Modeling the Unanticipated Side Effects of Successful Quality and Productivity Improvement Programs

Nelson P Repenning
Department of Operations Management and System Dynamics
Sloan School of Management, E53-336
Massachusetts Institute of Technology
Cambridge, MA, USA 02139
Tel: 617 253 0250 Fax: 617 253 6466
E-mail: nelsonr@mit.edu

Abstract

This paper describes a simple model of a manufacturing firm in which a successful productivity improvement program is implemented. This model is an attempt to generalize an earlier theory developed to explain one company's paradoxical experience with Total Quality Management (Kofman, Repenning and Sterman, 1994). The model describes a dynamic hypothesis concerning the firm's financial performance. In this model the Half-life Equation suggested by Schneiderman (Schneiderman 1988) is used to determine the maximum rate of improvement. The spread of skills and commitment is modelled as a diffusion process, and the allocation of resources to support that commitment is represented as a dynamic adjustment process with a multi-dimensional utility function and fixed resource constraint. This formulation, with the assumption of locally rational decision rules, results in differential rates of improvement in the capacity and demand generating areas of the firm. This differential, when coupled with traditional accounting, pricing, and human resource policies, can create unanticipated side effects that result in sub-standard performance or failure of the program.

Modeling the Unanticipated Side Effects of Successful Quality and Productivity Improvement Programs

I. Introduction[1]

The advent of such techniques as Total Quality Management (TQM) and Business Process Reengineering has increased the dynamic complexity of the manager's task. If the advocates of these and other programs are to be believed, managers can no longer be content with business as usual; rather, the successful firm of the future will be one that has developed the ability to continually improve the efficiency and effectiveness of its operations at a rate faster than that of its competitors [Demming 1986, Stata 1989, Hammer 1993]. Each of the aforementioned "improvement" programs has generated a number of high profile successes. However, as these improvement techniques have been disseminated to, and used by, an ever widening audience, it has become clear that the application of these techniques has not been universally successful. In particular, TQM, which is perhaps the most accepted "improvement movement", has been the subject of recent criticism in the popular business press [Newman 1991, The Economist 1992, Taylor 1992]. A recent study by Ernst and Young concluded few companies that use TQM experience significant change in profitability [Ernst and Young 1991]. A study of Baldridge award finalists by the US General Accounting Office concluded, that while it did significantly improve quality and productivity, TQM did little to improve the returns on sales or assets [GAO 1991].

Two competing theories have emerged to explain why these programs can be so successful in some organizations but not in others. The first focuses on failures of implementation [Kaufman 1992], while the second suggests that TQM is little more than another management fad [Harte 1992]. Unfortunately, however, to date these issues have received little formal analysis [Kim and Burchill 1992 is a notable exception]. As a result, the current debate on the usefulness of these programs has provided little guidance to the practicing manager who wishes to implement an improvement program, or to the academic who wishes to better understand the determinants of a program's success or failure.

Building on a previous paper which analyzed, in depth, one company's experience with TQM [Kofman, Repenning, and Sterman 1994], this paper develops a simple model to analyze the firm level effects of improvement, and the conditions under which this program is likely to be a success or failure. Section II presents a simple representation of the core improvement process. Section III augments this model by explicitly representing the diffusion of skills and the allocation of improvement resources. In the spirit of Morecroft [Morecroft 1985], these decision rules are formulated based upon boundedly rational policies. Section IV develops a model of a monopolistic firm which experiences exogenous improvement in its productive capabilities. Section V combines the models presented in sections III and IV to form a fully endogenous model of firm level improvement. Section VI presents final conclusions and thoughts on possible directions for future research.

II. Modeling the Improvement Process: The Half-Life Model

The actual process of improving any business operation is likely to be at least as complex as the operation itself. Since the present focus is on the consequences of improvement, rather than the improvement process itself, the model exploits an empirical regularity, first noticed by Schneiderman [Schneiderman 1988]:

> ...Any defect level, subjected to legitimate QIP{Quality Improvement Process} decreases at a constant rate so that when plotted on semi-log paper, it falls on a straight line.

Here a defect is broadly defined as any measurable, undesirable, component in the process of bringing a product to market. Under this definition defective products, late deliveries, and long

---

[1]. Many of the ideas and formulations presented here were developed jointly with John Sterman. Both the content and presentation of this paper have benefited from comments and suggestions by Edward Anderson, Rogelio Oliva, Elizabeth Saltonstall, and M. Anjali Sastry. All contributions are gratefully acknowledged.

product development times can all be considered defects. This relationship translates to a simple, first order, differential equation describing the behavior of defect measure i.

$$\frac{\partial D_i}{\partial t} = -\phi_i D_i, \quad 0 < \phi_i < 1 \tag{1}$$

Schneiderman labeled this relationship the Half-Life Model since the time required for any defect measure to fall by 50% is constant.

