

---

# From Continuous Improvement to Continuous Innovation

---

ROBERT E. COLE, UNIVERSITY OF CALIFORNIA-BERKELEY

© 2001, ASQ

*In this paper Cole explores many concepts, including continuous improvement, continuous innovation, discontinuous innovation, incrementalism, exploitation, and exploration. He reviews the many benefits of continuous improvement, as it is defined in traditional quality programs. Cole discusses many organizational challenges arising from hypercompetition, working at an accelerated pace, uncertainty, the infusion of new technology, and the impact of software and information technology. Above all, he focuses on alternative ways to creatively build quality improvement through continuous innovation into the development process. His chosen vehicle is the probe-and-learn process and how it can lead to higher quality and shorter product development cycles.*

*Three commentators further explore Cole's theses. Finster examines how continuous innovation is related to the creative process. Melton discusses the role of learning in continuous improvement and continuous innovation. And Weston explains how both continuous improvement and continuous innovation are necessary for business survival.*

*Key words: creativity, discontinuous innovation, error, innovation, organizational learning, probe and learn, problem solving, product development, sociotechnical system design, systematic improvement*

This paper aims to introduce some new ways of thinking about continuous improvement, drawing upon the experiences of leading firms in the high-tech sector. These involve a focus on surfacing and learning from error. First, the relationship between continuous improvement and innovation will be explored. Then, specific strategies and tactics for applying continuous improvement ideas to industries undergoing rapid technological change in uncertain environments will be discussed. As more industries find applications for information technology, the greater the applicability of the strategy and tactics to be described. The objective is to make clear the linkage of these strategies and tactics with the quality field and in particular with the job duties of quality professionals.

The current version of the American continuous improvement movement grew out of the Japanese quality movement as it developed in the late 1960s and evolved through the 1980s. It was brought to the attention of Westerners in the early and mid-1980s by Western observations of the corporate practices of leading Japanese companies, as well as authors such as Masaaki Imai (Cole 1999; Imai 1986). It combined ideas, developed earlier by leading Western authors like Shewhart, Deming, and Juran with Japanese innovations including, above all, large-scale worker participation and training in continual improvement initiatives.

It can be argued that the concept and tools of continual improvement have seen little evolution since the 1980s. To be sure, six sigma made a big splash in the 1990s, but the contribution of six sigma has not been in the tools it uses or revolutionary thinking but rather in its marketing of the central ideas of continuous improvement and its integration of these ideas with business incentives and objectives (Maguire 1999). Six sigma, like traditional continuous improvement, is dedicated to the reduction of error.

## THE BENEFITS OF CONTINUOUS IMPROVEMENT

The significance of continuous improvement goes far beyond the quality movement. Ultimately it is about organizational renewal and efforts to prevent organizational ossification. Drawing and elaborating on the analysis of “small wins” offered by Karl Weick and Frances Westley (1996, 454–455), the significant benefits associated with continuous improvement can be detailed as follows:

- Continuous improvement typically mobilizes large numbers of employees on behalf of organizational improvement in contrast to large-scale innovation efforts that often involve only selected experts. The contribution of such broad mobilization of employees is potentially large.
- As a corollary of these broad-based efforts, small wins in large systems can occur in parallel as well as serially, resulting in the aggregate, in large numbers of change efforts leading in turn to a magnification of results.
- A series of small wins often precedes and follows large changes, first paving the way for these changes by providing momentum and basic learning, and second by eliminating the impediments to optimizing the new processes or products. In this sense, small wins make large-scale change possible. Leonard Lynn’s study of the introduction of the basic oxygen furnace technology in steel making nicely illustrates this process (Lynn 1982).
- When many seemingly revolutionary changes are scrutinized, they are found to be based on a series of small wins. Consider the revolutionary impact that the step-by-step reduction of machine setup times and die change times, pioneered at Toyota, had on changing the economics of small lot production in the auto industry (Robinson 1991, 85–86).
- By being anchored in current practices, small wins encourage learning that is rooted in daily work routines—exactly the kind of learning that is most likely to be transformed into effective practice—

that is, to be retained and institutionalized. This process is at the heart of the Brown and Duguid’s (1991) elaboration of how communities of practice lead to innovation. The potential for institutionalization is particularly large when the changes are implemented by the same people who proposed them.

- Small wins by disparate improvement groups are opportunistic and widely distributed. As such they represent uncorrelated probes in an evolutionary system. This is a particularly valuable asset. The heterogeneous nature of these probes means that they are more likely to uncover unanticipated properties of the environment and promote beneficial learning. Probing and learning provide a valuable model of learning, problem solving, and improvement that is broadly applicable.
- Small process wins are often based on tacit knowledge that is not easily noticed and imitated by competitors. This contrasts with large-scale changes and product innovation, which are more likely to be based on explicit and codified language and available to competitors through reverse engineering. It is easier to sustain competitive advantage when the knowledge possessed is tacit process knowledge (Teece 1998, 63–64).

Overall, many researchers and R&D managers believe that it is the patient accumulation of small improvements that accounts for the bulk of technological progress (for example, Tushman, Anderson, and O’Reilly 1997, 48).

## CLARIFYING TERMINOLOGY

Many researchers contrast continuous improvement with innovation, continuous improvement with discontinuous innovation, incremental innovation with discontinuous innovation, and exploitation with exploration (Imai 1986; Sutcliffe, Sitkin, and Browning 2000; Tushman, Anderson, and O’Reilly 1997; March 1994). For now, the focus will be on the common distinction between continuous improvement versus innovation.

In the 1980s, a decade of seeming Japanese emergent supremacy, the continuous improvement approach was often held up as superior to innovation (Imai 1986; Florida and Kenney 1990; Gomory 1989). By the late 1990s and the beginning of the new millennium, the resurgence of American industry in high tech combined with the stagnation of the Japanese economy put renewed emphasis on the benefits of innovation (Brown and Eisenhardt 1998). A popular revisionist view was not content to argue for the weakness of the continuous improvement approach relative to breakthrough innovation. Rather, it argued that continuous improvement, slow and plodding, was downright un-American, inconsistent as it was with the American cultural emphasis on improvisation and innovation (Hammond and Morrison 1996).

