Meanings, Measures, Maps, and Models:

Understanding the Mechanisms of Continuous Change

Nelson P. Repenning

Department of Operations Management/System Dynamics Group
Sloan School of Management, E53-339
Massachusetts Institute of Technology
Cambridge, MA USA 02142
Phone 617-258-6889; Fax: 617-258-7579
E-Mail: <nelsonr@mit.edu>

November 2000

Version 1.0

Work reported here was supported by the MIT Center for Innovation in Product Development under NSF Cooperative Agreement Number EEC-9529140. Many thanks to Tim Tiernan, Bill Colwell, Laura Cranmer, Dave Lazor, Vic Leo, Frank Murdock, Roger Saillant and Ron Smith for their generous assistance. Thanks to Paul Carlile, Laura Black, Andrew Jones, Don Kieffer, Brad Morrison, Scott Rockart, Jenny Rudolph, John Sterman, Karl Weick, Wanda Orlikowski and seminar participants at MIT for helpful comments.

For more information on the research program that generated this paper, visit our World Wide Web site at http://web.mit.edu/nelsonr/www/.

Meanings, Measures, Maps, and Models:

Understanding the Mechanisms of Continuous Change

Abstract

There is now considerable controversy concerning the role that incremental change plays in the process of organizational transformation. Some scholars assert that incremental change is the primary source of resistance to more radical re-orientations, while others argue that on occasion, ongoing incremental change can produce dramatic transformation. To help reconcile these competing perspectives, in this paper I report the results of an inductive study of one firm's successful attempt to improve continuously and incrementally its core manufacturing process. The principal results of this effort are: (1) to challenge the current view of the source of change in process-oriented improvement initiatives; and (2) to offer an alternative characterization of the mechanisms through which competence-enhancing, incremental change actually occurs. The theory emerging from this analysis provides one path to resolving the dilemma posed by incremental change processes that can, on occasion, produce organizational transformation, but more often limit the organization's ability to adapt to its environment.

Meanings, Measures, Maps, and Models:

Understanding the Mechanisms of Continuous Change

Introduction

A recent development in the ongoing discourse on organizational transformation is the distinction between episodic and continuous models of change (Weick and Quinn 1999). While episodic models such as population ecology and punctuated equilibrium have historically dominated discussions of and theorizing about transformation (e.g. Hannan and Freeman 1984, Tushman and Romanelli 1985), recent studies suggest that transformational change can, on occasion, be the outcome of a more continuous process. In this view numerous small modifications to routines and structures accumulate over time to produce dramatic shifts in capabilities and orientation. Recent examples include Orlikowski (1996), who studies the changes induced by new information technology, and Brown and Eisenhardt (1997), who study the ability of firms to continuously adapt to the rapidly changing computer industry.

An interesting feature of this debate is that both views posit ongoing incremental change as a near constant feature of organizational life (Weick and Quinn 1999), but differ significantly on the question of whether these small changes can accumulate to produce dramatic transformation. Those favoring population ecology and punctuated equilibrium suggest that incremental changes targeted at improving effectiveness also reduce the ability of an organization to adapt to changes in its environment. In contrast, others argue that this is not necessarily the case. Incremental change, these scholars suggest, can on occasion accumulate to produce remarkable shifts in capabilities and orientations (e.g. Brown and Eisenhardt 1997, Orlikowski 1996). Thus, Romanelli and Tushman (1994:1144) hypothesize that "…small changes in individual domains of organizational activity will not accumulate incrementally to yield a fundamental transformation," while in contrast Orlikowski (1996: 89) argues that on occasion small changes can be "…repeated, shared, amplified, and sustained [and] can, over time, produce perceptible and striking transformation." In their review, Weick and Quinn (1999: 381) summarize this dilemma by writing, "Our review suggests both that change starts with failures to adapt and that change never starts because it never stops. Reconciliation of these disparate themes is a source of ongoing tension and energy in recent change research."

The competing perspectives on the role of incremental change highlight an important gap in the current discourse on transformation. While it plays a central role in theories like population ecology and punctuated equilibrium, until recently incremental change has received little attention from organizational scholars. As Brown and Eisenhardt (1997:1) write, "…incremental change is assumed to occur, radical change is the focus of interest." A particular limitation of the current literature is the dearth of studies that examine this phenomenon at the level of specific modifications to technology and routines. Historically, scholars have studied these changes from a more distant perspective, focusing on a set of changes or a change initiative, and have not delved deeply into the micro processes through which specific incremental changes are actually made (Orlikowski and Tyre 1994 and Macduffie 1997 are some notable exceptions). Until theorists develop a detailed understanding of the mechanisms that create incremental change, it will be difficult if not impossible to understand its role in either facilitating or impeding organizational transformation.

To help fill this gap, in this paper I report the results of an inductive study of one firm's successful attempt to improve continuously and incrementally its core manufacturing process. Process improvement provides a useful vehicle for studying continuous change for a number of reasons. First, process improvement advocates (principally Deming 1986) were among the first to recognize the utility of ongoing, incremental change. One of Deming's fourteen points is to "…improve continuously and forever the system of production." Thus, incremental change is close to the surface in such efforts, enabling a more direct examination of how managers attempt to guide it and how participants react to it. Second, such efforts, when they succeed, often produce change far more widespread than their promoters initially anticipate. Total Quality Management (TQM), for example, while originally envisioned as a system for controlling quality has, in some cases becomes a management system capable of inducing remarkable organizational transformation. Third, such techniques are primarily used in more stable industries by existing players. Thus, in contrast to the environments studied in Orlikowski (1996) and Brown and Eisenhardt (1997), process improvement allows the examination of ongoing incremental change in organizations residing in environments that do not necessarily require rapid adaptation.

The theory reported here was developed inductively based on an original case study of a process improvement effort within a division of a US automobile manufacturer. Understanding the sources of continuous change was, however, not the initial focus of the research effort. Instead, the data were initially collected with an eye towards understanding the failures of process-oriented change techniques such as TQM. A grounded-theory approach was chosen for this question based on its likelihood of producing novel insights into the organizational dynamics that conspired to prevent the successful use of process-focused techniques (Eisenhardt 1989). Underlying the initial research design was the implicit assumption that the mechanisms through which these techniques produced changed were fairly clear and well understood. As the analysis of the data progressed, however, it became clear that this assumption was not useful in explaining the observations. Thus, the study ultimately yielded two related pieces of theorizing, one focused on the understanding the failure of efforts to induce continuous, incremental change (from here forward referred to as the "companion study"), and the theory discussed in this paper, which outlines an alternative view of how such changes actually occur.

The principal result of this effort is to challenge the current view of the source of change in process-oriented improvement initiatives. The existing literature largely attributes the performance gains arising from process-focused change to the use of tools and structured problem solving methods associated with a given technique. This is evidenced both directly (e.g. Dean and Bowen 1994, Sitkin et al. 1994, Hackman and Wagmen 1995, Zbracacki 1999) and by the fact that tool usage is often a proxy for the success of a given effort (e.g. Westphal et al. 1997). In contrast, my data suggest that the actual tools and methods associated with a given technique play a less significant role in creating performance improvement. Instead, the primary driver of improvement appears to be a change in the set of assumptions that participants within a given process make about the determinants of its performance. Underlying each improvement technique lies a set of assumptions concerning the determinants of system performance. This model (not the associated tools which are simply manifestations of it) is the ultimate source of improved performance. Furthermore, the mechanism that creates change in this setting is not the rote application of a specific tool or method. Instead, improvement is more profitably viewed as a process of sensemaking in which confronting actual practice with that suggested by the alternative model reveals improvement opportunities which can then be exercised via distributed actions such as experimentation and improvisation.

This interpretation has implications far beyond the specific application to techniques like TQM. Most important, it suggests that theorists favoring the episodic view may have been too casual in asserting a strong trade-off between improving efficiency and maintaining the ability to adapt to changes in the environment. While this basic trade-off is at the heart of both the population ecology and punctuated equilibrium perspectives, its presence is prominently missing in both my data and the resulting theory. In fact, in the initiative that I study, improved effectiveness was accompanied by a reduction in the control structures, frequent and granular measurement schemes, and carefully delineated and monitored tasks normally associated with an “efficient” production environment. Efficiency and adaptability, rather than being substitutes, appear to be complements.