The parameter ϕ_i is specific to the process from which the defect measure is generated. Processes that are technically simple and require the input of only a few people will improve rapidly, implying a large ϕ_i, while processes that are technically complex and span numerous organizational boundaries will improve relatively more slowly. Schneiderman's results indicate that processes with short half-lives are likely to be associated with direct manufacturing, while areas such as product development and administration are likely to improve much more slowly [Schniederman 1987, Kaplan 1990].

Assuming the firm has **n** distinct areas, each containing **i** defect generating processes, two modifications are made to the "Half-Life" equation for the purpose of this analysis.

$$\frac{dD_i}{dt} = -\phi_i (D_i - D_{iMin}) \cdot C_n, \quad 0 < C_n < 1 \tag{2}$$

First, the rate of defect reduction in process i is proportional to the current level minus a minimum that is explicitly defined for each process. This is an important element for a formulation that must represent improvement efforts, such as reducing product development time, for which the theoretical minimum level is not zero. Second, the improvement rate is regulated by an index measure, restricted to the 0-1 interval, of the current commitment to, and use of, the given improvement methods in area n. The construct commitment, denoted C_n is an explicitly defined level which measures the percentage of the full time equivalent workforce in area n that has acquired, and is using, the skills and techniques dictated by the given improvement program.

III. Diffusion of Commitment and Allocation of Resources

Improvement does not begin immediately with the implementation of a given program. Rather, it requires the dissemination of the appropriate skills, which takes time, and the continued belief, on the part of both management and the workforce, that the methods in question actually work. As a result, the change in the commitment to a program can be decomposed into a "push" from management and a "pull" from results [Shiba, Walden, and Graham 1993].

$$\frac{dC_n}{dt} = \theta(C^* - C_n) + w_n C_n (1 - C_n), \quad 0 < \theta < 1, \ 0 < \mu < 1 \tag{3}$$

The first term on the right hand side of (3) represents management's push. Management sets a target, C*, for firm-wide commitment to the improvement program, assumed always to be 100% in this model. Absent pull effects, the actual commitment level approaches management's target via a first order information delay. The delay time, $1/\theta$, is the average time required for management to teach the workforce the new improvement tools and to enlist their participation in the program.

A diffusion process is used here, in the second term of (3), to determine the strength and sign of the "pull" effects [Bass 1969, Homer 1987, Paich and Sterman 1993]. Early in the program's life the accumulated intra-firm experience with the program will be low, consequently management's push will dominate. Later as experience accumulates, participants can evaluate the effectiveness of the program based on personal experience. The sign and strength of the pull effect is assumed to be determined by three factors: success of the program, p_n, adequacy of support resources, a_n, and perceived job security, s.

$$w_n = \omega_n [f_r\{p_n\} + f_a\{a_n\} + f_s\{s\}] \tag{4}$$

The parameter w_n represents the intensity of communication in the particular area.

The success of the program in area n, p_n, is defined as the ratio of the percentage change in productivity, P_n, to the predicted percentage, P_n*. The predicted change is determined by the half-lives associated with the processes in the given area.

$$p_n = \frac{P_n}{P_n^*} \tag{5}$$

The effect of results is weighted by the non-linear function, $f_r\{p_n\}$, shown in Figure 3.1. This function is negative when improvement is well below the prediction, representing the workforce's skepticism caused by poor results. As the improvement rate approaches the prediction, the effect becomes positive, asymptotically approaching one.

The addition of (5) creates the first important feedback loop in this system. By initializing commitment to zero, the defect level to 100, and assuming an improvement half-life of nine months, the introduction of an improvement program can be simulated by introducing, in month twelve, a unit step input to parameter C*, management's target commitment level (see Table 3.1 for additional parameter assumptions).