Masaaki Imai argued that continuous improvement worked best in a slow-growth economy, and innovation was more suited to a fast-growth economy (Imai 1986, 24). Yet, Japan was growing quite rapidly in the 1960s and 1970s, benefiting greatly at the same time from its continuous improvement activities. An alternative but more convincing explanation is that conventional continuous improvement works best when firms are playing catch-up; therefore, they know pretty much the direction they need to go by observing those ahead of them. Thus, continuous improvement fit large Japanese manufacturing firms, which, for most of the post-World War II period, were playing catch-up. When firms are operating on the frontiers of technological knowledge, however, more discontinuous innovation is required. During the late 1980s, as the Japanese moved to the frontiers in many industries, they found it difficult to shift gears to more discontinuous innovative change, and they have fallen behind in many key areas.

This discussion begs the questions of just how useful is the common categorization of continuous improvement versus innovation. The common assumption is that continuous improvement is small scale and that innovation is discontinuous and large scale. Yet, there is no logical reason to associate the term *innovation* with large-scale discontinuous change. Consistent with a dictionary definition, innovation is best associated with creative solutions, and

these can occur at a small as well as a large scale, and can be more, or less, discontinuous. Put more bluntly, there is plenty of innovation that occurs in the course of continuous improvement.

Typically, those who juxtapose continuous improvement versus innovation see them as *trade-offs* and/or as temporally sequenced. Sutcliffe, Sitkin, and Browning (2000, 316) summarize discussion of this perceived dilemma. Reflecting on the difficulty of combining exploitation and exploration, a closely related distinction, March notes that the difficulty of balancing the two is complicated by the fact that returns from the two options vary not only with respect to their present expected values but also with respect to their variability, their timing, and their distribution within and beyond the organization. The net result is that organizations have great difficulty in even understanding and specifying the appropriate trade-offs, much less defining and creating an appropriate balance between them (March 1994, 238–240).

What if, instead, continuous improvement and discontinuous innovation could be seen as complementary? That sounds good in principle but it suffers from the fact that some firms and industries are far better at one than the other. Some industry conditions give managers much stronger incentives, resources, and constraints to do the one rather than the other. As a result, their capabilities may be sharply skewed to one or the other. Yet, it is also clear that in many situations, those firms that can find a way to do both would be best off.

Thus, a number of scholars have tried to find some way to combine the two perspectives. Tushman, Anderson, and O'Reilly (1997, 19) call for an ambidextrous organization that combines efficiency and innovation, tactical and strategic, and large and small. Sutcliffe, Simkin, and Browning (2000) have made an effort to clarify what a complementary balanced approach might look like. They argue for a synergistic approach in which greater control (which they associate with continuous improvement) and exploration (which they associate with disjunctive change) are mutually reinforcing in that "each process facilitates and contributes to the effectiveness of the other" (Sutcliffe, Sitkin, and Browning 2000, 326). Their dis-

cussion of how this is to be effected is quite abstract focusing on the mutually reinforcing nature of reliability (associated with continuous improvement in their view) and resilience (associated with learning). They do provide two examples: the first one details the way in which failure mode effects analysis (FMEA)—presumably a control tool used in process improvement—was creatively applied to product development in one of the firms they studied. The second example involves the application of reliability and control methods to an R&D lab.

---

*Innovation is best associated with creative solutions, and these can occur at a small as well as a large scale, and can be more, or less, discontinuous. Put more bluntly, there is plenty of innovation that occurs in the course of continuous improvement.*

---

Seemingly contradictory to their understanding, FMEA analysis was developed in the early 1970s *specifically for product development* by Ford Motor Company engineers and fairly soon after, adapted for use in process improvement at Ford and elsewhere. Thus, viewed from this broader perspective, it was not all that creative for the firm in question to apply FMEA to product development. For an early use of FMEA for process improvement, see Ishiyama (1977).

The initial challenge is to see innovation as part of the continuous improvement process and then to see whether discontinuous innovation can be infused with a continuous improvement approach. It has already been argued that there can be a great deal of innovation built into continuous improvement efforts, or in the language of Brown and Duguid (1991, 53), incremental innovations grounded in work practices (communities of practice) occur throughout an innovative organization. Many creative solutions are associated with continuous improvement. On the other side, for large-scale discontinuous innovation to be successful, there has to be a great deal of continuous improvement surrounding it—before, during, and after.

The logic of the argument thus far, roughly consistent with Tushman, Anderson, and O'Reilly (1997),

suggests that instead of distinguishing between continuous improvement and innovation it might better to distinguish between continuous innovation and discontinuous innovation—with much of continuous innovation involving small-scale and local innovation. This is the terminology to be used hereafter. This usage encourages an understanding that, in practice, there is a continuum between continuous innovation and discontinuous innovation even if cases are coded only according to these two binary categories.

## THE NEW CHALLENGES

In this time of hypercompetition, organizations witness an accelerating pace of technological change—an acceleration of “clockspeed” in one industry after another (Fine 1998). There is an infusion of new technology even in traditional industries like furniture making and retail sales. Major vehicles for that infusion are the role of software and information technology (IT) in determining product functionality and facilitating logistics.

The speed at which firms develop and roll out new products has become an increasingly critical competitive issue. Consider that the product life cycles in the PC industry were approximately one year in the middle 1980s; by 1997, these were reduced to approximately three months (Curry and Kenney 1999, 8–9). Shorter product cycles mean that firms have less time to recoup their investments and be first to market with the right product, and quality confers major competitive advantage. Indeed, in the new economy, some go as far as to argue that in this world of increasing returns, those products and firms that get ahead, advance further over time as a result of a series of positive feedback loops. This is a world of winner-takes-all markets. This exaggerated view ignores the dynamism of emergent markets and technology. Nevertheless, there is clear evidence that in rapidly changing high-tech markets, being late to market significantly reduces profits (Vesey 1991). Every manager nowadays seeks to compress development, production, and delivery times and integrate these operations into as seamless a process as possible.