This interpretation does raise an important question: If efforts to improve effectiveness within an existing orientation are not necessarily at odds with the ability to change that orientation, why do organizations so often appear to sacrifice flexibility and adaptability in favor of improved productivity and efficiency? The resolution of this seeming paradox comes in recognizing that, viewed from the level of specific incremental changes, efforts to improve efficiency typically entail one of two distinct activities: exploiting opportunities in the structure of the physical process; and, appropriating the time, energy, and attention of those working in that process. Combining the results of this effort with the companion study suggests that, while efforts to exploit opportunities in a particular technology require little sacrifice in adaptability, a focus on exploiting people significantly limits an organization's ability to adjust to changes in its environment. By highlighting two different types of exploitation and their different mechanisms for producing improvement, the theory developed here provides one path to resolving the dilemma posed by incremental change processes that can, on occasion, produce organizational transformation, but more often limit the organization's ability to adapt to its environment.

The paper is organized as follows. In the next two sections I discuss the research methods and then briefly summarize the most striking features of the data. The theory that emerged from the analysis is then described. The paper concludes with discussion of and reflection on the implications of the theory for future work.

Methods

The initial focus of the research was to understand the causes and consequences of organizations' failed attempts to improve their core product development and manufacturing processes. Like much of the existing literature, at the outset the effort was premised on the (at the time, implicit) assumption that the mechanisms through which process improvement initiatives actually led to improved performance were well understood. Specifically, I assumed that the primary source of improvement lay in the use of the tools, routines, and practices associated with a specific initiative. As a consequence, much of the early inquiry focused on understanding why, in certain situations, the tools and methods associated with the initiative in question were either not used or not used properly. As the data analysis progressed, however, it became clear that the assumption of the tools as the primary source of improvement was contradicted by the data. Thus, the theory building was separated into two complementary components. The first focused on the initial question, "why do initiatives fail to produce change?" and is reported in the companion study. The second, reported in this paper, focused on developing an alternative theory of how process-focused techniques actually produce change. In the remainder of this section, I describe the initial research design as well as the events that led to the additional theorizing reported here.

Research Design

The research was conducted within a one division of a major automobile manufacturer. The division designs and manufactures electronic components that are later integrated into vehicles at separate body and assembly facilities. At the outset of the research, a team of people working at various levels within the division was assembled to assist in performing the research. The team included a division vice-president, people who had (or previously had) line responsibility in assembly plants and product development projects, and members of various internal consultancies that provided services to development projects and plants.

The research design focused on extreme or polar types (Eisenhardt 1989:537), and was begun by identifying several initiatives that were either dramatic successes or dramatic failures. Given the initial focus on the determinants of failure, the team expected that the processes of interest, those that prevented successful improvement, would be closer to the surface in dramatic successes and failures, and that comparison of these extreme examples might yield interesting insight. A large number of change initiatives were reviewed as candidates for study. Additional selection criteria included the size and scope of the effort, the level of investment in made by senior leadership, and the subsequent impact on organizing practice.

Extended discussions among team members (who typically had substantial experience with the initiatives under consideration) led to choosing three initiatives for study: a cycle time reduction/quality improvement effort in manufacturing (the MCT initiative); a cycle time reduction/productivity improvement initiative in product development (the PDP initiative); and an internal supplier certification initiative patterned after the Baldrige award (the Total Quality Excellence (TQE) initiative). The MCT and PDP initiatives provided a unique opportunity to study process improvement for a number of reasons. First, the MCT effort was dramatically successful, leading to among other things, a twenty-fold reduction in average manufacturing cycle time, while the PDP initiative failed to achieve most of its initial objectives, thus providing a sharp contrast in results. Second, despite these differing outcomes, the same senior executive (the general manager of the division) led the two efforts, and the second effort, PDP, was initiated in large part due to the success of the first. Thus, their comparison offers a rare opportunity to "control" for the impact of senior leadership and management style. Third, of the two, PDP received substantially more financial support, had a dramatically larger communication and training budget, involved a much larger group of people supporting and promoting the effort, and generally conformed more to the conventional wisdom concerning successful large scale organizational change.

After initial interviews, the third initiative was dropped from this component of the study because the data revealed that it was more focused on documentation than on actual process improvement. Readers interested in this initiative can consult Johnsson (1996).

Data Collection

The primary data collection method was semi-structured interviews. Over sixty interviews were conducted with participants in the two initiatives. I was given wide access to all levels within the organization; interviews included the most senior person in the division (the general manager), the executives in charge of both initiatives, plant managers whose facilities adopted the MCT program, product development vice-presidents and business unit managers, line supervisors in the manufacturing plants, managers whose projects piloted the PDP effort, operations and manufacturing engineers in plants, product engineers, machine operators, and material handlers. Interviews were conducted at two manufacturing facilities in North America, the product development and research center, and the corporate headquarters.

Interviews lasted between 45 and 90 minutes and were all recorded on tape. Interview subjects were each presented with a one-page outline of the topics I wished to cover. Each interview began with the subject describing his or her background with the organization and any relevant previous experience. Participants were then asked to give a detailed account of their experience with the initiative. Subjects were asked to assess the key successes and failures of the initiative and to offer their personal hypotheses for their causes. Finally, subjects were asked to describe any lessons learned and to speculate on what they would do differently if they were to participate in a similar initiative in the future. As soon as practically possible (usually the same evening), I listened to the interview tape, reviewed the notes taken during the interview, and wrote a detailed summary of the interview and initial reactions. Later the interview tapes were transcribed.

The interviews were supplemented with extensive collection of archival data. I was given access to a wide range of promotional and training material associated with each initiative including pamphlets, newsletters, instructional books, and video and audiotapes. Historical performance data were also collected. In the case of the MCT effort, extensive data on actual cycle times, product quality, productivity and other operational variables were available. Fewer data were available for the PDP effort.

Data Analysis

The data were analyzed using the traditional techniques suggested by Eisenhardt (1989) and Miles and Huberman (1984). I began by reviewing all the archival data, the interview transcripts, the notes taken during the interviews, and the post-interview summaries. From these sources, I developed a time line for each of the initiatives. The time line then provided the basic structure for writing two detailed case histories. The cases describe the history of the initiatives drawing on the quantitative data, archival materials, and recollections of participants, making extensive use of quotations from the recorded interviews. The objective was to describe the multiple views of the participants at critical junctures in the history of the initiative. After they were drafted, all interview subjects were given the case documents to review their quotations for accuracy. They were not allowed to change the content unless there was a factual dispute that could be resolved with additional data. Participants were also asked to review the entire case for accuracy. Case reviews often led to subsequent conversations in which the interviewee would provide additional background or history. The research team also reviewed the cases in detail, highlighting missing elements of the narrative and suggesting additional interview subjects or data collection. The cases are available from the author upon request.

With the cases in hand, I turned to developing a theory that explained the evolution of each initiative. The data appeared to be usefully divided into three time periods. First, there was the period prior to the MCT initiative in which there were several attempts to improve the manufacturing process, none of which yielded significant results. The second phase began with the MCT effort, which produced dramatic success. Finally, there was the third phase, which started with the PDP effort. Initially, the most useful insights into the dynamics of failed efforts to improve–the initial research question–arose from comparing the pre-MCT phase (during which many improvement efforts failed) and the PDP phase. The data showed striking similarities across the two phases despite the vast differences in context. Building on this analysis, working with a co-author, I developed the theory of process improvement failure reported in the companion study.

Given its striking success, an important component of the initial theorizing was to explain how the MCT effort was able to overcome the failure modes suggested by the emerging framework. As this portion of the theory building proceeded, however, it became clear that the existing understanding of the dynamics of successful change was inadequate. At this point, I focused my attention, again following standard inductive practice, on developing a more complete theory to explain MCT's success. This phase was driven by three components. First, there was the MCT data. Second, there was the existing organization theory on process improvement (principally TQM, see Dean and Bowen 1994, Hackman and Wagemen 1995, Westphal et al. 1997, and Zbaracki 1999), and, third, there were the prescriptive writings of the proponents of various process-based improvement techniques such as Deming on TQM (1986), Goldratt on theory of constraints (Goldratt and Cox 1986), and Hammer and Champy on re-engineering (Hammer and Champy 1993). Starting from an alternative premise suggested by the data, I iterated between the academic literature, the prescriptive literature, and the MCT narrative to develop an alternative framework for thinking about the mechanisms that create incremental change. In what follows, I briefly describe the most striking features of the MCT experience and then outline the theory that emerged to explain them.