Figure 3.1

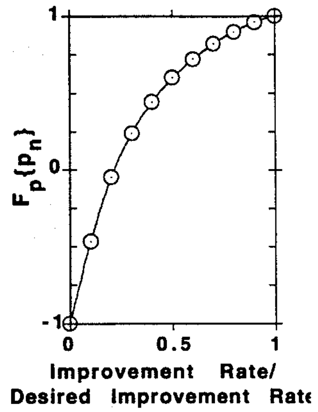

Figure 3.2

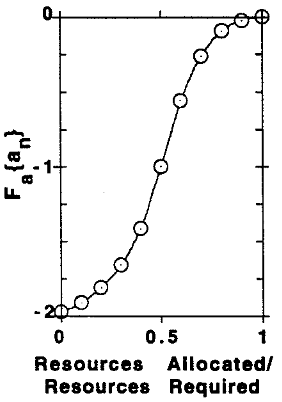

Figure 3.3

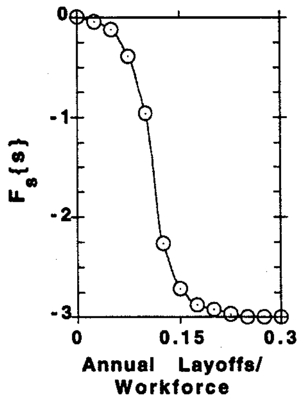

Table 3.1

Parameter	Value
f	.077 (1/months)
D Initial	100 (defects)
D Minimum	10 (defects)
q	.084 (1/months)
w	.5
C Initial	0
C*	0 until time=12, then 1
$F_a\{\cdot\}$	0
$F_s\{\cdot\}$	0

The initial result is a rising commitment to improvement and a reduction in defects. Figure 3.4 shows that with this choice of parameters, management's initial push results in a small increase in commitment, as initial worker skepticism retards diffusion.

Figure 3.4

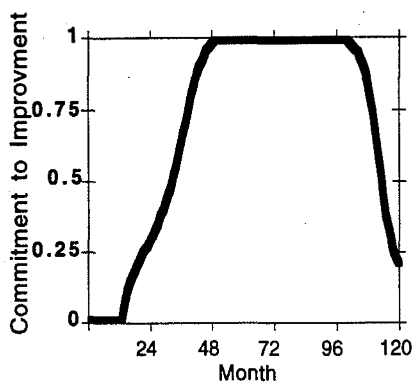

Figure 3.5

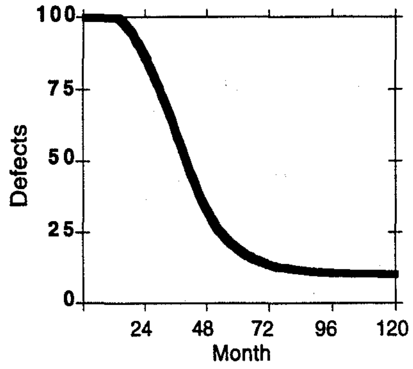

However, as the initial push begins to generate noticeable improvement, commitment increases. The reduction in defects further increases commitment creating positive feedback. As the defect level approaches the minimum level, improvement eventually must fall to zero.

The second term inside the brackets on the right-hand side of (4), $f_a\{a_n\}$, represents the effect of resource adequacy on worker commitment. The model assumes that workers, in order to participate effectively in an improvement program, require resources in the form of management's attention and a reduction in their current responsibilities. The total resource requirement in area n, r_n, is the product of the number of people in the area, L_n, the improvement resource requirement per person in area n assuming full participation, r_n, and the current commitment level in that area, C_n.

$$r_n^* = L_n \rho_n C_n \tag{7}$$

The adequacy of resources, a_n, is the level of resources currently allocated to the area, r_n, divided by the resource requirement.

$$a_n = \frac{r_n}{r_n^*} \tag{8}$$

The non-linear weighting function, $f_a\{\}$, is shown in Figure 3.2. If the adequacy of support resources falls significantly below the required level, the resulting negative effect on word of mouth will dominate any positive effects resulting from the improved productivity.

The final factor affecting the change in workforce commitment to improvement is job security. As will be shown shortly, under certain circumstances, a successful improvement program may result in excess manufacturing capacity. If this is the case, management may be tempted to cut costs by firing some portion of the workforce. Based on the assumption that the workforce can correctly connect the productivity improvements it makes with the resulting excess capacity, the model assumes that commitment to improvement will be reduced if workers believe that further improvement will increase the probability of lay-offs or firings. As a result, the change in commitment can only be positive if the workforce believes management's commitment to no lay-offs. The workforce's belief in this commitment is assumed to be solely a function of management's past actions. The change in commitment is modeled as a decreasing function of the workforce's "memory" of past lay-offs. A non-linear memory structure is used to represent this process. When the annual percent lay-off rate, S, is greater than the workforce's current memory, s, the memory is updated very quickly, while when the converse is true, the memory is updated very slowly .

$$\frac{ds}{dt} = n(S - s);\quad n = \eta \text{ for s > S, } n = \nu \text{ for S > s, } \nu >> \eta \tag{9}$$

The "memory" of lay-offs is weighted by the non-linear function $F_S\{.\}$ shown in Figure 3.3. The function decreases quickly as the memory of lay-offs increases. The decline is of sufficient magnitude to dominate the other two effects, $F_p\{.\}$ and $F_a\{.\}$, if perceived job security is low.

