirement systems such as SAP require in addition to their schooling several months of in-class training time followed by up to a year of on the job training. Accordingly, we will postulate two classes of employees for the implementation firm. Experienced employees $x(t)$ will have completed both their class-room and on-the-job training and will be fully productive. In contrast, the productivity of new employees $n(t)$ who are still undergoing training will be discounted to reflect their lesser productivity (Equation 1). In fact, in many instances their effective productivity will be negative because they require mentoring from experienced employees.

$$c(t) = x(t) + \delta n(t) \tag{1}$$

where $|\delta|$ is assumed to be less than unity, and $c(t)$ represents the capacity of the firm (measured in equivalent experienced employees) to meet their customers' implementation requirements.

Following, the "fractional-flow" assumption common to most staffing papers, a constant fraction of new employees will become trained each period and move into the experienced rank. The inverse of this fractional training rate, that is τ, can be interpreted as the length of time required for training the new employees (Equation 2). Similarly, a constant fraction will retire each period or leave through attritition. The parameter α can be interpreted as the average tenure time for experienced employees (Equation 2).

$$\frac{dx(t)}{dt} = \frac{n(t)}{\tau} - \frac{x(t)}{\alpha} \tag{2}$$

For simplicity's sake "raiding" gains from and losses to other companies' experienced personnel are neglected because these should tend to cancel out for the average firm.

Capacity requirements *r(t)* expressed in experienced-employee equivalents are assumed to grow exponentially:

$$r(t) = r_0 e^{\kappa t} \tag{3}$$

where κ and r_0 are assumed nonnegative. The constant ro the requirements of the firm in capacity at time $t=0$, and κ represents the instantaneous fractional requirement growth rate.

Formulating the objective function is less straightforward than the constraints. In particular, a policy that maintains stationary ratios over time between the capacity requirements and the actual capacity would seem intrinsically appealing. The standard formulation for this sort of problem in the literature following the HMMS model, however, is (neglecting hiring, firing, and promotion costs):

$$\min_{n(t)} \int_0^T \left\{\alpha\left[r(t) - c(t)\right]^2 + \beta\left[x(t) + n(t)\right]\right\} dt \tag{4}$$

This formulation minimizes the sum of two costs over a finite time interval. The first, quadratic penalty reflects the escalating costs associated with increasing capacity shortages. For instance, a small shortfall of employees can be compensated for with overtime. However, employees can only work so many hours per week, so a large employee capacity shortfall will force the implementing firm to either decrease implementation quality, lengthen implementation times, or turn down possible business. The second, linear term reflects the salary costs of both the new and experienced employees. This mixed linear-quadratic formulation is, however, problematic under sustained growth. The most important problem is that under requirement growth the quadratic shortfall penalty term necessarily increases with time relative to the linear salary term. Hence, either the quadratic shortfall term will be relatively insignificant at the interval's beginning or it will dominate near the end. A second problem is that both shortfall and salary penalties will escalate exponentially with time. Accordingly, the end requirements will exercise a relatively greater influence than the near-term's. As the cost of capital is expensive and long-term growth less certain than near-term, a different formulation seems reasonable here. The alternative used in this paper is presented below.

$$\min_{n(t)} \int_0^\infty e^{-\rho t} \left\{\zeta\left[r(t) - c(t)\right]^2 + \beta\left[(1-\delta)n(t)\right]^2\right\} dt \tag{5}$$

where ρ, ζ, and β are assumed nonnegative

The first modification of Equation 5 from Equation 4 is to add a discount factor into the penalty expression. This will discount the value of future penalties by an annual fractional rate ρ. In principle, this new objective can be solved for any time interval, however, if $\rho \geq 2\kappa$ then we can avoid end effects by integrating to infinity. The second feature of this formulation is the quadratic penalty function for the difference between the total number of workers employed—both experienced and new—and the number of experienced workers equivalents required. It seems reasonable to assume that most firms will be willing to employ at least as many employees as they would in a non-growth situation. However, any employees in excess of this will accrue salary costs which in some sense are not directly contributing to current profits. In fact, as the excess employees increase ever further, the firm's competitiveness may decline precipitously. Hence, the quadratic assumption for the salary penalty also seems reasonable.