Contradicting Existing Interpretations

A newly hired general manufacturing manager (GM) launched the MCT initiative. He had previously worked at a major computer manufacturer, and thus was familiar with manufacturing electronic components. The MCT initiative was dramatically successful. In the course of approximately five years, the division was able to reduce its average manufacturing cycle time from over twenty days to less than one day. The reduction in cycle time alone proved to be an important source of competitive advantage as it allowed the plants to react to changes in customer requirements much more rapidly. The accompanying reduction in inventory cut carrying costs and freed up so much floor space in existing plants that at least two of five planned new facilities were not needed, thereby saving hundreds of millions of dollars in capital expenditures. Further, product quality, revenue and profit also improved dramatically during the same period. The effort was so successful that it spawned a company-wide effort to duplicate the success in other divisions that is on going as of this writing.

Yet, despite these successes, the history of the MCT effort as reported by participants offers at least three striking contrasts to both the existing academic and prescriptive literature on process improvement. First, while existing analyses of TQM (there are few treatments of other improvement techniques) posit the usage of the associated tools and structured problem solving methods as the source of improvement (e.g. Dean and Bowen 1994, Hackman and Wageman 1995, Zbaracki 1999), the early phases of the MCT effort showed an almost complete lack of their use. Instead, the early efforts used few tools or structured problem solving methods and were highly informal. For example, one team leader characterized his early efforts at inventory and cycle time reduction:

> We would look at the piles sitting around and ask “do we really need these?” and the answer was usually no. It really didn’t involve that much effort.

Similarly, another explains her approach:

> ...we always would focus on “ look around you, where do you see parts and why are they there?” We didn’t spend a lot of time on the exact calculation of buffer sizes. We didn’t worry about whether the buffer should be x units or y units. It looked like a lot of units, so we started by trying to cut the pile in half. Then we would try to cut it in half again, and so on.

In both cases despite the informality of the approach, dramatic improvement ensued. The ease and simplicity of making improvements in the early phases of the initiative was widely reported.

One plant manager summarized the early efforts in his plant (in which cycle time was cut by over 50%) by saying, "In the first year, things came pretty easily."

Second, while both interpretations of TQM (e.g. Sitkin et al. 1994) and the larger literature on episodic models of change (e.g. Tushman and Romanelli 1985) have argued that improved efficiency and effectiveness come largely through standardizing routines and exerting more control over production processes, there was little evidence of such efforts in the MCT experience. Instead, while there were eventually significant improvements in product quality, in the early phases of the initiative there was scant emphasis on the consistent use of methods or process standardization. In fact, one plant manager reported, “At that time” he recalls ”...we didn’t standardize anything except how we made the measurements.” Further, while new techniques are often explicitly contrasted with a previous approach to demonstrate their superiority (for example see Hammer and Champy's 1993 comparison of reengineering and TQM), the MCT effort successfully utilized numerous techniques and the eventual process represented a combination of different approaches to managing manufacturing. For example, one engineer describes how his plant's approach to manufacturing evolved out of the combination of two different methods:

> ...[we have] developed a hybrid Kan Ban TOC [Theory of Constraints[1]] system. TOC says ‘keep the constraint running at all times’, Kan Ban says ‘limit the amount of material you have on the floor’. We try to do both.

Third, while proponents of process-focused techniques often depict their methods as steady-state approaches that are eternally useful, the MCT effort and each of its constituent approaches showed a definite life-cycle. Initially the use of a specific technique was ad hoc and informal, but often led to dramatic improvement. Later the usage often became increasingly sophisticated: measurements systems became more

[1] . Theory of Constraints is a production scheduling and improvement methodology discussed in Goldratt and Cox (1986). It will be discussed in more detail later in the paper.

detailed; mapping procedures became institutionalized; and complex methods such as the theory of constraints were used for production scheduling. Despite the increasingly sophisticated usage of various approaches, however, the rate of improvement slowed, causing some to question its utility for creating additional change.

Thus, in the analysis of the data and the focus on failure modes, I was confronted with an initiative that produced dramatic success, yet contradicted both the academic interpretation of improvement and the prescriptions of improvement gurus on at least three important dimensions. Improvement did not appear to stem primarily from the use of tools or structured problem solving methods. Improvement did not come through standardizing routines or increasing control over the activities of the workforce. No single technique or philosophy was used, nor did it become the "steady state." Instead there was an overall lifecycle to the effort and there were smaller cycles associated with each specific technique. In the next section I describe the theory that emerged to understand and explain these features of the MCT experience.

Theory

Alternative Conceptions of System Performance

To understand the success of the MCT effort, I begin with the most obvious and significant difference between it and the previous efforts to improve performance: the new general manager. More specifically, the data strongly suggest that at the time of his arrival, there was a significant difference between the GM's approach to understanding manufacturing and that of his new colleagues. Prior to the MCT effort, the division's manufacturing operations were operated in a manner similar to that of other companies whose business requires substantial capital investment and labor expense. Supervisors were charged with keeping each piece of equipment and each member of the workforce busy at all times. For example, one operations manager recalled:

> Before [MCT] if you were to walk out onto the floor and ask a supervisor how things were going, he would say “Great, all my machines are running” and you would see tons of WIP sitting around.

The measurement and incentive systems supported this approach by emphasizing direct labor performance (roughly defined as the number of units produced per person per day) and managers imposed heavy penalties on those who missed their daily targets. As an operations manager at one plant recalled, “...supervisors would always hit their exact targets, if the goal was 200, they would pack [produce] 200, never 198, never 202.” Another noted “...[supervisors] would make sure everybody was busy all the time to make labor efficiency” and that “... supervisors who missed their targets knew they were going to get ‘beat up’ by their managers.”

Due to its long history in the auto industry, the utilization-based approach had become deeply embedded in the technology, routines, and behaviors that constituted the manufacturing system. Further, and importantly, the data also suggest that, through a continually evolving set of self-reinforcing adaptations, participants at all levels within the organization played a role in both creating and maintaining that focus. Those in the finance and accounting function developed increasingly sophisticated schemes to measure utilization, while Vice Presidents legitimated and reinforced this focus by reviewing those measurements, sometimes at the level of specific machines, on a regular basis. Similarly, supervisors and operators made numerous local modifications to the process to insure that those objectives were achieved. Most significantly, they held significant stocks of work-in-process inventory to ensure that breakdowns and quality problems at upstream machines did not interrupt production at downstream machines. As one supervisor said:

> Supervisors at that time were evaluated on labor performance on a daily basis. It didn’t take long for them to develop a buffer in front of their line so that if the schedule called for 700 and their line was fully utilized at 800, they could still run 800 units every day, and still make their labor performance.

Plant managers and manufacturing engineers also supported the utilization-based approach by designing and installing increasingly sophisticated automated storage and retrieval systems to handle the growing piles of work in process. One plant was forced to erect a tent to house excess inventory that could not be held within the main building. The tent, initially planned as temporary storage, quickly became a permanent part of the plant's operation. Finally, as the piles of work-in-process inventory grew and the plants resorted to increasingly expensive methods to handle them, the cost of operating the plants increased, thus justifying an increased emphasis on fully utilizing those assets, and thereby perpetuating the cycle.

The GM also reports both a focus on utilization and a significant gap between his approach to understanding manufacturing and that of his new colleagues. He recalls his initial reaction to the management approach at the division:

> ...in the first series of meetings that I attended, the topic of discussion was the utilization of a particular piece of equipment in our [particular] facility. You know, a particular piece of equipment!…I couldn't figure out why? To me, it was nonsensical, although the executive VP who was asking didn't view it that way.....from [my previous employer, a major computer manufacturer] I was used to monthly reviews on inventory performance, etc. Here each and every plant was being measured on machine utilization, and nobody was looking at inventory, or the time spent actually adding value.