The allocation of resources to different areas in the firm plays an important role in determining the final outcome of the improvement program. The model assumes a fixed resource constraint R that is not sufficient to support each improvement effort in each area at 100% commitment. The issues surrounding changing the total amount of resources are not considered. The amount of resources allocated to area n, r_n, is determined by the percent allocation parameter x_n.

$$r_n = x_n R \qquad (10)$$

This coefficient is determined by a two-dimensional attractiveness function that exponentially weights each area's current resource requirement, measured as a fraction of firm's total resource requirement, r*, and that area's productivity improvement rate, P_n .

$$x_n = \frac{P_n{}^{\alpha} q_n^{\beta}}{\sum_N (P_n{}^{\alpha} q_n{}^{\beta})}, \quad \alpha,\beta > 0 \qquad (11)$$

$$q_n = \frac{r_n^*}{\sum_n r_n^*} \qquad (12)$$

(11) and (12) add additional feedback loops which significantly affect the dynamics of diffusion and performance in a firm with multiple areas. For the results presented below, there are assumed to be n=2 areas involved in the improvement effort. All parameter assumptions are identical to those made in the previous section, except that Area Two is assumed to have an improvement half-life of thirty-six months, while the half-life for improvement in Area One remains nine months (additional parameter assumption are listed in Table 3.2). As in the previous section, the introduction of an improvement program is simulated via the introduction of a unit step in the parameter C* at month twelve.

The initial results are similar to the single area, resource-unconstrained case: Commitment to the improvement effort jumps initially and then begins to grow exponentially as the positive feedback of the diffusion process begins to dominate. Due to its faster improvement rate, Area One begins to receive a greater share of the available improvement resources.

Table 3.2

Parameter	Value
f_1	.077 (1/months, 9 month Half-Life)
f_2	.02 (1/months, 24 month Half)
D_1,D_2 Initial	100 (defects)
D_1,D_2 Minimum	10 (defects)
q_1,q_2	.084 (1/months)
w_1,w_2	.5
C Initial	0
C*	0 until time=12, then 1
L_1,L_2	100 (people)
r_1,r_2	1 (resources/person/month)
a	25
b	1
Fs{.}	0

Figure 3.6

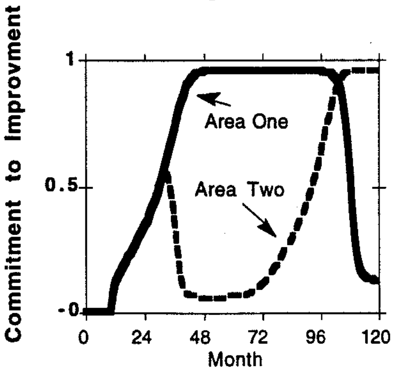

Figure 3.7

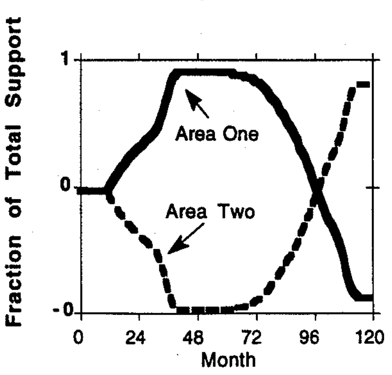

Figure 3.8

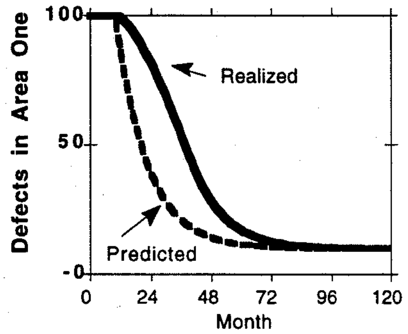

Figure 3.9

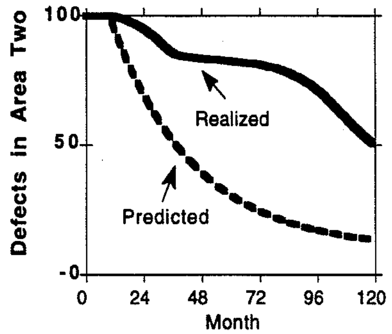

Subsequently, commitment in Area Two diminishes quickly due to lack of support, further strengthening management's commitment to supporting Area One. Eventually, as measurable improvements become more difficult to make, commitment declines in Area One. Management support is then re-focused to Area Two, which experiences a subsequent recovery in commitment. Area One significantly outperforms Area Two in terms of defect reduction. It is important to note that this difference is much larger than that predicted by the simple half-life model