The common element in all this is speed. Is it, however, compatible with continuous improvement—par-

ticularly in an environment of great uncertainty? The tools of continuous improvement were developed in fairly slow-moving industries like the automotive industry. The problem-solving protocols have stressed that in order to solve a problem, one must systematically go through a set of elaborate steps. One must first plan and decide on what the right problem is, clarify the reasons for selecting that problem, assess the present situation, collect all the relevant data, sort it, analyze it, decide on the cause of the problem, develop and implement a corrective measure, and evaluate the results. If the evaluation is positive, then one must standardize and act to prevent regression. This is impressively systematic but also very time consuming! It is not an approach that works well in a rapidly changing environment.

Indeed, even in relatively slow-moving industries, some companies have tried to come to terms with the need for speed in their problem-solving efforts by developing new tools. Ford Motor Company developed a streamlined version of its problem-solving process for a selected subset of problems. Ford calls the approach the Rapids Program. Developed in 1993, it is Ford's adaptation of GE's work-out process. Short workshops are developed around problem areas that are seen as immediate improvement activities or what Ford calls "quick hitters." Selected teams then meet to identify their concerns, generate ideas for action, and recommend solutions; and they do it all in the space of a day or even hours. A planning team then develops the implementation plan and tracks and revises actions as needed. Other auto companies, like Volkswagen, have their own version of the Rapids Program.

What about the really fast-moving industries where managers are under incredible pressure to accelerate the pace of development, production, and delivery and to integrate them in a seamless process? The literature suggests that in competitive technology-intensive global markets, advantage is built and renewed through more discontinuous forms of innovation—through the creation of new families of products and businesses. One can contrast this with continuous incremental product line extensions and improvements that are essential for maintaining leadership. These maintenance activities are significant, but come into play only

after leadership has first been established through discontinuous forms of innovation (Lynn, Morone, and Paulson 1996, 9).

The question to be addressed, however, is does continuous innovation have a contribution to make toward the promotion of discontinuous innovation? Examining recent scholarship dealing with speed and product development, speed is found to be associated with an emphasis on concurrent engineering involving the use of overlapping product development stages and parallel processing, design without delay (shorten lead times by taking out all unnecessary delays), and design for manufacturability (Flynn et al. 1999, 247; Brown and Eisenhardt 1995; Clark and Fujimoto 1991). The focus of these efforts is on streamlining and simplification.

---

*The initial challenge is to see innovation as part of the continuous improvement process and then to see whether discontinuous innovation can be infused with a continuous improvement approach. . . . On the other side, for large-scale discontinuous innovation to be successful, there has to be a great deal of continuous improvement surrounding it—before, during, and after.*

---

The field of quality has long stressed the importance of applying quality principles to the new product development process. Thus, Armand Feigenbaum outlines first the 16 sequenced steps in new product development and then shows how quality principles mesh into this sequence. The four principal quality routines that he sees integrated into these product development process steps are as follows:

1. Establishment of the quality requirements of the product: This involves the creation of customer-oriented specifications and standards.
2. Design of a product that meets these requirements: This involves the establishment of the detailed drawings for the product and preparation of the related engineering instructions and all associated testing and simulations.
3. Planning to assure maintenance of the required quality: This involves the formal activation of the

details of the quality program that covers control of purchased material, maintenance of quality during processing, and production and assurance of quality during field installation and product servicing.

4. Preproduction review of the new design and its manufacturing facilities and formal release for active production: This involves the planned formal evaluation of the designed product at the various stages of the complete design process to assure its capability of meeting its warranties and guarantees in actual use (Feigenbaum 1991, 626–628).

Reviewing this model, its linear character is seen first, followed by its focus on systematizing and rationalizing the product development process. The quality field has traditionally been focused on planning, simplification, systematization, and streamlining as the basis for insuring that the product development process will yield high-quality products (Hargadon and Eisenhardt 2000).

---

*Examining recent scholarship dealing with speed and product development, speed is found to be associated with an emphasis on concurrent engineering involving the use of overlapping product development stages and parallel processing, design without delay (shorten lead times by taking out all unnecessary delays), and design for manufacturability. . . . The focus of these efforts is on streamlining and simplification.*

---

In Joseph Juran's view, structured processes such as those outlined by Feigenbaum are not enough to insure new, high-quality products. Often firms still need to increase speed, improve the competitiveness of their products, and deal with chronic wastes that are created. Juran sees these problems as resulting mostly from weaknesses in the quality planning processes and as ones requiring continuous innovation. In particular, he focuses on the importance of eliminating chronic waste and increasing the annual rate of quality improvement faster than competitors (Juran and Godfrey 1999, 5; 3–5). Thus, the key to achieving high

levels of quality in the product development process is to eliminate chronic waste (for example, rework) through better planning, simplification, systematization, and streamlining. Generally speaking, then, the key is still rationalization of the development process.

In the light of this discussion, the two examples of synergistic alliance of greater reliability and control (continuous improvement) with greater learning and exploration (innovation), as proposed by Sutcliffe, Sitkin, and Browning, can be reexamined. The first one details the way in which failure mode effects analysis (FMEA)—presumably a control tool—was creatively applied to product development. The second one involves the application of reliability and control methods to an R&D lab. Both examples involve integrating traditional tools of the quality movement into the R&D and product development process. As such, they are quite consistent with the conventional approaches to quality improvement in the product development process that stress rationalization of the development process.

## PROBE AND LEARN AS AN ALTERNATIVE APPROACH

There are, however, alternative ways of thinking about how to creatively build quality improvement through continuous innovation into the development process. It is to this end that the rest of the analysis is devoted. The first step in such thinking is to understand that product development in turbulent sectors, like high tech, is an emergent process in which the premium is on learning and rapid incorporation of that learning into subsequent as well as previous development processes. This severely limits the contribution of conventional planning, so much the hallmark of the traditional approach to incorporating quality into the product development process. It also paradoxically encourages the successive generation of error, early and often, as part of the learning process.