3 CONTROL STRATEGY

Given the constraints in Equations 1 to 3, we can minimize the objective in Equation 5 by using optimal control theory (Bertsekas 1987). The resulting optimal control policy is presented below.

(Without loss of generality we will assume for the remainder of the paper that the shortfall penatly coefficient ζ is equal to unity.)

Theorem 1

The optimal solution to the problem described above will be of the form:

$$x(t) = Ar(t) + (x_0 - Ar_0)e^{-Bt} \tag{6}$$

$$n(t) = \left(\kappa + \tfrac{1}{\alpha}\right)\tau Ar(t) - \left(B + \tfrac{1}{\alpha}\right)\tau (x_0 - Ar_0)e^{-Bt} \tag{7}$$

where x_0 and r_0 are respectively the initial values of experienced employees and capacity requirements, and where:

$$A = \frac{1 + \tau\delta(\rho - \kappa)}{\tau^2\kappa(\rho - \kappa)\left[\delta^2 + \beta(1-\delta)^2\right] + \tau\delta\rho + 1} \tag{8}$$

and

$$B = -\frac{\rho}{2} + \frac{1}{2\tau}\left[\rho^2\tau^2 + 4\frac{\tau\rho\delta + 1}{\delta^2 + \beta(1-\delta)^2}\right]^{\frac{1}{2}} \tag{9}$$

Furthermore, B is guaranteed to be positive.

The proof is in the Appendix. ■

The solutions for *x(t)* and *n(t)* are both composed of a first, steady-state term that increases with time and a second, transient term that decreases with time. Additionally, *n(t)* contains a term replacing experienced worker attrition. As required, the coefficient of the steady state term A reflects both the capacity and the salary penalties significantly over the entire time interval. Furthermore the coefficients of the steady-state solutions are stationary. This leads to the following lemma:

Lemma 1

For any control policy of the form $x(t) = Ae^{\kappa t}$ the ratio of the numbers of new to experienced employees is stationary with respect to time and is solely a function of training time τ, the attrition rate the requirement growth rate. That is

$$\frac{n_{ss}(t)}{x_{ss}(t)} = \left(\kappa + \alpha^{-1}\right)\tau \tag{10}$$

Note that the attrition rate $1/\alpha$ is often negligible in comparison with the growth rate κ. The proof is in the Appendix. ■

To increase our intuition as to the behavior of this policy under typical real-world conditions, we develop a numerical example using the following stylized values for the parameters. These values are taken from interview data for a typical ERP installation company.

Table 1: Example Parameters and Initial Conditions

Requirement Growth Rate	κ	20%	per year
Attrition rate	α	0%	per year
Discount Rate	ρ	50%	per year

Training time	τ	1.5	years
Relative productivity of new employees	δ	-25%	exp. empl. equivalents
Relative cost of Salary to Capacity penalties	β	50%	
Initial Requirements	r_0	100	exp. empl. equivalents
Initial Experienced Employees	x_0	90	experienced employees

