

Enlightened Experimentation: The New Imperative for Innovation

by Stefan Thomke



Harvard Business Review

Reprint R0102D



ILLUSTRATION BY A. SKWISH

The high cost of experimentation has long put a damper on companies' attempts to create great new products. But new technologies are making it easier than ever to conduct complex experiments quickly and cheaply. Companies now have an opportunity to take innovation to a whole new level – if they're willing to rethink their R&D from the ground up.

Enlightened Experimentation

The New Imperative for Innovation

by Stefan Thomke

EXPERIMENTATION LIES AT THE HEART OF EVERY company's ability to innovate. In other words, the systematic testing of ideas is what enables companies to create and refine their products. In fact, no product can be a product without having first been an idea that was shaped, to one degree or another, through the process of experimentation. Today, a major development project can require literally thousands of experiments, all with the same objective: to learn whether the product concept or proposed technical solution holds promise for addressing a new need or problem, then incorporating that information in the next round of tests so that the best product ultimately results.

In the past, testing was relatively expensive, so companies had to be parsimonious with the number of experimental iterations. Today, however, new technologies such as computer simulation, rapid prototyping, and combinatorial chemistry allow companies to create more learning more rapidly, and that knowledge, in turn, can be incorporated in more experiments at less expense. Indeed, new information-based technologies have driven down the marginal costs of experimentation, just as they have decreased the marginal costs in some production and distribution systems. Moreover, an experimental system that integrates new information-based technologies does more than lower costs; it also increases the opportunities for innovation. That is, some technologies can make existing experimental activities more efficient, while others introduce entirely new ways of discovering novel concepts and solutions.

Millennium Pharmaceuticals in Cambridge, Massachusetts, for instance, incorporates new technologies such as genomics, bioinformatics, and combinatorial chemistry in its technology platform for conducting experiments. The platform enables factory-like automation that can generate and test drug candidates in minutes or seconds, compared with the days or more that traditional methods require. Gaining information early on about, say, the toxicological profile of a drug candidate significantly improves Millennium's ability to predict the drug's success in clinical testing and, ultimately, in the marketplace. Unpromising candidates are eliminated before hundreds of millions of dollars are invested in their development. In addition to reducing the cost and time of traditional drug development, the new technologies also enhance Millennium's ability to innovate, according to Chief Technology Officer Michael Pavia. Specifically, the company has greater opportunities to experiment with more diverse potential drugs, including those that may initially seem improbable but might eventually lead to breakthrough discoveries.

This era of "enlightened experimentation" has thus far affected businesses with high costs of product development, such as the pharmaceutical, automotive, and software industries. By studying them, I have learned several valuable lessons that I believe have broad applicability to other industries. As the cost of computing continues to fall, making all sorts of complex calculations faster and cheaper, and as new technologies like combinatorial chemistry emerge, virtually all companies will discover that they have a greater capacity for rapid experimentation to investigate diverse concepts. Financial institutions, for example, now use computer simulations to test new financial instruments. In fact, the development of spread-

sheet software has forever changed financial modeling; even novices can perform many sophisticated what-if experiments that were once prohibitively expensive.

A System for Experimentation

Understanding enlightened experimentation requires an appreciation of the process of innovation. Namely, product and technology innovations don't drop from the sky; they are nurtured in laboratories and development organizations, passing through a *system* for experimentation. All development organizations have such a system in place to help them narrow the number of ideas to pursue and then refine that group into what can become viable products. A critical stage of the process occurs when an idea or concept becomes a working artifact, or prototype, which can then be tested, discussed, shown to customers, and learned from.

Perhaps the most famous example of the experimental system at work comes from the laboratories of Thomas Alva Edison. When Edison noted that inventive genius is "99% perspiration and 1% inspiration," he was well aware of the importance of an organization's capability and capacity to experiment. That's why he designed his operations in Menlo Park, New Jersey, to allow for efficient and rapid experimental iterations.

Edison knew that the various components of a system for experimentation—including personnel, equipment, libraries, and so on—all function interdependently. As such, they need to be jointly optimized, for together they define the system's performance: its speed (the time needed to design, build, test, and analyze an experiment), cost, fidelity (the accuracy of the experiment and the conditions under which it is conducted), capacity (the number of experiments that can be performed in a given time period), and the learning gained (the amount of new information generated by the experiment and an organization's ability to benefit from it). Thus, for example, highly skilled machinists worked in close proximity to lab personnel at Menlo Park so they could quickly make improvements when researchers had new ideas or learned something new from previous experiments. This system led to landmark inventions, including the electric lightbulb, which required more than 1,000 complex experiments with filament materials and shapes, electromechanical regulators, and vacuum technologies.