In contrast to a utilization-based approach, the GM believed that a focus on reducing cycle time provided the path to improved performance. He explains:

> The way you make money is to make product and the more product you make, the more money you make… if you can speed up the process, that's making more product. If you can add more value, that's making more product… what I want to know is, can you make more product and can we make more money…I don't want to know about machine utilization at that level.

Other participants, particularly the plant managers, also noted significant differences between the dominant approach to managing manufacturing and that offered by their new boss. One plant manager recalled his reaction:

> The concept of cycle time and value added were not unheard of at [this company]....however, [the GM] had a very different vision about how those measurements could be used to really drive the operation and improve productivity.

While another recalls:

> …[the MCT effort] really began to cause us to think differently about our manufacturing process…with our traditional metrics, we focused on inventory levels and turns, but the speed of our operation was never part of our focus in the plants. MCT was the first time we started to think about the speed of our operation.

Importantly, the data also suggest that the differences in philosophy between the GM and the rest of the division resided largely at the level of the manufacturing *system*. That is, the GM took no issue with the execution of specific tasks or the choice of particular technologies. In fact, given his role as an outsider, he had little to offer on these dimensions. Indeed, he felt that the focus on improving the execution of individual tasks was at the heart of the division's difficulties. He recalls:

> I felt that nobody was looking at our manufacturing facilities as a system. They were looking at pieces, and, as a result, spent their time trying to optimize each piece. The guy running the solder equipment might be running boards through there like nobody's business, but did that improve the output of the system? I'm not sure anybody really knew that.

The data thus suggest that one key component to understanding the success of the effort lies in the GM's alternative perspective on managing the manufacturing system. How that alternative perspective was translated to actual performance improvement was highlighted in the first phase of the MCT initiative.

Phase 1a: Alternative Measurements

The existence of a manufacturing system in which every participant, from the executive vice-president in charge of the division to material handlers on the shop floor, plays a role in continuously re-creating an environment built on a utilization-based model has interesting implications for any attempt to introduce changes derived from another approach. On one hand, changes derived from an alternative model can be potentially powerful if that perspective captures dimensions of performance not present in prior conceptions. On the other, however, such changes can be difficult to induce for at least two reasons. First, they are difficult precisely because they are not consistent with the dominant performance model. Many people in the division had worked under the utilization-based approach for more than twenty years. Second, because every participant played a role in maintaining a focus on utilization, any intervention targeted at one organizational level (e.g. supervisors) could be thwarted by participants at another (e.g. vice-presidents). For example, one participant in the initiative discussed the failure of a previous effort to improve that had focused on those in the materials planning function:

> After working with the materials managers for a little while, it became obvious to me that they were not able to control [the level of inventory]. The manufacturing manager would make decisions that caused [inventory] to go up that the materials manager could not control.

To understand the mechanisms through which the GM overcame these barriers and translated his alternative approach into improved performance, it is helpful to conceptualize process improvement as a process of sensemaking (Weick 1995). Weick (1995) argues that there are two generic antecedents to sensemaking: *uncertainty*, which is created by cues for which there are no existing interpretations, and *ambiguity*, which is created by cues for which there are multiple interpretations. In launching the effort, the GM successfully used *measurement* to create uncertainty and thereby induce participants to reconsider their approaches to managing manufacturing. He describes his initial intervention:

> We analyzed [for a sample product] the time elapsed between when a part came in the back dock until the time it left the shop floor, and asked the questions "How long did it take?", and "What was the value added [fraction of time in which function or feature was being added to the product]?" We found out [for this product] it took 18 days to make the product and we were adding value to the product 0.5% of the time. When I laid this out for everybody...they were astonished.

The results of this analysis were universally reported to be surprising. A number of participants recalled that, prior to this analysis, they estimated the value-added percentage to be on the order of ten percent, a factor of twenty greater than the actual value. Thus, by measuring the system's performance in an entirely new way, the GM confronted managers with a set of cues that challenged their conception of the manufacturing systems they oversaw. For example one plant manager recalled:

> ...we had a gut feel that our cycle times were going to be pretty long...but what really got us was that even with the very crude definitions of value add time we were using–they are much stricter now–we had astoundingly low cycle efficiencies [the ratio of value add to total production time].

The uncertainty created by the new analysis was brought further to the fore when the GM subsequently required each plant to report cycle time on a monthly basis, but provided them with little information as to how it was to be measured. As a plant manager said, “[The GM] didn’t give us a lot of the details...he probably knew more about it than he led us to believe, but I think he wanted us to take a fresh look." Thus, initial efforts in the plants were targeted, not at improvement, but at developing the measurement scheme.

As a sensemaking perspective suggests, however, plant managers, manufacturing engineers, and supervisors were not passive observers of the manufacturing process. While they were charged with developing a measurement system, they were also responsible for running and maintaining the process on a daily basis. The structure of the process, both the physical technology and the standard routines and operating procedures, arose in part from their daily actions. Thus, as the need to develop measurements caused them to understand the process in terms of cycle time, they also began to manage it in terms of cycle time. Improvement followed almost immediately. A plant manager describes the early activity in his facility:

> ...in the first year we started with simple counts at different times during the day, and we started to plot them and to try and understand what was happening. Very quickly our creative engineering personnel came up with clever ways to control the buffers that helped make big improvements.

Nothing in the GM's approach or the measurement scheme itself suggested how to change specific steps within the process to reduce cycle time, yet performance improved, often dramatically. The previous quote hints at the mechanism through which this change occurred. In trying to measure cycle time, which required counting inventory levels at various time intervals, engineers quickly discovered which buffers were large (thus contributing to long cycle times) and which weren't. In other words, a focus on cycle time and the process of attempting to measure it revealed *improvement opportunities*, elements of the process that, if changed in some fashion, would improve performance.

While a focus on cycle time quickly revealed opportunities for reducing it, nothing in the GM's approach or the measurement scheme detailed how to capitalize on those opportunities. Instead, as the manager above suggests, participants "came up" with them. In the following quote, he explains the source of the knowledge required to make the improvements discussed above:

> During the first two years almost everything we tried we picked up from our own people...by giving them the free reign to start to do some things, they were able to go back to textbooks, business cases, friends they had in other areas that were trying similar things, etc. We tried everything from the Toyota Production System’s Kan Ban to doing statistical process control on buffer sizes.

While the alternative model suggested improvement opportunities, due to its system-level focus, it provided little guidance as to how to capitalize on those opportunities. Instead, by revealing opportunities to reduce cycle time, it suggested ways that participants could use existing knowledge and skills in service of the new system wide objectives.

Thus, improvement and change in the MCT case can be understood as a process of sensemaking in which participants, as they began to incorporate elements of the cycle time-based model in their own conceptions of system performance, noticed improvement opportunities, which were then exercised using locally accumulated, task specific knowledge. In the first phase, the process was catalyzed by new measurements (and the subsequent charge to create a system for reporting them) that created uncertainty. While many participants were initially skeptical of the cycle time-based approach, early improvements confirmed its utility and caused participants to take the perspective more seriously. Over time, participants began to both interpret the manufacturing system in terms of cycle time (sensing) and, through improvement activities, enact that perspective within the system's structure (making).

The basic cycle of confronting existing conceptions with an alternative model of system performance, the revelation of opportunities, and the distributed use of locally accumulated knowledge to exercise those opportunities was repeated at least three additional times during the initiative. The first of these started soon after that the measurement system was put in place and provides a second example of the proposed change mechanism at work. It also highlights the role of the second generic antecedent to sensemaking, ambiguity.

Phase 1b: Alternative Meanings

Following the initial data analysis and the subsequent installation of the new measurement system, the GM, in his own words, "…spent 70% of my time on the road. I spent all my time talking to plant managers, letting them know why this was good stuff, and making sure they knew what my expectations were." During these early visits, the GM challenged the meanings of common terminology, thereby creating ambiguity around the existing interpretations of the manufacturing process. For example, he attached a very different definition to the word "cycle time" than did his new colleagues. He recalled:

> One of the first debates we had was over how to measure cycle time. Many people thought of cycle time as the cycle time of the equipment. They were looking at reducing the time a part spent on a particular piece of equipment from 20 seconds to 10 seconds. My feeling was when you are at 18 days who gives a rat’s ass about the cycle time of specific machines.