IV. Firm Level Effects of Improvement

This section develops a simple model of the firm to help evaluate the consequences of the unbalanced improvement rates identified above. The firm is assumed to be a monopolist facing a constant price elasticity demand curve, (13), with elasticity e and scale parameter A.

$$Q_D = AP^{-\varepsilon} \tag{13}$$

The firm is also assumed to require two factor inputs, direct labor, L_1, and indirect labor, L_2, and to be endowed with productive technology represented by the Leonteif production function, (14).

$$Q_s = Min[\alpha_1 L_1, \alpha_2 L_2] \tag{14}$$

The first type of labor, L_1, is assumed to be associated with direct manufacturing. The current stock of this type of labor is equal to the time integral of hiring, L_h, attrition, L_a, and lay-offs L_d.

$$L_t = \int_t (L_h - L_a - L_d)dt \tag{15}$$

Hiring, which is constrained to be positive, is the attrition rate plus a fractional correction for the difference between the current labor stock and the desired level, L^*.

$$L_h = Max\left[\zeta_h\left(L^* - L\right) + L_a, 0\right] \tag{16}$$

The desired stock of labor is set to the profit maximizing value, the long-run optimal quantity demanded, Q^*_{LR}, divided by the productivity parameter a_1.

$$L_1{}^* = \frac{Q^*_{LR}}{\alpha_1} \tag{17}$$

The attrition rate is equal the current labor stock divided by the average career length, T_l.

$$L_a = \frac{L}{\tau_L} \tag{18}$$

The rate of autonomous work force reduction, or lay-offs, is formulated similarly to hiring with the addition of a multiplicative constant l, which measures management's willingness to fire or lay-off excess labor.

$$L_d = \lambda \bullet Max\left[-\zeta_d\left(L^* - L\right), 0\right] \tag{19}$$

The second type of labor, L_2, is assumed to perform tasks that are not directly associated with manufacturing. It is further assumed that this type of labor can be immediately hired or fired with no additional consequences to the firm. As a result, L_2 is always equal to the short run optimal level, the profit maximizing short-run demand, Q^*_{SR}, divided by the productivity parameter a_2.

$$L_2 = \frac{Q^*_{SR}}{\alpha_2} \tag{20}$$

The final element required to completely specify the model is the mechanism required to determine pricing. Since the chosen technology exhibits constant marginal costs, the profit maximizing price can be calculated via a simple rule which takes the convenient form of marginal cost multiplied by a fixed mark-up [Varian 1992].

$$P^* = \frac{mc}{1 - \frac{1}{\varepsilon}} \tag{21}$$

With completely flexible factor acquisition, the marginal cost of producing one unit is simply the sum of each unit factor cost divided by its respective productivity parameter. However, due to the delays in the acquisition and reduction of direct labor, and in an effort to more realistically represent the price setting operation, a more complex pricing structure is used. First, it is assumed the equilibrium contribution to marginal cost of the direct labor, $\frac{p_{L1}}{\alpha_1}$, is known with certainty. Second, when the firm's stock of direct labor is above the optimum level, $L_1{}^*$, the marginal contribution is scaled downward by a non-linear, decreasing, function of the current level of direct labor utilization[2], f{u}.

$$u = \mathrm{Min}\left(\frac{\alpha_1 L_1}{Q_D}, 1\right) \tag{22}$$

Third, it is assumed that the productivity of indirect labor is difficult to observe, and as a result, difficult to calculate. The standard management accounting solution to this problem is to estimate a product's total cost based upon its direct labor content. The structure used here approximates this practice by calculating expenditure on direct labor as a fraction of the total expenditure, and then using this quantity to scale the direct marginal cost accordingly. Finally, the actual cost adjustment, M, is assumed to be an exponentially weighted average of the adjustment indicated by the mark-up rule.

$$\frac{dM}{dt} = \xi\left(\frac{f\{u\} L_1 p_{L_1}}{L_2 p_{L_2} + f\{u\} L_1 p_{L_1}} - M\right), \quad 0 < \xi < 1 \tag{23}$$

The result of these assumptions is a rule for calculating marginal cost that takes the following form;

$$mc = \left(\frac{p_{L1}}{\alpha_1} \cdot f\{u\}\right) \cdot \frac{1}{M} \tag{24}$$

---

[2] The actual optimal policy in this case is to consider direct labor a fixed factor, hence the marginal cost is simply the contribution of the indirect labor. This policy can be closely approximated by making f{u} decline very quickly for value of u<1.

Marginal cost is equal to the equilibrium direct marginal cost adjusted for utilization and multiplied by a scaling factor, 1/M, to account for indirect costs. It should be noted this structure results in the correct calculation for marginal cost, and, as a result the profit maximizing price, when the model is in equilibrium.