Implicit in this description, is that product development in a turbulent environment requires a nonlinear process, with both backward and forward movement occurring as the development team often revisits past

decisions based on new information and changing circumstances (Hargadon and Eisenhardt 2000). These conceptions suggest that the task of infusing continuous innovation into the processes of discontinuous change goes well beyond the traditional approach of quality experts, which is to figure out how to apply conventional quality improvement tools to rationalize and streamline the discontinuous change process. It requires understanding conceptually what is meant by continuous innovation and developing the tools to implement those new understandings.

Conventional marketing is particularly useless in identifying desirable product features and uses of discontinuous products and services. This is because of the enormous uncertainty surrounding the development of such new products and services. Typically the industry is evolving, the market is ill defined, and the infrastructure for delivering the still-developing technology to the yet-undetermined market is nonexistent. There are timing uncertainties: time required to develop the new technology; time required for the market to emerge; and time required for complementary technologies to emerge; and these uncertainties interact (Lynn, Morone, and Paulson 1996, 10).

Under these rapidly changing circumstances with high levels of uncertainties and complex interaction effects, problems and errors are inevitable (Morgan 1997, 94). This is in stark contrast to conventional quality thinking that stresses seeking the holy grail of prevention. Specifically, the standard thinking that has been pounded into quality professionals is that an organization should aim to prevent errors and defects upstream through designing quality in. Failure to do so will lead to an inevitable compounding of error that results in heavy reliance on repairing and reworking defective products downstream (see Flynn et al. 1999, 250). Much of the success of post-World War II Japanese quality movement can be read as arising from the shift from downstream detection to upstream prevention.

Yet, one of the challenges of product development under conditions of rapid change, high uncertainty, and complex interaction effect, is precisely to surface error early and often! It is not only inevitable, but desirable! Prevention, of course, is still a goal but it occurs

either through concerted efforts to continually uncover and then remove error or through efforts to eliminate as early as possible those specific errors that do not contribute to learning. In this sense, the simple view that the quality movement has historically evolved from detection to a focus on prevention (Garvin 1988, 19) is, in turbulent uncertain and interactive environments, incorrect. Rather what is seen is a much more complex equation in which the generation and detection of error plays a renewed *and desired* role. Of course, quality leaders have always advocated learning from error. For example, as part of traditional product quality control, one often carries out accelerated life testing under simulated field conditions: Find the location (where) and timing (when) at which an error (failure) in the tested products can be generated. The purpose is to use that information to control, reduce, or prevent subsequent error.

---

*The first step in such thinking is to understand that product development in turbulent sectors, like high tech, is an emergent process in which the premium is on learning and rapid incorporation of that learning into subsequent as well as previous development processes. This severely limits the contribution of conventional planning, . . . It also paradoxically encourages the successive generation of error, early and often, as part of the learning process.*

---

The discussion here, however, is about something quantitatively and qualitatively different. It is about intentionally and successively generating errors throughout the product development process and in interaction with downstream customers from whom lessons can be learned. This implies a distinction between desirable error from which lessons can be learned (which should be encouraged) and unnecessary error (which does not lead to learning and should be prevented). Error is particularly desirable when it can cause employees to question the underlying values and policies of the organization, leading to new, more efficient, and/or effective behavior (Argyris 1992).

The focus is on meeting customer needs through discovery that enables heightened performance and

new features, not increased reliability through control. This is especially the case with getting new technologies to customers as quickly as possible as firms seek to create new markets and carve out market leadership positions with potentially long-term positive consequences. As Geoffrey Moore (2000, 11–12; 152) says, companies need to “go ugly early.” He emphasizes that getting bad reviews for product features and quality performance is better than getting no reviews at all. Being first to market with new technology is a time in which customer intimacy and operational excellence are not appropriate targets. Rather one should be learning from one’s mistakes through having the product in the field and then building that feedback into the next version of the product or service. One should be careful with this argument. If a firm has a brand that it wants to protect, for example, it must weigh the impact of early failures on reputation relative to the benefits of getting this imperfect product to market early. Generally, it is important to explain to early adopters the risks that they may be taking in their use of an early version of a product.

Compare this situation to evolving practices in a slower-moving industry like automobiles. It was not uncommon in the mid-1970s for dealers to tell new car owners to collect up all the problems they had with the car in the first month and then bring it in for the dealer to fix. Superficially this looks like the same approach—using the customer in the field to test the new product. There is, however, an important difference. The feedback that the dealer acted on to make repairs was seldom forwarded to the manufacturer with the expectation that the manufacturer would redesign away the problems found by users. Even today, U.S. automobile manufacturers are notoriously unwilling to make design fixes for automobile models already in production (MacDuffie 1999). Moreover, these dealer requests were not made just for new models but for existing models—many of which had the same problem in their fifth year as they had in their first year. These dealer practices disappeared in the 1980s under the pressure of Japanese competition that produced high-quality vehicles with a minimum of start-up problems for the new owner. That is, the Japanese were more able to achieve operational excellence from the

start. As mentioned, however, the automobile industry has been a relatively slow-moving industry technologically speaking. In industries where there are emergent interacting technologies with uncertain trajectories and in an uncertain environment, it is unwise from the start to make operational excellence one’s top priority. Rather, customer feedback on quality, performance, features, and subsequent redesign is more desirable.

How does a firm manage in this environment? It does so by focusing on the front end of a redefined development process to demonstrate the relevance of continuous innovation even in the case of discontinuous product development. By examining how companies that have developed successful products operate in this space, a probe-and-learn process is revealed (Brown and Eisenhardt 1998; Morgan 1997, 273; Lynn, Morone, and Paulson 1996).

Essentially companies develop their products by probing potential markets with early versions of the products, learning from their mistakes, modifying their product, and probing again. In effect, they run a series of market experiments—introducing prototypes into a variety of market segments. When using this approach, the initial product is not the culmination of the development process as it is in traditional organizations. Rather, the initial product is just the first step in an improvement process! This first step in the development process is, in and of itself, less important than the learning and the subsequent better-informed steps that follow (Lynn, Morone, and Paulson 1996).