A plant manager recalls a similar reaction to his early conversations with the GM:

> …he was using words like cycle time and cycle efficiency and at that time I don't think any of us really totally understood what he was talking about.

Just as the uncertainty created by alternative measurements challenged the conventional view of the determinants of performance, the ambiguity created by alternative definitions also caused participants to rethink their approach to managing manufacturing. As cycle time was redefined to encompass the entire process rather than individual machines, work-in-process inventory, a major contributor to cycle time, took on an increasingly negative connotation. While inventory had always been treated as costly by the finance group, it had little influence on the operations of the plants. As one engineer said "...[before MCT] we basically completely ignored the fact that buffers were expensive. We weren't measured against it, so we didn't think about it." Making a similar point, a materials manager said "...we (the materials, planning, and logistics staff) had always been concerned about inventory...however, it was never really embraced outside the material control area." Further, other negative effects of excess inventory (e.g. long cycle times) were absent from the prevailing practice. In contrast, the GM viewed excess inventory as public enemy number one in the fight against long cycle times. He discusses his strategy for communicating his alternative perspective while on his frequent visits to the plants:

> ...I would always go out to the production floor and to the warehouse. I wanted to show them examples of what I was talking about. I might look at the shipping labels in the warehouse. If it were May, I would usually find parts that had been received the previous August, and I would ask, "if you aren't using this stuff until May, why has it been sitting here since last August?" It costs us cash, space, and maybe even quality problems.

By challenging the meanings of common terminology (and therefore creating ambiguity), these interventions also initiated the basic improvement mechanism of confronting existing practice with an alternative model of system performance. For example, as the meaning attached to inventory began to shift towards something undesirable, participants rapidly discovered ways that it could be eliminated. One internal consultant discusses her approach to facilitating improvement:

> First we organized all the inventory in totes...then we color-coded it by good stuff, bad stuff, and stuff we didn't know about. Then, finally, we could see how much inventory we really had, and the condition of that inventory...we also found inventory that some

> people knew about that others didn't. It was shocking…big success came just by making people aware of the piles.

While these changes were exceedingly simple to execute (they required little task specific knowledge), they again highlight the proposed change mechanism at work: an alternative model suggests that if inventory can be reduced, system performance will improve; excess inventory becomes an improvement opportunity when the model is propagated via the redefinition of inventory as undesirable; participants use existing knowledge (in this case a trivial application) to capitalize on those opportunities.

During the first phase, which lasted approximately two years, many plants were able to cut their average cycle time by more than fifty percent. Despite these dramatic successes, there was little use of structured problem solving methods or sophisticated improvement tools. Instead, changes were informal and often involved the trivial application of task specific knowledge. As mentioned above, this experience stands in stark contrast to both standard conceptions of efficiency–which rely on well-defined standard procedures and routines–and current interpretations of process improvement–which posit the use of improvement tools and structured problem solving methods as the primary source of change. In the theory developed so far, standard routines, improvement tools, and structured problem solving methods play little role. The relationship between tools and structured methods and the improvement mechanism discussed so far is highlighted in the second phase of the MCT effort.

Phase 2: Alternative Maps

In the second year of the initiative, the GM created a small corporate staff to oversee the effort. To continue the pace of improvement, the newly created group began by formalizing the process of calculating the value-added percentage that the GM had used in his initial analysis. In this phase the plants were required to create process maps and report the resulting value-added percentages.[2] A staff member explains the approach:

> ...[to calculate the value-added percentage] you had to walk through the shop floor and ask the question, “Is this value added?” for every step in the process. By the time you were finished you had flow-charted the entire process and really highlighted all the value add stations....After calculating [the value added percentage], we really started to understand the process flow of our products. We knew where value was being added, and, more importantly, where value was not being added.

Similar to the changes induced by the new measurements and meanings, the mapping process revealed further improvement opportunities. As the quote above suggests, it highlighted additional buffer inventories that could be reduced and non-value added steps in the process that could be eliminated. It further caused people to challenge the logic underlying the existing production technology. A plant manager offers one example:

> The process made us challenge specifications and engineering requirements that we had previously taken as given. For example, it caused us to [ask] why did we need to protect a circuit board from the outside environment when it [resided within the car]? We finally decided after much thought and experimentation that we didn’t, so we eliminated it [thus saving twelve hours].

The GM also offered his own example:

> We had a line that required a radio to be "burned-in"[operated at elevated temperature] for 16 hours. I asked “Why?”, and they told me because “That’s what engineering told us to do”. Well, we looked at the failure distribution and it seemed like all the failures came in the first two hours. So first we cut it to eight hours, then to four and so on. This saved time, equipment, and money.

As with the measurement and meaning phase, nothing in the mapping process suggested which steps could be eliminated or how the plants might go about choosing them. Instead, it simply highlighted major contributors to cycle time. As the initiative proceeded improvement opportunities became harder to find and the changes necessary to capitalize on them required

---

[2] . The actual measure was called Manufacturing Cycle Efficiency (MCE) and represented the ratio of time in which a function of feature was being added to the product to the total cycle time.

increasingly formal approaches. Thus, the sensemaking activity, which in the previous phase was highly unstructured, became more disciplined and relied on a more sophisticated tool (the process map) to propagate the model. Further, exercising the newly revealed improvement opportunities required more sophisticated methods such as running formal experiments.

Importantly, the utility of the mapping process cannot be understood without reference to the overarching focus on cycle time. This connection is highlighted by the fact that this was not the first time the plants had mapped their processes. For each product, the plants had *already* established routing diagrams capturing the sequence of production steps required to assemble a given product. These documents also constituted a map of the production process. In contrast to the maps created during the MCT initiative, however, the routing diagrams told the reader nothing about the time products spent *between* production steps; they contained no information about the cycle time of the production process. Thus, the mapping tool in this context was a more formal manifestation of the cycle-time based model. Mapping had been used in the past, but until MCT the underlying perspective did not include cycle time as a relevant concern.

Phase 3: Alternative Models

The third phase of the initiative represents the fourth instance of the proposed change mechanism and further confirms the role of tools and structured methods as manifestations of the underlying improvement model. The third phase also highlights a subtle shift in the emphasis of the guiding model and a change in the authority structure of the plants.

In the third year of the effort, just as the previous informal approach began to suggest fewer opportunities, the pace of improvement began to slow. At this point, some of the senior participants felt the need for a more sophisticated method to guide future improvements since many of the easy opportunities provided by the new model had been exploited. In particular, the leader of the effort felt the need to move beyond static representations like process maps towards a better understanding of the dynamics of the production process. To that end, he turned to the philosophy Theory of Constraints (TOC) (Goldratt and Cox 1986). The appeal of the TOC training was both the utility of the underlying model, and the fact that its founder, Eli Goldratt, had developed a training program that relied on computer simulators to develop new intuition about the dynamics of production systems. Once selected, TOC was spread rapidly throughout the division. Each plant manager sent two people to a two-week course offered by the Goldratt Institute. In the ensuing six months, those who attended the course created a two-day version that was given to virtually every manufacturing engineer and supervisor. Following that, a four-hour hands-on version (called the "cookie factory") was developed and delivered to almost every machine operator and material handler in the division.

The TOC phase was interesting for a number of reasons. First, it maintained the focus at the system level. TOC is based on the assumption that within in any manufacturing line one machine (the constraint) has the least capacity and, thus, both provides the basis for production scheduling and helps focus improvement activity. Nothing in the TOC training tells the user how to improve the throughput of the constraint machine. Rather it suggests only that the constraint machine should be the focus of improvement activity. The manager in charge explains the appeal of the TOC approach:

> It [TOC] allowed you to step back and understand the shop floor as a system rather than as a bunch of process areas, particularly if you worked inside of one. Even though your training would lead you to make decisions one way, it led you to a new intuition that helped you make decisions differently.