It is now possible to simulate this portion of the model in isolation. The parameter assumptions are given in Table 4.1. For the purpose of this simulation, the productivity of direct labor, measured by a_1, is assumed to make an approximately four-fold improvement in productivity, while the productivity of indirect labor, a_2, remains constant. The firm is also assumed to have a policy of no lay-offs.

Table 4.1

Parameter	Value
A	25,600
e	2
a_1,a_2 Initial	.25 (1/units produced)
z_h	.083 (1/months)
z_d	.5 (1/months)
T_L	240 (months)
x	.083 (1/months)
p_{L1}	.5 (dollars/person/month)
p_{L2}	1.5 (dollars/person/month)

Figure 4.1

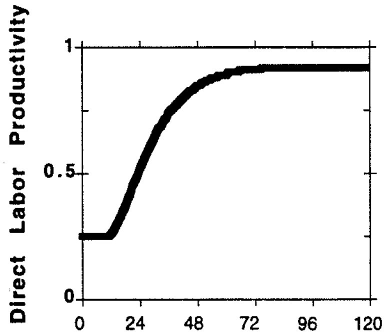

Figure 4.2

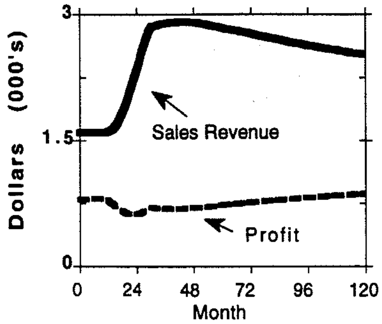

The result of this improvement is a substantial increase in unit sales and sales revenue. However, profit increases much more slowly and profit margin actually decreases. The explanation for this somewhat surprising behavior is as follows. First, in this case, although the effective cost of direct labor is falling due to the productivity improvement, the demand for direct labor is actually falling. This appears to violate conventional logic concerning the downward slope of demand curves; however, it does not. Demand curves represent possible equilibrium combinations of quantity and price, while in this case the combination of productivity improvement and rigidity in the stock of labor have moved the model into a disequilibrium situation. While the productivity improvement has resulted in an increase in demand due to a lower marginal cost, it has also resulted in an even larger increase in capacity.

Figure 4.3

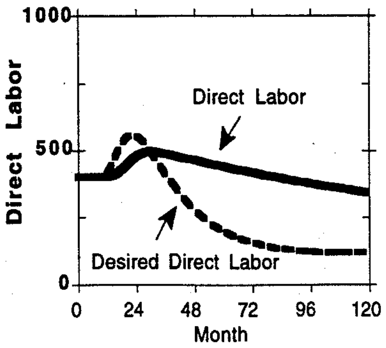

Figure 4.4

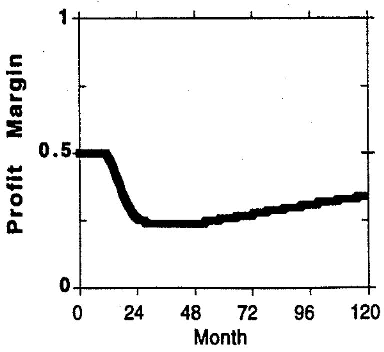

Additional simulations demonstrate that this result depends on two parameters, the price elasticity of demand, and the fraction of total cost that results from direct labor (see Table 4.2). If the elasticity is large and direct labor occupies a large fraction of total cost, the demand for the improving factor will increase. However, if the elasticity is low and the improving factor occupies a small fraction of the total cost, demand will actually decrease[3].

Table 4.2

% Change in Equilibrium Demand for Direct Labor due to Productivity Increase

% Direct Labor of Total Cost	Price Elasticity e=1.5	Price Elasticity e=2	Price Elasticity e=3	Price Elasticity e=4
25%	-63%	-59%	-50%	-39%
50%	-46%	-33%	+6%	+68%
75%	-11%	+33%	+194%	+553%

As a result of this dynamic, the firm finds itself in the position of having excess direct labor. If it is not willing to lay-off this labor then the short-run marginal costs is equal to that contributed by indirect labor The resulting optimal strategy is to price below the long run equilibrium in order to better use the excess capacity. As direct labor is reduced to the desired level through attrition, price can be commensurably increased back to the long run level. The result of this price cutting is that profit grows, but at a rate much slower than that of unit sales or sales revenue.

These effects are further exacerbated by the delays involved in updating the mark-up rule used to account for indirect costs. As the productivity of direct labor increases, the ratio of indirect to total cost increases. The marginal cost, however, is underestimated because the direct labor driven management accounting system only recognizes this change with a delay. This causes price to fall below the short run equilibrium. Additional sensitivity analysis, not presented, indicates that an increase in this adjustment delay further degrades profitability.