Probing markets with immature versions of the product only makes sense if it serves as a vehicle for learning. It can be used to learn about the technology, and whether and how it can be scaled up. It can be used to learn about the market and which applications and market segments are most receptive to particular configurations of product features. It can be used to learn about the influence of exogenous factors like government regulations and what needs to be done to satisfy them. Probing and learning is, above all, an experimental, iterative process. The firm enters an initial market with an early version of the product, learns from the experience, modifies the product, and adjusts the marketing approach based on what was learned. Then it tries again and again, as necessary. In summa-

ry, development of a discontinuous innovation becomes a process of successive approximation, probing, and learning again, each time trying to take a step closer to a winning combination of product and market (Lynn, Morone, and Paulson 1996).

Lynn, Morone, and Paulson (1996) go on to document a series of four prototype products emanating from Motorola on the way to developing its first commercial portable cellular phone. It took over eight years before it got to the point of a product that was designed to be manufactured in mass quantities. Each generation represented a step forward in performance, based on lessons learned in the previous step. Each step created greater understanding of the market and acceptance of a product that had been initially met with great skepticism both inside and outside the company. It all came together in 1984 when Motorola came to the market with several products aimed at specific markets. For a discussion of comparable probe-and-learn strategies from Charles Schwab and Sun Microsystems, see Brown and Eisenhardt (1998, 147–148).

---

*The focus is on meeting customer needs through discovery that enables heightened performance and new features, not increased reliability through control. This is especially the case with getting new technologies to customers as quickly as possible.*

---

In conventional product development there is a single launch into which all the accumulated knowledge is put, and product designers can only hope that they will be successful: All their eggs are in one basket. If they succeed, the payoff may be quite large but they are making a very big bet and losing will be very costly. Because the process takes so long, they run the risk that the market and the technology may have changed from the time they made their initial judgments. Thus, the probability of failure and the cost of failure increase (Stalk and Hout 1990).

With probe and learn, there is no single launch of the new product as in conventional innovation but rather a series of launches sometimes over many years. Each new launch leads to a modification of the target. In so doing, the firm reduces uncertainty and thereby

reduces the financial risks for the next launch because of what it has learned at each stage. This is all about continuous innovation (Lynn, Morone, and Paulson 1996).

National Semiconductor engages in a variety of activities that emulate this probe-and-learn process in its interaction with key customers as it develops chip applications for cellular phones and automotive products. These activities range from encouraging customers to jointly develop products with the company, to passing out free product samples to customers, and then revising the product design according to their feedback. National Semiconductor has software in place to do Web meetings with customers to facilitate these exchanges. It gives customers access to its database so they can monitor National Semiconductor's development process and provide rapid feedback. All these activities are designed to create an iterative probe-and-learn process for product development, and, as such, involve continuous improvement. National Semiconductor reports that these kinds of processes work less well in cost-sensitive markets like PCs where customers don't have the resources to perform joint development. (The author is indebted to Nien-Tsu Chen, director of quality, National Semiconductor Corp., for this account.)

At first glance, however, this version of continuous innovation doesn't seem to be the kind of continuous improvement quality specialists have been accustomed to thinking about. First, it occurs in the early stages of product development, while most of the applications with which quality experts are familiar are operational improvements in manufacturing environments. Second, it facilitates discontinuous technological innovation, while most applications with which quality experts are familiar deal with continuous or incremental innovation. Third, it uses the customer as the driving force for the learning process, while most of the applications with which quality experts are familiar involve internally generated data. And fourth, it doesn't use the typical continuous improvement tools that evolved out of the quality movement.

On closer examination, however, it can be seen that the probe-and-learn process does lie at the heart of continuous innovation. In fact, it captures the essence

of continuous improvement. Probe and learn is based on a series of continuous small gradual steps. If well done, it is experimental in the best sense of embodying fact-based management. Probe and learn is focused on process not results, like all continuous improvement activities. The process of successively honing in on the right product through a series of iterative steps that take firms closer to a successful commercial product is very consistent with the spirit of a continuous improvement. Probe and learn—more accurately put probe, test, evaluate, and learn (and refine)—is essentially an accelerated plan-do-check-act (PDCA) cycle, just as is Ford's Rapids Program. Probe and learn can be seen as a new form of PDCA suitable for dynamic environments. Unlike conventional PDCA, the probe-and-learn process underweights plan, and overweights do. It stresses the rapid-fire learning that comes from evaluation of the do phase and the iterative nature of the process. Finally, probe and learn is about organizational renewal and thus totally consistent with the ultimate objective of continuous improvement. Yet, it is also associated with quick learning and the acceleration of the product development process, a prime requirement in this era where firms operate on Internet time. In summary, the probe-and-learn process embodies the principles of continuous improvement.

## PROTOTYPING

Already noted was the practice of probing potential markets with prototypes but, of course, prototypes have a much broader role in the product development process. Prototypes are analytical or physical models that are used to test or verify aspects of the product design at different stages of the development process. They are useful in early design phases to assess the size and feel of a product; at later stages, comprehensive physical prototypes can reveal interferences among components and whether everything works when connected (Rao et al. 1996, 516–517). Through the use of successively comprehensive prototypes, the same accelerated PDCA cycle, which characterizes the description of the probe-and-learn process, is seen.

Prototypes can be used to produce a model of the whole product or some small component. While pro-

ducing virtual objects via computer-aided design (CAD) has become standard practice, the production of multiple physical prototypes ranging from simple cardboard-and-glue models to sophisticated stereolithography (SLA)-produced models is recognized to add considerable value and speed to the development process. The central contributions of prototyping are its acceleration of learning and coordination throughout the development process, across diverse functional groups or geographically dispersed groups within and outside the firm. Prototyping focuses attention on problem areas needing improvement, clarifies sources of different views, and confirms common areas of understanding and agreement. Prototyping facilitates communication across cross-functional groups (inside and outside the firm) and contributes to the development of a common language (Leonard-Barton 1991). Thus, prototyping can be used for streamlining the flow of a total CAD/CAM/molding/assembly operation among multiple-production partners. It accomplishes this by seeing to it that the original equipment manufacturer and all tiers of suppliers are tuned in to a common understanding of what has been accomplished and what still needs to be done.