Second, TOC represented an increasingly sophisticated variation on the theme of cycle-time reduction. In the early days of the effort, cycle time could be easily reduced by eliminating piles of unneeded inventory. As the initiative progressed, however, the inventory that remained was increasingly important in maintaining the throughput of the plant. Further reductions could not be made without a more sophisticated understanding of how inventory affected plant performance. The TOC model provided one method for taking that next step. Thus, as the MCT effort progressed, the system level model guiding practice, while still focused on the basic theme of reducing cycle time, became increasingly sophisticated.

Third, the change mechanism observed earlier again repeated itself except that, whereas in the earlier phases of the initiative many of the changes were made by supervisors and manufacturing engineers, in the TOC phase improvements were often made by the operators. An improvement consultant (who ran the cookie factory course) recalls:

> After the training people would pull me aside and say “hey, cookie factory lady look at this”, and they would show me a pile of inventory that could be reduced or some other opportunity for improvement.

The role the TOC model (combined with the distribution of authority) in generating improvement was further highlighted by another supervisor who had recently been transferred to a production line struggling to hit its daily production schedule:

> We started the effort by trying to find the constraint. It turned out that the previous supervisors didn’t know which machine was the constraint. As a result, a lot of [operator's] time was spent keeping non-constraint machines running while the constraint was idle.

In this case, the TOC model allowed participants to use locally accumulated knowledge about keeping machines running more effectively by focusing that knowledge on the machine that limited the throughput of the process. By incorporating TOC notions in their improvement efforts, the area in question was able to improve its schedule performance from less than seventy percent to almost one hundred percent in less than six months. The supervisor running this area during the period of these improvements attributes much of the success to giving operators a thorough understanding of the TOC concept:

> Some of the big improvements did not come until we got the mind-set into the operators and the materials handlers....We started by teaching each of the work teams how to

> manage their line using TOC...the classes were useful, but I felt the real learning came from working with them on their lines on the floor.

Similarly, among those interviewed, there was widespread consensus on the utility of delivering the training to all levels in the division. One engineer remarked:

> The extensive training made a big difference. Nobody in this plant could say that they didn't have the opportunity to learn about TOC. I thought, at the time, it was a waste to train 1,500 hourly workers, but it really helped.

Finally, the role of the TOC model in both revealing improvement opportunities and facilitating the alternative use of existing knowledge is further highlighted by the fact that the division employed a group of manufacturing simulation specialists who had already developed scheduling and coordination strategies that *outperformed* those provided by Goldratt. Further, the assumption of a fixed constraint machine was clearly incorrect for many of the division's production lines. Yet, the TOC training was highly influential and produced significant improvements. TOC produced results where the more sophisticated systems failed because it could be easily communicated to those actually running the production line, the supervisors and the machine operators. As the leader of the simulation group said:

> I called it 'Shop Floor Scheduling and Coordination Awareness 101'. If you wanted to concentrate in three days everything you would want to understand about the dynamics of the shop floor and how to keep the line running, this was it.

Phase 4 and Beyond: The End of the Initiative

As the division's cycle time approached one day, additional improvements became increasingly difficult to make. After more than five years of steady improvement, the changes suggested by a focus on cycle time and the theory of constraints created only marginal improvement. The highly simplified frameworks used to guide improvement activity abstract away from much of the system's structure, and, as the process improved, those features left out of the model became increasingly constraining. A manufacturing engineer provides an example,

> Where do you want your MCT to go? You want to be as responsive as possible to your customers.... At [our plant], if you give us the raw material in the morning, the parts will

> be done that night....However, the biggest problem right now is how fast do we turn over raw materials. Our logistic system works in one day buckets, so cutting MCT any further isn't really going to help us serve the customer any better. There are other measurables which might drive bigger improvements now.

The declining utility of the cycle-time based model was compounded by the fact that, with the success of the effort, it became increasingly embedded in the organization's structure, routines and practices. For example, cycle time is still a key measurement that plants are required to report. Another supervisor explains the dilemma created by the success of the effort:

> Sometimes these objectives start to take on a life of their own. People quit thinking about them, or asking whether they make sense. Now when people ask for further reductions I always ask "Why?"... We need to keep it in the back of our minds so it doesn't start to grow, but we shouldn't use it as a driver anymore.

Another expressed a similar sentiment:

> We can't target MCT directly anymore. We've reduced all the obvious inventory. Now we need to attack the problem from a different angle.

Discussion

The theory that emerged from the analysis of the MCT initiative can be summarized as follows: The five year process of eliminating excess inventory, reducing cycle time, improving quality and increasing productivity began with the GM's introduction of an alternative conception of system performance. By replacing the then-dominant utilization-based approach with a new model based on cycle-time, he offered a radically different perspective on managing manufacturing. The propagation of that model and the ensuing improvement can be usefully represented as a process of sensemaking in which participants, as they came to share elements of the alternative conception, noticed improvement opportunities, and exercised them through improvisation, experimentation, and the alternative use of locally accumulated knowledge. Due to the significant gap between the cycle-time based approach and previous practice, improvement activities were initially informal and undisciplined, but highly productive. Later, to continue the pace of improvement, increasingly formal and disciplined applications of the underlying model were used to identify additional opportunities and increasingly sophisticated methods were required to capitalize on them. Further, due to its success, the cycle-time based approach became increasingly embedded in the technology and routines that constituted the manufacturing system. Finally, as factors outside the cycle-time based approach became increasingly constraining, the rate of improvement slowed.

Beyond the implications for understanding process improvement (many of which have already been highlighted), the analysis also offers a potential contribution to the larger discourse on modes of organizational change. As highlighted in the introduction, the question of whether or not organizations have the ability to continuously transform themselves turns on the role of incremental changes targeted at improving efficiency. Those favoring episodic models suggest that ongoing incremental improvements to effectiveness also create structural inertia and convergence, thus limiting the ability of organizations to adapt to changes in their environments. In contrast, those favoring a more continuous view assert no such trade-off and thus argue that, on occasion, incremental changes can accumulate to produce transformation. In what follows, I use both the MCT data and the resulting theory to first challenge an assumption central to the episodic view–that efficiency and adaptability necessarily represent a strong trade-off–and then, building on a distinction arising from the companion study, propose an alternative premise that provides one path to reconciling these competing perspectives.

The argument underlying the efficiency/adaptability trade-off is a familiar one. Organizations that wish to survive in competitive niches must develop high degrees of effectiveness. Effectiveness results from the ability to regularly and reliably produce a specific output (Hannan and Freeman 1984), and this is achieved through: (1) highly standardized, well defined, institutionalized, and tightly enforced routines and procedures; and (2) a tightly coupled network of resource commitments and value relationships (Tushman and Romanelli 1985). Thus, incremental change in this view is largely a process of "…achieving a greater consistency of internal activities with strategic orientation (Tushman and Romanelli 195: 178) and of "…institutionalization and…creating highly standardized routines (Hannan and Freeman 1984: 154)." Of course, the argument continues, standard and institutionalized routines are also antithetical to adaptability and, thus, survival requires sacrificing any possibility of ongoing adaptation (Hannan and Freeman 1984). The logic underlying these arguments is sufficiently compelling that theorists rarely give it significant attention. Yet, when applied to specific instances of refining a complex production technology, there are some important subtleties.

The MCT initiative was clearly targeted at improving effectiveness within an existing orientation, and it did so dramatically. Cycle time fell by a factor of twenty, quality and delivery performance improved, and productivity grew dramatically. Yet, despite these improvements in efficiency and effectiveness, there is little evidence that the organization became commensurately less adaptable. Indeed, using the logic of the population ecology and punctuated equilibrium, a strong argument can made that the organization's ability to accommodate changes in its environment was enhanced as a consequence of the effort. For example, during the initiative a complex, computer-based production scheduling system run by a central authority was replaced with a more flexible, visual system run by the operators themselves. Similarly, the measurement system became less granular and the daily activities of operators were put under less scrutiny. In both of cases, closely specified, tightly monitored, and highly institutionalized procedures were replaced with simpler, more flexible systems and more distributed discretion. Furthermore, the widely shared, deeply embedded, and stable belief in the value of maintaining a high degree of utilization was replaced with a looser, evolving focus on cycle time reduction. Thus, in the data, despite the dramatic improvements in performance there is little apparent trade-off between efficiency and adaptability.