V. The Firm with Endogenous Improvement

This section integrates the models discussed in the sections III and IV. The parameter assumptions are identical to those made in the previous sections (see Tables 3.1 and 4.1). Direct labor is assumed to work in Area One, while indirect labor is in Area Two. The respective defect levels are translated to productivity parameters via the following equation, where A_i is the gross (i.e. defective units included) production per unit of labor type i (A_i. is assumed to equal 133 units/month for subsequent simulations).

$$\alpha_i = 1 - \left(\frac{D_i}{A_i} \right) \tag{26}$$

To close the final feedback loop, one additional piece of structure is required. Autonomous reduction in direct labor, L_d, is a function of the difference between the desired and actual labor force, and a parameter l, management's willingness to lay-off workers. While this has been set to

---

[3]. It is also possible to show this result analytically. Contact the author for more details.

zero in previous simulations, in one case presented below, it is assumed that if management perceives a decline in profitability, it will be willing to lay-off some fraction of the excess labor. l is determined by a non-linear function of the ratio of actual to desired profit margin. That is, as the profit margin falls farther below the desired level, l increases from 0 to 1.

The results of two simulations are presented below. The first, as in previous examples, assumes that management maintains its commitment to no lay-offs, and the second assumes l rises quickly if profitably begins to fall. Like in the previous example, the profit margin begins to fall as productivity improves (Figure 5.1). Although profit is increasing, sales revenue is increasing at a faster rate. In the first case, the decline in profit margin is taken to be a sign of poor cost control; management reacts by laying-off the excess labor (Figure 5.3). This makes a small but immediate improvement in profit, and a larger improvement in the profit margin. However, this short term gain comes at the expense of the firm's long run success. By resorting to lay-offs, to improve profitability, management has effectively ended the improvement program. Commitment in both areas falls quickly after the lay-off. Further, due to the nature of the memory process assumed, commitment recovers very slowly, as it takes a long time for management to regain the trust of the workforce.

Figure 5.1

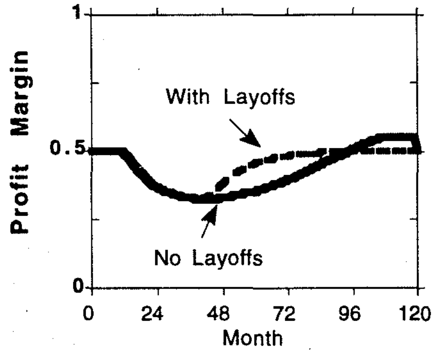

Figure 5.2

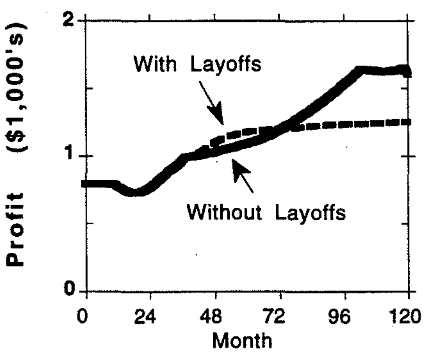

Figure 5.3

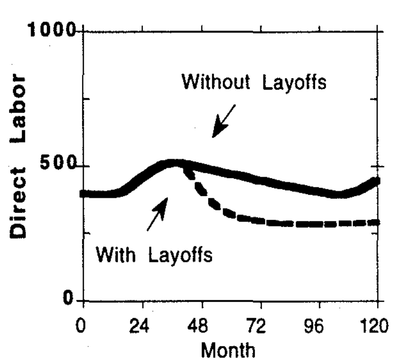

Figure 5.4

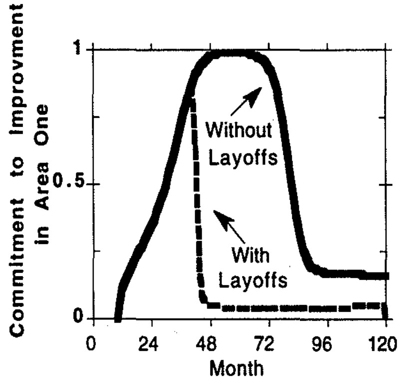

It is obvious that under the specified conditions, the firm does better in the long run by maintaining its commitment to no lay-offs. However, the firm is faced with a classic "worse before better" situation that is frequently found in complex systems [Forrester 1969]. A company embarks upon an improvement program in an effort to improve its competitiveness and profitability. Unless a manager understands the dynamics described above, she is likely to favor the faster improving areas in direct manufacturing, and misinterpret the subsequent decline in profit margin as a sign that the program is not living up to its promise. Further, publicly held firms must face the scrutiny of the financial analysts who are even less likely view a decline in profit margin as a possibly necessary consequence of a successful productivity improvement program. As a result, a firm is probably likely to resort to downsizing in an effort to improve profitability and to demonstrate to the capital markets that they are "serious" about controlling costs. The result of these actions is, of course, the failure of a program the could have provided a significant long term benefit.