Prototyping directly improves the quality of the product through early identification of error, and multiple iterations continually test the designers' assumptions about the product, leading to improved redesigns (Hargadon and Eisenhardt 2000; Leonard-Barton 1995). The very incompleteness of early prototypes guarantees the generation of error. Because of their ability to facilitate early detection of error and thereby reduce engineering changes, prototypes can reduce design iterations. The accelerating development of rapid prototyping technologies, along with the emergence of computer-aided design and engineering tools, have increased the speed and lowered the cost at which multiple prototyping iterations can occur, thereby speeding up the development process itself. At the same time, the cost of successive design iterations is reduced (Wright 2001, 130–170; Thomke and Reinertsen 1998, 20–21). The learning process associated with rapid prototyping grows out of a process of prototype, test, evaluate, and refine the product; it embodies the probe-and-learn pattern that has been documented.

In the light of these observations, it is quite remarkable how little attention has been paid to the benefits of prototyping in the quality literature. In Juran's comprehensive *Quality Handbook* (5th edition), there is no entry under prototyping. In Feigenbaum's extensive treatment of new product development in his renowned book, *Total Quality Control* (3rd edition), there is only a brief 11-line paragraph, over half of which is taken up with warnings of how performance of handmade prototypes may differ from those made under actual production conditions (Juran and Godfrey 1999; Feigenbaum 1991, 243). This suggests that the probe-and-learn functions of rapid prototyping are underestimated in the quality community. In part, this may be because there were few prototyping tools available in the past. It may also be that this underestimation occurs because prototyping is less focused on the simplification, systematization, and streamlining functions of traditional quality improvement. Instead, it is focused on coordination, learning, and exploration.

### Beta Testing

Probe and learn is also being implemented at the middle and latter stages of the product development process. Notable is the growing use of beta testing. There has been an explosive growth of beta testing over the last decade in the United States. It is a practice that began in the computer industry, but by 1994, it was estimated that 50 percent of *Fortune* 1000 companies had participated in beta testing and 20 percent were said to use it regularly (Daly 1994, 37). One can only presume that the number is still higher today. One of the most dramatic examples of the use of beta testing was by Microsoft. Its Windows 2000 release was said to have 500,000 prerelease customers participating in the beta testing (Wright 2001, 414).

Originally beta testing referred to the exercise and evaluation of a complete product working in the operating system environment; it would typically precede announcement and release. In recent years, however, the concept has been expanded to include customer evaluation and input prior to formal release of the product (Paul 1999, 17, 19). In that sense, it is about exposing users to incomplete products full of errors.

Users are motivated by getting the opportunity to try out and use early versions of the product with the understanding that they will report back to the manufacturer on their experiences.

---

*... the probe-and-learn process does lie at the heart of continuous innovation. In fact, it captures the essence of continuous improvement. Probe and learn is based on a series of continuous small gradual steps. If well done, it is experimental in the best sense of embodying fact-based management. Probe and learn is focused on process not results, like all continuous improvement activities.*

---

Customers often want to participate because of the potential competitive advantage that comes from being first to install a working model and getting an early look at new technology. Product developers realize that they can get very useful feedback from customers about possible performance problems and about the functionality of the product (Does it work as intended in a diverse user environment? Does it have the most desirable product features?). The knowledge thus gathered can then be incorporated, as deemed relevant, into subsequent iterations of the product. The firm producing the product may also use it to promote its ties with those valued customers that get a first look at the new technology and to promote the sales of the product by using early positive experiences to promote sales to subsequent customers (Dolan and Matthews 1993, 318–330). Some companies may release successive betas as a mode of successively approximating what the customer really wants.

Companies with products that have low manufacturing costs have refined this practice to a fine art. Software companies, in particular, have used the Internet to rapidly collaborate with application developers over successive beta iterations. A case in point: When Netscape developed Navigator 3.0 in just seven months, it went through six beta iterations, learning each time from the feedback of developers over the Internet, and incorporating what was learned into the next modified product version.

If well designed with careful selection of beta testers

to reflect the user community, structured questions for the initial users, and with an action plan to quickly address issues raised by the testers, beta testing represents an opportunity for rapid learning about new products. It is an articulation of the probe-and-learn process used in the middle to latter part of the product development process. However, selecting the “wrong” customers and/or applications for beta testing can lead to false inferences and erroneous decisions. Similarly, there can be significant risks for the firms that agree to provide sites for beta testing, and they need to carefully consider whether and under what conditions are in their interests. Often contractual language is vague, and users are not aware of the risks with producers failing to give beta testers appropriate warnings. Users continue to detest error; they can accept it only if they really understand the level of risk and accept being part of the development process. The type of customers becomes a factor as well; all things being equal, technical enthusiasts are more likely to be forgiving of error than mainstream customers.

---

*...prototyping is less focused on the simplification, systematization, and streamlining functions of traditional quality improvement. Instead, it is focused on coordination, learning, and exploration.*

---

At the extreme end of the continuum is open source software where continuous innovation takes place throughout the development process. It is an extended application of probe and learn. The Linux kernel development process, for example, is one of continuous improvement with none of the releases ever being final. The kernel of the operating system schedules the tasks, which include the execution of end-user applications, such as Web browsers, word processors, and database management systems, by allocating the computer's system resources to the programs in execution. In particular, the kernel controls the hardware and manages the flow of information and communication among various components of the computer.

In one of the e-mail discussions on the Linux kernel mailing list in August 1999, Linus Torvalds, the original creator of Linux, elaborates on the merit of let-

ting “people see what's going on” and argues forcefully for the importance of frequent submission of small patches with incremental change. He wrote,

The point of open development is that people see what's going on [underlining in original]. You don't get that if people see just the end result after a year. You want to have random people just see small updates—because they will often catch silly mistakes. Now, with huge megapatches, people just go numb... With the regular “let's release this as it is developed” support, there have been Web sites with commented patches, people who read the incremental stuff and comment on stupid things [that] I and others do (LINUXCARE 1999, 11).