The resulting conception of incremental change also suggests a weak relationship between efficiency and adaptability. If the primary source of improvement is the gap between a system-level perspective on performance and actual practice, then standard routines and institutionalized procedures play only a supporting role in producing improvements in effectiveness. They provide a reliable baseline from which to identify gaps, and they reduce the cognitive load by allowing participants, at any given moment, to reconsider only a subset of all the activities in on-going improvement activity. In contrast to the population ecology and punctuated equilibrium views, in which reproducible routines themselves are the *source* of improved effectiveness, here they are little more than temporary placeholders or null hypotheses. Further, if the primary engine of improvement is the alternative use of locally accumulated knowledge, then there is little value in an organizational structure that enforces tight compliance to those standards that do exist. In contrast to the episodic view, in this conception the factors that contribute to ongoing improvement in efficiency–simple routines and procedures that are constantly subjected to revaluation and a distributed authority structure that emphasizes local action–rather than preventing adaptation to environmental change, would seem to very much support it. More generally, if a process of confronting existing practice with alternative possibilities drives competence-enhancing change, then the distinction between internally and externally focused change (which is central to the punctuated equilibrium view) loses some of its analytical power. Whether it is driven by a new technology (Orlikowski 1996), turbulence in the competitive environment (Brown and Eisenhardt 1997), or an alternative perspective on internal performance, in all cases change arises from a sense that opportunities for improving effectiveness exist coupled with an organizational structure that allows people the latitude to capitalize on those opportunities.

Positing a weak link between on-going improvements in efficiency and the ability to adapt to change in the environment does, however, raise an important question. While it is now well documented that on occasion ongoing incremental change can accumulate to produce transformation rather than convergence, it is also equally clear that most organizations, most of the time, are not able to maintain this desirable state of affairs. A large and growing body of data suggests that most attempts at even incremental change fail (e.g. see Easton and Jarrell 1998 for a study of attempts to implement TQM and Pfeffer and Sutton 2000 for a more general discussion). Few organizations are able to constantly adapt and avoid the need for periodic radical change (Romanelli and Tushman 1994). These dynamics were prevalent prior to the MCT effort when numerous attempts at change were made and none succeeded. Seriously entertaining the argument that there is only a weak trade-off between efficiency and adaptability (thus allowing the possibility of continuous and ongoing transformational change) requires explaining why in most cases a focus on efficiency and effectiveness appears to be antithetical to more dramatic transformation.

The discussion so far has centered on the processes through which organizations exploit opportunities to improve their production technologies. This usage of the word "exploitation" (and attendant concepts such as refinement, improvement, and efficiency) corresponds closely to both the definition outlined by March and colleagues (e.g. Levinthal and March 1981, March 1991), and that which has appeared widely in organization studies. A key result of the companion study, however, is to suggest that, while its use and meaning are widely accepted, the word "exploitation" should probably be used with more care than it currently receives. Two types of exploitation are present in the data: that directed at finding and exercising improvement opportunities in the physical technology (exploitation in the sense of March) and that directed at appropriating the time and energy of those working within that technology (exploitation in the sense of Marx, see Marx 1887).

The distinction is important because the two modes appear to be fundamentally incompatible. As was observed prior to the MCT effort, when managers choose appropriating peoples' time and energy as the path to improved effectiveness, likely actions include developing more frequent and granular measurement schemes, setting more aggressive objectives for those measurements, and imposing more stringent penalties for missing those objectives. In addition, managers may also impose more direct controls on the workforce's activities via monitoring, increasingly detailed reporting requirements, and technology such as software that monitors the keystroke rate of those entering data. Each of these amounts to restricting the freedom that people have in executing their tasks. Pathological self-reinforcing dynamics often arise in this situation because many of these actions, at least in the short run, produce modest gains in performance. Early gains provide powerful evidence confirming the managers' initial belief that exploiting people is a viable path to improved performance. This “self-confirming attribution error” creates a powerful reinforcing cycle that drives managers to implement increasingly rigid control structures.

The dynamics generated by a focus on exploiting people create the look of structural inertia and convergence. The production technology remains largely unchanged. The attention of managers is directed at standardizing tasks and then monitoring compliance to those standards. Increasingly sophisticated technologies, ancillary to the main production flow, institutionalize those standards. Authority and latitude for local change are systematically diminished. Importantly, however, while these activities produce an organization that is highly resistant to adaptation and change, they have little to do with exploiting opportunities to improve or refine the underlying production technology. Instead, they are quite effective in preventing them. Due to highly decomposed tasks, closely monitored standard modes of execution, and highly institutionalized routines and procedures participants have little opportunity to experiment or use local knowledge in new ways, thus eliminating any possibility of improvements to the underlying technology. A focus on exploiting people effectively eliminates the organization's ability to exploit its technology.

Ironically, while the managerial interventions arising from a focus on exploiting people are targeted at eliminating the freedom to make localized change (thus focusing more attention on production activities), more often than not, they simply redirect its focus. The data from both initiatives clearly shows that the increasingly aggressive objectives and increasingly constraining measurement systems that arise from a focus on exploiting people eventually place participants in an untenable position; they cannot achieve their objectives within the imposed constraints. When this occurs, localized changes are increasingly directed at circumventing the measurement and control systems, thus giving the appearance that objectives are being achieved when in fact they are not. In the pre-MCT days, machine operators often surreptitiously accumulated secret stocks of work-in-process inventory to be used on days when low yield rates might have otherwise prevented them from achieving their daily performance objectives. Similarly, in the product development arena, engineers routinely failed to alert senior managers to significant design problems in ongoing projects so as to appear to conform to a newly instituted development process.

Once attention is shifted in this direction, localized change, which can be the principal engine of transformation, plays a central role in the feedback processes that generate inertia and convergence. Attempts to circumvent the measurement system and depart from standard procedures, once discovered by managers, provide compelling evidence that the stringent control systems were not only needed, but require further augmentation, creating a reinforcing cycle of increasing mistrust, stronger controls, and more elaborate departures from them. These feedbacks structure an environment all too familiar to students of management. Managers and workers are engaged in protracted and ongoing conflict. Supervisors get "beat up" for missing increasingly detailed and aggressive objectives and engineers get "shot" for missing product launch dates. The activities of the workforce are tightly controlled and, as a consequence, they have little interest in doing anything beyond following those specifications to the letter. The departures from standard procedure that do occur are largely targeted at manipulating or subverting measurement and evaluation systems. It should come as little surprise that organizations caught in this cycle have a limited ability to adapt, often require radical change, and frequently die.

Thus, when exploitation is focused on exercising opportunities latent in a new technology (either a "hard" technology as in Orlikowski 1996 or soft technology as in MCT), incremental changes can accumulate to produce transformation. But, when exploitation turns to appropriating peoples' time, energy, and attention, incremental changes made by both managers and those that work for them combine to create structural inertia and convergence on a stable orientation. Put differently, when exploiting technology, changes made by line staff and managers can be mutually reinforcing, thus creating the possibility of dramatic transformation. In contrast, when the focus turns to exploiting people, changes made by managers and the people that work for them are much more likely to offset each other, thus producing inertia and convergence.

So, in the end, the theory developed here reaches a conclusion similar to that found in episodic models–organizations will frequently converge on operating modes and find it difficult to incrementally adapt–but does so by reversing the underlying logic. Rigid routines, closely monitored standards, and tightly controlled activities, rather than preceding improvement to internal effectiveness, are, instead, often its unintended consequence. No single conception of performance will produce improvement forever, and when its ability to reveal improvements begins to decline, the organization reaches a critical juncture. Ideally, it would move to an alternative conception capable of revealing new opportunities and thus maintain a focus on exploiting opportunities present in the physical technology.

Unfortunately, managers are unlikely to recognize that a particular perspective has run its course. Instead for a host of reasons (discussed in the companion study), they are more likely to attribute a declining rate of improvement to the attitudes and dispositions of their employees, and, having done so, resort to increasingly standardized routines and institutionalized procedures as the path to further improvement. Of course, these routines will be consistent with the dominant performance model that has so successfully produced improvement in the past. Thus, for example, during the pre-MCT period, managers having successfully used the utilization-based model in the past, continued to increase both the frequency and granularity of their measurements until manufacturing facilities were required to collect and report utilization on a per machine, per day basis. The irony of this situation is that through standardizing and institutionalizing routines, the performance model becomes increasingly embedded in the physical and authority structure of the organization as a *consequence* of its declining utility in suggesting improvement opportunities. If this shift in emphasis occurs, from capitalizing on opportunities in the physical technology to appropriating peoples' time and energy, then ongoing transformation stops and convergence begins.