Figure 1: Optimal Capacity vs. Time

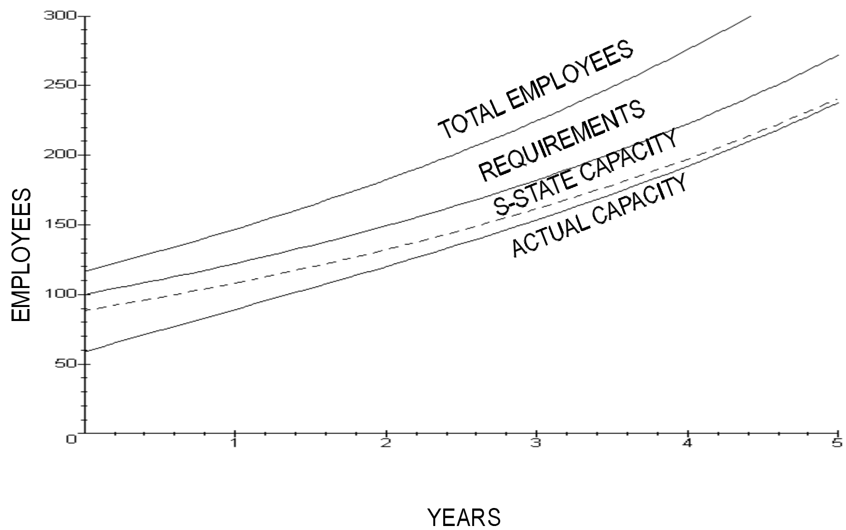

Figure 1 shows the time-series for the total employees, the capacity requirements, and the actual and steady-state capacity from years 0 to 5. (As a reminder, the capacity is defined as the number of experienced employees, plus the number of new employees *discounted* by the relative productivity of new employees *(δ)* (Equation 1). The result will be measured in equivalent experienced employees and will always be less than the *total* number of employees.) The time-series reveal several interesting points typical of the real-world situations. Firstly, the effect of initial conditions upon the track of the actual capacity vanishes quite rapidly. Thus, the steady-state solution dominates the transient term after Year 2. Secondly, the steady-state capacity is less than the requirements. In this particular case, using the stylized parameters from Table 1, only 89% of the requirements are met. This deficit comes about because the capacity requirement penalty alone would cause the capacity to exactly meet requirements. If there is, in addition, a positive salary penalty cost then the marginal benefit of each new employee will be less than in the no-salary-penalty case. Hence, there will be a smaller capacity as the salary penalty coefficient β increases. Finally, the steady-state total of new and experienced employees exceed the capacity requirements by 25%. While such a surplus is not, as we will show later, a mathematical necessity, this behavior is fairly typical of real-world parameters.

In comparison, hiring only enough warm bodies irrespective of experience level to meet the capacity requirements would be quite problematic. Recall that the result from Lemma 1 is necessary for any hiring policy which maintains a constant ratio of employees to requirements. Hence, using the "warm-body" policy would ensure that only $1/(1+\kappa\tau)$ = 77% of the capacity requirements would be met. This is 12% less than the optimal policy. Thus, in many situations such a "warm-body" strategy may risk increasing the installation time or decreasing its quality. Pursuing this over the long term could lead to a loss in market share compared with the optimal strategy.

Figure 2: Optimal Employee Levels vs. Time in Years

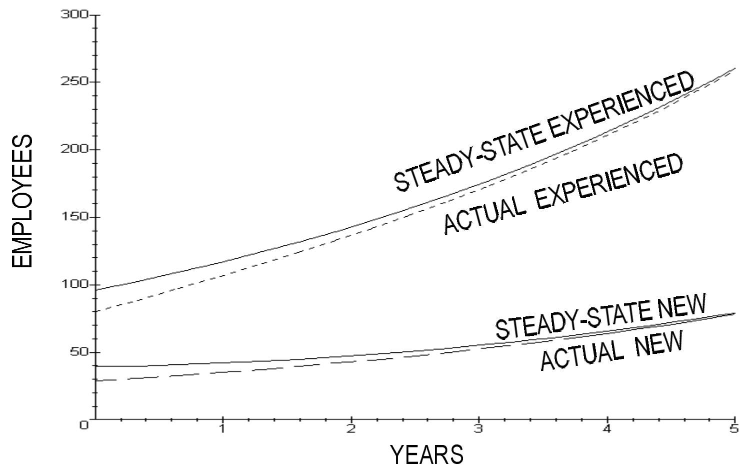

Figure 2 shows the optimal employee levels over time. Note that the actual new and experienced employees converge quickly upon their respective steady-state behaviors (which are 96 and 29% of the capacity requirements). However, note also the excess number of new employees at Year 0 required to overcome the deficit of experienced employees. This shows that catching up once behind the optimal curve can be expensive in terms of salary costs. But even once this transient behavior dies out, there still remains a large steady-state fraction of new employees. In fact, from Lemma 1 the new employee fraction will be 23%. Thus, the need to train employees and the large growth rate in requirements combine to create quite a substantial new employee fraction. If implementation quality is not only a function of the percentage of requirements met, but also the of the average employee experience level, then training these employees may create quite a significant quality risk.