VI. Conclusion

In an attempt to explain the failure of improvement programs such as Total Quality Management and Business Process Re-Engineering this paper has developed a simple model of a firm that attempts to implement such a program. As a result of the diffusion process used to model commitment and the policy of starting with early successes, the results presented suggest that improvement is likely to be very unbalanced, with the relatively simple direct manufacturing processes improving very quickly, while the more complex indirect processes, due to a lack of management attention, improve more slowly. If the percent change in productivity is larger than the resulting change in demand, unbalanced improvement will result in excess direct capacity. In the face of excess capacity, the profit maximizing firm will price below the long run equilibrium level, resulting in a substantial increase in unit sales, a small increase in profit, and a decline in profit margin. The decline in profit margin, if it is misinterpreted as poor cost control, may induce management to lay-off excess labor in an effort to improve its bottom line and demonstrate that it is "serious" about cost control. The lay-off effectively ends the improvement program.

Future research on this topic might be profitably focused on two areas. First, the analysis should be extended to models of the firm which contain more complex representations of technologies and market structures. The monopolistic model should be extended to a competitive environment. Second, the model contains a number of hypotheses that could be tested empirically. Foremost among these, the analysis suggests that firms whose production costs are largely a function of direct labor content will benefit more from the current batch of improvement programs, than firms whose costs are derived largely from indirect sources such as product development.

References

—— (1992).The Cracks in Quality. *The Economist.* 322, 67.

Bass, F. (1969). "A New Product Growth Model for Consumer Durables", *Management Science*, Vol. 15, No.5, January 1969.

Demming, W. E. (1986). *Out of the Crisis.* Cambridge: Cambridge, Ma: MIT Press.

Ernst and Young (1991). "International Quality Study – Top Line Findings" and "International Quality Study – Best Practices Report" Ernst and Young/American Quality Foundation.

Forrester, J. (1969). *World Dynamics.* Cambridge: The MIT Press.

General Accounting Office (1991). US companies improve performance through quality efforts. GAO/NSIAD-9-190 (2 May).

Hammer, H. and J. Champy (1993). *Re-engineering the Corporation.* New York, N.Y.: Harper Collins.

Harte, S. (1992).*Corporate Style.* Atlanta Journal & Constitution. 11 October, R1.

Homer, J. (1987). A Diffusion Model with Application to Evolving Medical Technologies, *Technological Forecasting and Social Change*, 31, 197-218.

Jacob, R. (1993). TQM: More than a dying fad? *Fortune.* 18 October, 66-72.

Kaufman, R. (1992). Why Operations Improvement Programs Fail: Four Managerial Contradictions, *Sloan Management Review*, Fall, 83-93.

Kaplan, R. (1990a). *Analog Devices: The Half-Life System*, Case 9-191-061, Harvard Business School.

Kim. D. and G. Burchill (1992). System Archetypes as a Diagnostic Tool: A Field Based Study of TQM Implementation. *Proceedings of the 1992 International System Dynamics Conference.* Utrecht: University of Utrecht, 311-320.

Kofman, F., N. Repenning and J. Sterman (1994). Unanticipated Side Effects of Successful Quality Programs: Exploring a Paradox of Organizational Improvement, Working Paper #3667-94-MSA, Sloan School of Management, Cambridge, MA 02142.

Morecroft, J. (1985). Rationality in the Analysis of Behavioral Simulation Models. *Management Science* 31 (7): 900-916.

Newman, G. (1991). The case against quality. *Across the Board* (The Conference Board). June, 58.
Paich, M. and J. Sterman. (1993). Boom, Bust, and Failures to Learn in Experimental Market *Management Science*, 39(12), 1439-1458.
Schneiderman, A. (1988). Setting Quality Goals, *Quality Progress.* April, 55-57.
Shiba, S, D. Walden, and A. Graham (1993). *A New American TQM. Four Practical Revolutions in Management.* Portland, OR: Productivity Press.
Stata, R. (1989). Organizational Learning— The Key to Management Innovation, *Sloan Management Review*, 30(3) Spring, 63-74.
Taylor, P. (1992). Such an Elusive Quality. *Financial Times.* 14 February, 9.
Varian, H. (1992). *Microeconomics Analysis.*, New York, N.Y.: W.W. Norton & Co.