Conclusion

Change continues to be a major topic of discussion in organization studies research. Through exploring the micro-processes that lead to incremental change, this analysis contributes to this discourse by offering an alternative characterization of the processes through which such changes actually occur. The primary driver of improvement in this conception is not the use of specific tools, structured problem solving methods, or standard procedures. Rather, change ultimately stems from an alternative conception of system performance capturing elements absent from the approaches currently in use. Tools, methods and procedures are simply manifestations of the underlying approach. These ideas have implications both for future studies of process improvement and more general theories of organizational change.

Given its exploratory nature, this research has a number of serious limitations. Studying two initiatives in a single organization puts severe limits on the generalizability of the results.

Further, process-focused improvement efforts are only one type of efficiency improvement. Thus, the most important implication of this work is to suggest that understanding the true nature of efforts to improve organizational effectiveness at the level of specific changes is a potentially profitable area for future inquiry. If, however, the ideas developed here survive subsequent test and consideration, they may provide one path towards a more general conception of organizational change that incorporates the currently competing episodic and continuous views.

References

Brown, S. L. and K. M. Eisenhardt (1997). "The Art of Continuous Change: Linking Complexity Theory and Time-Paced Evolution in Relentlessly Shifting Organizations," *Administrative Science Quarterly*, 42: 1- 34.

Carroll, J., J. Sterman, and A. Markus (1997). "Playing the Maintenance Game: How Mental Models Drive Organization Decisions." R. Stern and J. Halpern (eds.) *Debating Rationality: Nonrational Elements of Organizational Decision Making*. Ithaca, NY, ILR Press.

Dean, J. W. and D. Bowen (1994). "Management Theory and Total Quality: Improving Research and Practice Through Theory Development, *Academy of Management Review*, 19(3): 392-418.

Deming, W. E. (1986). *Out of the Crisis.* Cambridge: MIT Center for Advanced Engineering Study, Cambridge, MA.

Easton, G. and S. Jarrell (1998). "The Effects of Total Quality Management on Corporate Performance: An Empirical Investigation," *Journal of Business*, 71, 2.

Eisenhardt, K.M. (1989). "Building Theories from Case Study Research," *Academy of Management Review*, Vol. 14, No. 4, 532-550.

Giddens, A. (1984). *The Constitution of Society: Outline of the Theory of Structure*. Berkeley, CA: University of California Press.

Goldratt, E. M. and J. Cox (1986). *The Goal: A Process of Ongoing Improvement* (revised edition). Croton on Hudson, NY: North River Press.

Hackman, J. and R. Wageman (1995). "Total Quality Management: Empirical, Conceptual, and Practical Issues," *Administrative Science Quarterly*, 40: 309-342.

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

Hannan, M. T., and J. Freeman (1984). "Structural inertia and organizational change," *American Sociological Review*, 49: 149-164.

Johnsson, Fredrik (1996). Sustainable Improvement Programs: Total Quality Excellence. Unpublished MS Thesis, Alfred P. Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Ma.

Kanter, R.M., T.D. Jick, and R. A. Stein (1992). *The Challenge of Organizational Change*. New York, Free Press.

Levinthal, D. and J. March (1981), "A model of adaptive organization search," *Journal of Economic Behavior and Organization*, 2, 307-333.

Macduffie, J.P. (1997). "The Road to Root Cause: Shop-Floor Problem Solving at Three Auto Assembly Plant," *Management Science*,43(4): 479-502.

March, J.G. (1991). "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2, 1:71-87.

Marx, K (1887), *Capital*, Ch. 10, section 2, available at http://www.marxists.org/archive/marx/works/index.htm

Miles, M. and A. M. Huberman (1984). *Qualitative data analysis*, Newbury Park, Ca: Sage.

Orlikowski, W.J. (1992). "The Duality of Technology: Rethinking the Concept of Technology in Organizations," *Organization Science*, Vol. 3, No. 3.

Orlikowski, W.J. (1996). "Improvising Organizational Transformation over Time: A Situated Change Perspective," *Information Systems Research*, 7(1), 63-92.

Orlikowski, W.J. and M.J. Tyre (1994). "Windows of Opportunity: Temporal Patterns of Technological Adaptation," *Organization Science*, 5, 1:98-118.

Romanelli, E. and M. L. Tushman (1994). "Organizational Transformation as Punctuated Equilibrium: An Empirical Test," *Academy of Management Journal*, 37, 5:1141-1166.

Sitkin, S., K. Sutcliffe, and R. Schroeder (1994). "Distinguishing Control from Learning in Total Quality Management: A Contingency Perspective," *Academy of Management Review*, Vol. 19, No. 3:537-564.

Sterman, J., N. Repenning, and F. Kofman (1997). "Unanticipated Side Effects of Successful Quality Programs: Exploring a Paradox of Organizational Improvement." *Management Science*, April, 503-521.

Tushman, M. L. and C. A. O'Reilly (1996). "Ambidextrous Organziations: Managing Evolutionary and Revolutionary Change," *California Management Review*, 38, 4, 8-30.

Tushman, M.L. and E. Romanelli (1985). "Organizational Evolution: A metamorphosis model of convergence and reorientation." In L.L. Cummings and B.M. Staw (eds), *Research in Organizational Behavior*,7:171-222.

Weick, K. E. and R. E. Quinn (1999). "Organizational Development and Change," *Annual Review of Psychology*, 50: 361-386.

Weick, K.E. (1998). "Improvisation as a Mindset for Organizational Analysis," *Organization Science*, 9, 5: 543-555.

Weick, K.E. (1995). *Sensemaking in Organizations*, Sage, Thousand Oaks, CA.

Westphal, J. D., R. Gulati, and S. M. Shortell (1997). "Customization or Conformity? An Institutional and Network Perspective on the Content and Consequences of TQM Adoption," *Administrative Science Quarterly*, 42: 366-395.

Wruck, K.H., and M.C. Jensen (1994). "Science, Specific Knowledge, and Total Quality Management," *Journal of Accounting and Economics*, 18, 247-287.

Wynne, B. (1988). “Unruly Technology: Practical Rules, Impractical Discourse and Public Understanding”, *Social Studies of Science*, 18, 147-167.

Zbaracki, M.J. (1998). "The Rhetoric and Reality of Total Quality Management," Administrative Science Quarterly, 43: 602-636.

Bone Yard

Organizing for Process Improvement

This interpretation of process-focused improvement efforts stands in stark contrast to that in the current literature. Most importantly, while the current literature typically focuses on the use of the tools and structured problem solving methods associated with a given approach (e.g. controls charts and statistical process control), in the framework developed here such tools and methods are relegated to a supporting role. The primary content of each of techniques lies in an alternative conception of the performance of the production system. The distinction is important for at least two reasons. First, it suggests, as was observed in the MCT effort, that dramatic improvement can often occur in advance of the use of sophisticated tools and methods. Further, as is often observed in other initiatives, it also suggests that the rote application of supporting tools and methods uninformed by a detailed understanding of the underlying model is unlikely to produce significant change. Thus, not surprisingly, wide spread, hands-on training that focuses on developing the intuition underlying the associated tool was central to both the success of MCT and other participatory initiatives (see Carroll, Marcus and Sterman for another example).

Second, it suggests that successful improvement, since it requires the use of local knowledge, also requires an organization structure that allows participants considerable latitude in the execution of their individual tasks. Within this second point lies what might be the fundamental paradox of successful process focused-improvement: increasing control over the production process (in terms of its ability to reliably produce output) often requires reducing control of those that work within that process. In such a world there is little role for highly granular measurement schemes focused on individual activities. Given the role of local knowledge, the ability of managers to directly evaluate individual contributions is limited. Instead, process focused measures built on a consistent underlying system-level model provide a more effective means of coordinating distributed action.