4 SYSTEM BEHAVIOR

To sharpen our intuition further with respect to managing knowledge workers under demand growth, we will attempt to more completely characterize the behavior of the steady-state solutions. To this end we will develop several propositions concerning the comparative statics for measures of interest and again provide a numerical example of a typical firm. For the following section, we will assume for simplicity that the attritition rate is negligible in comparison with the growth rate.

4.1 Service Level

First we will define the implementation service level S to be the ratio of the actual capacity to requirements in the steady state. That is:

$$S = \frac{c_{ss}(t)}{r(t)} \tag{11}$$

Proposition 1

If the training time (τ) is less than some critical value $\tau_U(\delta,\rho,\kappa)>0$, then the steady-state service level S will decrease with increasing training time τ.

Additionally, if the relative productivity of new employees (δ) is greater than some critical value $\delta_L(\beta,\tau)<0$, then the service level S will decrease with increasing capacity requirement growth (κ).

Furthermore,

$$\tau_U \geq \frac{1}{2\delta^2\kappa(\rho - 2\kappa)} \text{ if } \delta \geq 0$$
$$\tau_U \geq \frac{2}{3\kappa} \text{ if } \delta < 0 \tag{12}$$

and

$$\delta_L \leq \frac{2\beta}{1-\beta} \leq -\frac{1}{2} \tag{13}$$

Note that the third term in Equation 13 is only true if $\beta<1$. The proof is in the Appendix. ■

Note that the constraints put on the critical values above will include the majority of real-world cases with which the author has come into contact. Thus, we can state that in many situations the implementation service level will decline with increasing training times and growth rates. Intuitively the reason behind this deterioration is that there are simply not enough productive workers in order to do all the implementation tasks required in each period. Some tasks in at least some implementations will either need to be postponed or left undone. This will result in a degradation of service to the customer in the form of increased project completion time or lesser implementation quality.

4.2 Employee Experience

Another possible contributor to service quality is the average experience of the employees implementing a software package. Similar observations have been made in other service industries (Oliva 1996). If a task is done by a new employee, as is typical of on-the-job training, it is less likely to be done correctly. Some of these mistakes will be found, others likely will not (Ford 1995). Hence, the greater the fraction of new employees in a firm, the lesser the quality of implementation. In order to measure the experience level, we define the experienced employee ratio X to be the steady-state ratio of new employees to total employees. That is:

$$X = \frac{x_{ss}(t)}{x_{ss}(t) + n_{ss}(t)} \tag{14}$$

Proposition 2

The experienced employee ratio will decrease with increasing employee training times.

The proof follows immediately from Proposition 1 and Lemma 1. ■

In essence, the results of Proposition 2 lead to the conclusion that the quality of individual tasks and hence the entire implementation project will decline with increasing training time.

4.3 Productivity and Cost

Finally, it would be instructive to learn how operating costs are influenced by training times. We will define the productivity P to be the ratio of the effective capacity measured in experienced employee equivalents to the total number of employees. Hence:

$$P = \frac{c_{ss}(t)}{x_{ss}(t) + n_{ss}(t)} \tag{15}$$

Actually, this is a somewhat conservative measure because it does not account for the increases in administrative overhead necessary to maintain a larger firm. However, it is sufficient for to prove the following proposition.

Proposition 3

Steady-state productivity will decrease with increasing employee training times and demand growth rates.

The proof is in the Appendix. ■

Finally, it would be interesting to see the effect of training time and growth on operating costs. To this end we define the salary cost *d(t)* to be the product of the total number of the employees and the annual salary per employee γ:

$$d(t) = \gamma\left[x_{ss}(t) + n_{ss}(t)\right] \tag{16}$$

This definition leads to the following proposition.

Proposition 4

If the conditions describe in Proposition 1 hold, then salary costs will increase with employee training time and demand growth rates.

The proof follows immediately from the definition of productivity and Propositions 1 and 3. ■

Again, this is a conservative estimate for operating costs because it does not include the cost of quality. Proposition 4 tells us that the cost per task actually performed will increase with training time or demand growth *under all conditions*. Hence, many tasks will remain either undone or botched. The only reason that the salary cost drops at all in Proposition 5 is that the under certain extreme conditions, the firm may decide to increase its service level.