Peer review is a critical feature of open source development. One way to encourage peer review is to increase product release frequency and shorten the product cycle. The sooner the feedback is incorporated, the more developers are encouraged to contribute. As such, quick responses keep developers engaged. Compared to commercial software, Linux is a continuously evolving product of a higher update frequency. Most commercial software companies release their products and/or follow-up upgrades only every few years, and the releases are often delayed. Although commercial firms use “daily build” to update progress, the released information is only circulated in the firm internally.

Since the first release of Linux, there has been on average one new version of the system released every week. On average, the product cycle is on the order of weeks. The development version is where developers can experiment with advanced technology and try new ideas. When developers are active, there are as many as three new development kernel releases a day, a much shorter cycle than that of the stable version. New features are tested in the development version first and then become included in the stable version. In 1996 alone, there were 30 official releases of the stable version while there were 80 releases of the development version (Lee and Cole 2000). In short, this process embodies the essence of continuous innovation.

## Implementing a Probe-and-learn Strategy

How does one implement a probe-and-learn strategy? Brown and Eisenhardt (1998, 155–156) provide a set of recommendations for how to create a wide variety of low-cost probes. Their recommendations are listed and elaborated as follows:

- Vary the time frames for the variety of low-cost probes being pursued. This involves creating both short- and long-term probes and encouraging them to emerge from different parts of the business. Insofar as they come from different parts of the organization, they are consistent with traditional continuous improvement through their broad-scale involvement of personnel in the improvement effort.
- Choose some risky probes even if they have a high probability of failure, especially small failures. These are opportunities for learning.
- Select some probes that require implementation and measure their results.
- Solicit concrete feedback because it is a very effective mode of learning.
- Use more probes when the marketplace is highly volatile.
- If large probes are unavoidable, then seek to break them into a series of small options that serve as opportunities to learn and also provide a chance to cut losses should overall failure become evident over time.
- Place more probes in areas that represent the most likely future, whether it be market segment or emergent technology.
- Select some unrelated probes in unfamiliar areas. Random probes are more likely to reveal the unexpected and the unanticipated. (See Weick and Westley (1996) discussion earlier in this paper.)
- When feasible, build on successful probes to create a knowledge base for emergent strategies.
- Know when to quit a series of probes in one area when diminishing returns set in and commit to others.

## CONCLUSION

If continuous improvement is conventionally considered, then it is likely best in slow-moving industries and in industries where firms are playing catch-up to a future that is laid out before them. These are industries where exploitation rather than exploration is required for success. If one's understanding of continuous improvement is widened to think in terms of continuous innovation, then there is a place for it in the process of exploration and discontinuous innovation. This has been the thrust of the previous analysis. Whenever it occurs, continuous innovation of the kind that has been described is not a natural process that automatically occurs in organizations (Tyre and Orlikowski 1994). It requires constant, active management and engagement with workers in an effort to initiate and sustain momentum. Probe and learn, insofar as it takes places in different parts of the organization at different times, through multiple initiatives, has the potential to serve as a sustained energizing force.

---

*Beta testing represents an opportunity for rapid learning about new products. It is an articulation of the probe-and-learn process used in the middle to latter part of the product development process.*

---

Probe and learn, applied to the product development process, captures the essence of continuous innovation. It is a process well suited to fostering discontinuity and innovation. It is an experimental iterative process that operates to successively solve problems in markets characterized by turbulence, uncertainty, and complex interactions. Probe and learn teaches that the generation of error is part of a productive learning process and should not always be avoided or suppressed. How firms learn to manage error in the new economy provides an important indicator of their success. This is a special challenge for the quality discipline—a discipline that has grown up viewing deviance and error as the enemy. For if quality professionals don't learn how to manage error in a dual fashion, other disciplines will take up the slack.

In discussing the four manifestations of probe and learn—distributing early versions of products to selected markets, prototyping, beta testing, and open source development—the intent was not to suggest that this exhausts the utility of probe and learn in the new product development process. To the contrary, the examples were only meant to show the broad potential for applying probe and learn in the product development process. The challenge for quality practitioners and scholars is to develop a set of tools that will improve the deployment and optimization of probe-and-learn strategies. It is no longer enough to simply look for areas within the product development process to which traditional quality improvement tools can be applied to rationalize and streamline the process. Finally, if there is a place for continuous innovation in discontinuous product development, surely there is a place for it throughout the production chain.

---

### ACKNOWLEDGMENT

The first draft of this article was presented as a keynote speech to the 3rd International (Euro) CI Net Conference at Ålborg University in Ålborg Denmark September 18, 2000. I am indebted to Prof. Frank Gertsen for his support and comments on the original draft. I would also like to thank Tito Conti and Eva Chen for their comments and suggestions.