4.4 Remarks

The net effect of the above propositions is that in many if not most real-world situations, both implementation quality and productivity will decrease in a market dominated by increasing demand of knowledge workers. But how serious will the impact be?

To give a numerical flavor for the impact of training time on these operational measures, examine Figure 3, which uses the stylized parameters presented earlier in Table 1. Training times range up to five years as periods in excess of this seem quite rare in most knowledge worker staffing problems (with the exception, of course, of academic researchers!).

.

Figure 3: Operational Measures vs. Training Time

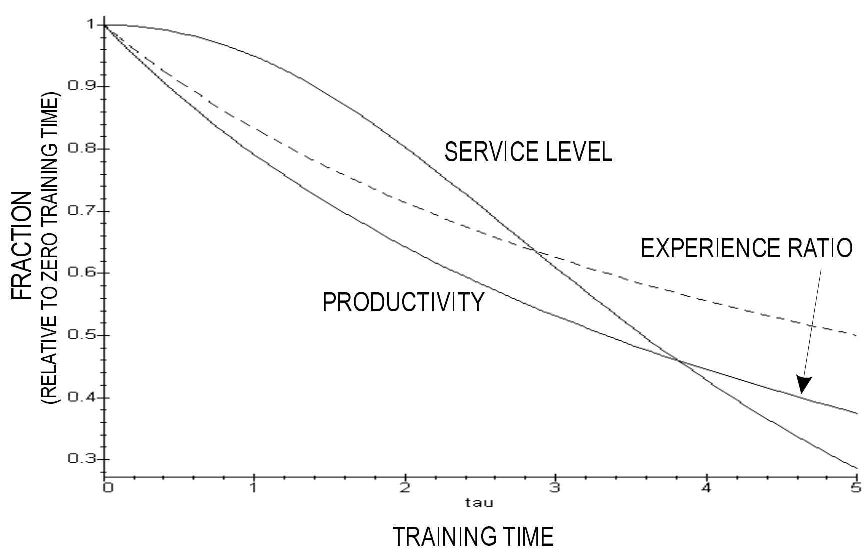

Figure 3 shows how increasing training times degrade operating measures even under a staffing policy optimal with respect to the objective function stated earlier. (It should be mentioned that the basic trends of Figure 3 also seem to be true for other sorts of similar objective functions such as linear salary penalties which were discussed earlier.) As far as the actual measures are concerned, the service level S decreases significantly only for training times over 1 year. The experience ratio X, being inversely proportional to training time, falls along the expected hyperbolic path. Productivity also declines in a convex manner. Assuming the earlier stylized training time of 1.5 years, service, experience, and productivity levels decline respectively to only 89, 77, and 71% of the no-training-time case. Given the relatively sedate rate of market growth (20% per annum), these are fairly serious degradations in the operating characteristics, and must be accounted for when grappling with these sorts of situations. Especially, as in the real world, unlike the current model, such degradation will likely lead to a decline in market growth.

5 DISCUSSION

To summarize, this paper has developed a dynamic programming formulation for the problem of staffing knowledge workers under demand growth. The novelty in this formulation is its insensitivity to the typical problems associated with staffing policies formulated using an HMMS-style objective. The resulting formulation was solved for an optimal policy characterized by both a steady-state and a transient component. An illustration of the policy in the setting of software implementers of enterprise requirement planning systems was developed. Additionally, the behavior of the system to changes in training time and market growth rates was examined both analytically and numerically.

There are a number of important results for this paper, including the actual policy developed. However, the most important result from a theoretical standpoint is the degradation in operating measures in many circumstances induced by the training of new employees under conditions of market growth. For example, this paper has shown that average knowledge-worker productivity will decline with growth. This result, which represents the effect of individual learning on planning, contrasts nicely with those involving learning that is embodied in the aggregate productive technology or organizational procedures of the firm (Gerchak, Parlar, and Sengupta 1990). This path is particularly interesting as some of the larger implementation firms such as CSC are beginning to develop databases of common implementation problems for their employees to use for troubleshooting. In this way, implementation firms may be able to both increase the productivity of their employees directly and reduce the required employee training period. One useful line of research would be to develop a staffing model including both types of learning.