---

### REFERENCES

- Argyris, Chris. 1992. *On organizational learning*. Cambridge, England: Oxford University Press.
- Brown, John, and Paul Duguid. 1991. Organizational learning and communities of practice. *Organization Science* 2, no. 1:40–57.
- Brown, Shona, and Kathleen Eisenhardt. 1995. Product development: Past research, present findings, and future directions. *Academy of Management Review* 20, no. 2:343–378.
- . 1998. *Competing on the edge*. Boston: Harvard Business School Press.
- Clark, Kim, and Takahiro Fujimoto. 1991. *Product development performance: Strategy, organization, and management in the world auto industry*. Cambridge, Mass.: Harvard Business School Press.
- Cole, Robert E. 1999. *Managing quality fads*. New York: Oxford University Press.
- Curry, James, and Martin Kenney. 1999. Beating the clock: Corporate responses to rapid change in the PC industry. *California Management Review* 42, no. 1:8–36.
- Daly, John. 1994. For beta or for worse. *Forbes ASAP*, 5 December, 36–40.
- Dolan, Robert, and John Matthews. 1993. Maximizing the utility of customer product testing: beta test design and management. *Journal of Product Innovation Management* 10, no. 4:318–330.
- Feigenbaum, Armand. 1991. *Total quality control*. 3rd edition. New York: McGraw-Hill.
- Fine, Charles. 1998. *Clockspeed*. Reading, Mass.: Perseus Books.
- Florida, Richard, and Martin Kenney. 1990. *The breakthrough illusion*. New York: Basic Books.
- Flynn, B., J. Flynn, S. Amundson, and R. Schroeder. 1999. Product development, speed, and quality: A new set of synergies. In *Perspectives in total quality*, edited by Michael Stahl. Oxford, England: Blackwell Publishers.
- Garvin, David. 1988. *Managing quality*. New York: Free Press.
- Gomory, Ralph. 1989. From the ‘ladder of science’ to the product development cycle. *Harvard Business Review* 67, no. 6:99–105.
- Hammond, Josh, and James Morrison. 1996. *The stuff Americans are made of*. New York: Macmillan.
- Hargadon, Andrew, and Kathleen Eisenhardt. 2000. Speed and quality in new product development: An emergent perspective on continuous organizational adaptation. In *The quality movement and organization theory*, edited by Robert E. Cole, and W. Richard Scott. Thousand Oaks, Calif.: Sage Publishing.
- Imai, Masaaki. 1986. *Kaizen*. New York: McGraw-Hill.
- Ishiyama, Takayuki. 1977. On system (sic) for applying FMEA and outline of its applications. *Reports of Statistical Applications and Research*, (Union of Japanese Scientists and Engineers) 24:40–50.
- Juran, Joseph, and A. Blanton Godfrey. 1999. *Juran’s quality handbook*. 5th edition. New York: McGraw-Hill.
- Lee, Gwen, and Robert E. Cole. 2000. The Linux kernel development as a model of knowledge creation. Paper presented at Strategic Management Society 20th Annual International Conference. Vancouver, British Columbia, 18 October 2000.
- Leonard-Barton, Dorothy. 1991. Inanimate integrators: a block of wood speaks. *Design Management Journal* (summer): 61–67.
- . 1995. *Wellsprings of knowledge*. Boston: Harvard Business School Press.
- LINUXCARE. 1999. Code freeze; ISDN perennial lateness. *Kernel Traffic*, August 3–10, 1999, (44 posts): Re: no driver change for 2.4k?; <http://kt.linuxcare.com/kernel-traffic/kt1999081931.ep1#9>.
- Lynn, Gary, Joseph Morone, and Albert Paulson. 1996. Marketing and discontinuous innovation: the probe and learn process. *California Management Review* 38, no. 3:8–37.

---

## From Continuous Improvement to Continuous Innovation

---

- Lynn, Leonard. 1982. *How Japan innovates*. Boulder, Colo.: Westview Press.
- Maguire, Miles. 1999. Cowboy quality. *Quality Progress* 32, no. 10:27–34.
- March, James. 1994. *A primer on decision making*. New York: Free Press.
- Moore, Geoffrey. 2000. *Living on the fault line*. New York: HarperCollins.
- Morgan, Gareth, 1997. *Images of organization*. 2nd edition. Thousand Oaks, Calif.: Sage Publishing.
- Paul, Gerald. 1999. Project management and product development. In *Juran's quality handbook* (5th edition), edited by Joseph Juran, and A. Blanton Godfrey. New York: McGraw-Hill. 17.1–17.20.
- Rao, Ashok, L. Carr, I. Dambolena, R. Kopp, J. Martin, F. Rafii, and P. Schlesinger. 1996. *Total quality management: A cross-functional perspective*. New York: John Wiley & Sons.
- Robinson, Alan. 1991. *Continuous improvement in operations*. Cambridge, Mass.: Productivity Press.
- Stalk, George, and Thomas Hout. 1990. *Competing against time*. New York: Free Press.
- Sutcliffe, Kathleen, Sim Sitkin, and Larry Browning. 2000. Tailoring process management to situational requirements: Beyond the control and exploration dichotomy. In *The quality movement and organization theory*, edited by Robert E. Cole, and W. Richard Scott. Thousand Oaks, Calif.: Sage Publishing.
- Teece, David. 1998. Capturing value from knowledge assets: The new economy, markets for know-how, and intangible assets. *California Management Review* 40, no. 3:55–79.
- Thomke, Stefan, and Donald Reinertsen. 1998. Agile product development: Managing development flexibility in uncertain environments. *California Management Review* 40, no. 5:8–30.
- Tushman, Michael, Philip Anderson, and Charles O'Reilly. 1997. Technology cycles, innovation streams, and ambidextrous organizations: Organizational renewal through innovation streams and strategic change. In *Managing strategic innovation and change*. New York: Oxford University Press
- Tyre, Marcie, and Wanda Orlikowski. 1994. Windows of opportunity: temporal patterns of technological adaptation in organizations. *Organization Science* 5, no. 1:98–118.
- Vesey, J. T. 1991. The new competitors: They think in terms of speed and market. *Academy of Management Executive* 5, no. 2:23–33.
- Weick, Karl, and Francis Westley. 1996. Organizational learning: Affirming an oxymoron. In *Handbook of organization studies*, edited by Stewart Clegg, Cynthia Hardy, and Walter Nord. London: Sage Publications.
- Wright, Paul. 2001. *21st century manufacturing*. Upper Saddle River, N.J.: Prentice Hall.

---

### BIOGRAPHY

**Robert E. Cole** is the Lorraine Tyson Mitchell II Chair in Leadership and Communication and the co-director of the Management of Technology Program at the Haas School of Business at the University of California–Berkeley. He is the author of *Managing Quality Fads: How American Business Learned to Play the Quality Game* (1999) and is the co-editor, with W. Richard Scott, of *The Quality Movement and Organizational Theory* (2000).

Cole earned a Ph.D. in sociology from University of Illinois. He may be contacted as follows: Haas School of Business, University of California–Berkeley, Berkeley, CA 94720-1900; 510-642-4295; Fax: 510-642-2826; E-mail: cole@Haas.Berkeley.EDU .