Another important point of the production planning literature for learning curves is that of market interaction. In Spence's (1983) classic paper, he showed that a firm will tend to produce more products under an aggregate learning curve when the market size is dependent on the product cost. While not directly comparable, this paper suggests that, as the knowledge required for an industry increases, so too will the costs of growth increase. In particular, the increasing cost and decreasing service level resulting from training employees could result in a loss of customer utility. Hence, all other things being equal, a market characterized by knowledge workers may develop more slowly than one that is not. Thus, the presence of an experience curve embodied at the individual level may possibly have consequences opposite from an experience curve embodied at the firm level. An integrated market-production model of knowledge workers under market growth with utility-driven market growth to investigate this proposition would accordingly also be of interest.

Finally, there has been recently a great deal of interest in viewing supply chains as networks of interlinked core competencies or knowledge bases. This paper, being essentially an operational model of competency management, may also have an implication for these "technology supply chains" (Fine 1996). As has been shown, managing knowledge workers under market growth is an expensive business and requires deep pockets, which are beyond the means of many third-party implementation firms. On the other hand, the pockets of the software firms, who have the most to gain from rapid market growth, should be somewhat deeper. Thus, some situations might warrant direct involvement of software firms in the implementation business. One example of this has been the recently announced alliance between SAP

AG and CSC in order to provide a reliable source of implementers for SAP installation projects (SAP 1996). This suggests that in certain situations of market growth, it may make sense for a supplier of product to support a supplier of knowledge workers. Further research in such questions would be quite valuable as we shift ever more to a knowledge economy.

6 APPENDIX

Proofs can be found in a fuller technical version of this paper available from the author.

7 REFERENCES

Bartholemew, D. J. (1973). *Stochastic Models for Social Processes*. 2nd ed. London: John Wiley & Sons.

Bertsekas, Dimitri P. (1987). *Dynamic Programming: Deterministic and Stochastic Models*. Englewood Cliffs, NJ: Prentice-Hall.

Ebert, Ronald J. (1976). "Aggregate Planning with Learning Curve Productivity." *Management Science* 23 (2):171-182.

Fine, Charles H. (1996). "Industry Clockspeed and Competency Chain Design: An Introductory Essay." Paper read at Manufacturing and Service Operations Management Conference, June 24-25, 1996, at Dartmouth College, Hanover, New Hampshire.

Ford, David N. (1995). *The Dynamics of Project Management: An Investigation of the Impacts of Project Process and Coordination on Performanace*. Doctoral, Dyanamic Engineering Systems, Massachusetts Institute of Technology, Cambridge, MA.

Gaimon, Cheryl, and Gerald L. Thompson (1984). "A Distributed Parameter Cohort Personnel Planning Model that Uses Cross-Sectional Data." *Management Science* 30 (6):750-764.

Gerchak, Yigal, Mahmut Parlar, and S. Sankar Sengupta (1990). "On Manpower Planning in the Presence of Learning." *Engineering Costs and Production Economics* 20:295-303.

Grinold, Richard C. (1977). *Manpower Planning Models*. New York: North-Holland.

Gross, Neil (1997). "Year 2000: The Meter's Running." *Businessweek*, December 29, 1997.

Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon (1960). *Planning Production, Inventories, and Work Force*. Englewood Cliffs, NJ: Prentice-Hall.

Oliva, Rogelio (1996). *A Dynamic Theory of Service Delivery: Implications for Managing Service Quality*. Doctoral, Management Science, Massachusetts Institute of Technology, Cambridge, MA.

Parker, Geoffrey Goodrich (1997). Building Supplier Capability through Project Choice. Cambridge, MA: Massachusetts Institute of Technology.

SAP (1996). CSC Announces Global Alliance with SAP.

Spence, A. Michael (1983). "The Learning Curve and Competition." *The Bell Journal of Economics* 12 (1):49-70.

Sterman, John D. (1988). People Express Management Flight Simulator: Software and Briefing Book. Cambridge MA: Massachusetts Institute of Technology.