Edison's objective of achieving great innovation through rapid and frequent experimentation is especially pertinent today as the costs (both financial and time) of experimentation plunge. Yet many companies mistakenly view new technologies solely in terms of cost cutting, overlooking their vast potential for innovation. Worse, companies with that limited view get bogged down in the confusion that occurs when they try to incorporate new technologies. For instance, computer simulation doesn't

Stefan Thomke is an associate professor of technology and operations management at Harvard Business School in Boston.

simply replace physical prototypes as a cost-saving measure; it introduces an entirely different way of experimenting that invites innovation. Just as the Internet offers enormous opportunities for innovation—far surpassing its use as a low-cost substitute for phone or catalog transactions—so does state-of-the-art experimentation. But realizing that potential requires companies to adopt a different mind-set.

Indeed, new technologies affect everything, from the development process itself, including the way an R&D organization is structured, to how new knowledge—and hence learning—is created. Thus, for companies to be more innovative, the challenges are managerial as well as technical, as these four rules for enlightened experimentation suggest:

1. Organize for rapid experimentation.

The ability to experiment quickly is integral to innovation: as developers conceive of a multitude of diverse ideas, experiments can provide the rapid feedback necessary to shape those ideas by reinforcing, modifying, or complementing existing knowledge. Rapid experimentation, however, often requires the complete revamping of entrenched routines. When, for example, certain classes of experiments become an order of magnitude cheaper or faster, organizational incentives may suddenly become misaligned, and the activities and routines that were once successful might become hindrances. (See the sidebar “The Potential Pitfalls of New Technologies.”)

Consider the major changes that BMW recently underwent. Only a few years ago, experimenting with novel design concepts—to make cars withstand crashes better, for instance—required expensive physical prototypes to be built. Because that process took months, it acted as a barrier to innovation because engineers could not get timely feedback on their ideas. Furthermore, data from crash tests arrived too late to significantly influence decisions in the early stages of product development. So BMW had to incorporate the information far downstream, incurring greater costs. Nevertheless, BMW’s R&D organization, structured around this traditional system, developed award-winning automobiles, cementing the company’s reputation as an industry leader. But its success also made change difficult.

The Essentials for Enlightened Experimentation

New technologies such as computer simulations not only make experimentation faster and cheaper, they also enable companies to be more innovative. But achieving that requires a thorough understanding of the link between experimentation and learning. Briefly stated, innovation requires the right R&D systems for performing experiments that will generate the information needed to develop and refine products quickly. The challenges are managerial as well as technical:

1) Organize for rapid experimentation

- Examine and, if necessary, revamp entrenched routines, organizational boundaries, and incentives to encourage rapid experimentation.
- Consider using small development groups that contain key people (designers, test engineers, manufacturing engineers) with all the knowledge required to iterate rapidly.
- Determine what experiments can be performed in parallel instead of sequentially. Parallel experiments are most effective when time matters most, cost is not an overriding factor, and developers expect to learn little that would guide them in planning the next round of experiments.

2) Fail early and often, but avoid mistakes

- Embrace failures that occur early in the development process and advance knowledge significantly.
- Don’t forget the basics of experimentation. Well-designed tests have clear objectives (what do you anticipate learning?) and hypotheses (what do you expect to happen?). Also, mistakes often occur when you don’t control variables that could diminish your ability to learn from the experiments. When variability can’t be controlled, allow for multiple, repeated trials.

3) Anticipate and exploit early information

- Recognize the full value of front-loading: identifying problems upstream, where they are easier and cheaper to solve.
- Acknowledge the trade-off between cost and fidelity. Experiments of lower fidelity (generally costing less) are best suited in the early exploratory stages of developing a product. High-fidelity experiments (typically more expensive) are best suited later to verify the product.

4) Combine new and traditional technologies

- Do not assume that a new technology will necessarily replace an established one. Usually, new and traditional technologies are best used in concert.
- Remember that new technologies emerge and evolve continually. Today’s new technology might eventually replace its traditional counterpart, but it could then be challenged by tomorrow’s new technology.

Today, thanks to virtual experiments—crashes simulated by a high-performance computer rather than through physical prototypes—some of the information arrives very early, before BMW has made major resource decisions. The costs of experimentation (both financial and time) are therefore lower because BMW eliminates the creation of physical prototypes as well as the expense of potentially reworking bad designs after the company has committed itself to them. (Physical prototypes are still required much further downstream to verify the final designs and meet safety regulations.) In addition, the

rapid feedback and the ability to see and manipulate high-quality computer images spur greater innovation: many design possibilities can be explored in “real time” yet virtually, in rapid iterations.

To study this new technology’s impact on innovation, BMW performed the following experiment. Several designers, a simulation engineer, and a test engineer formed a team to improve the side-impact safety of cars. Primarily using computer simulations, the team developed and tested new ideas that resulted from their frequent brainstorming meetings.

The Potential Pitfalls of New Technologies

New technologies can slash the costs (both financial and time) of experimentation and dramatically increase a company’s ability to develop innovative products. To reap those benefits, though, organizations must prepare themselves for the full effects of such technologies.

Computer simulations and rapid prototyping, for example, increase not only a company’s capacity to conduct experiments but also the wealth of information generated by those tests. That, however, can easily overload an organization that lacks the capability to process information from each round of experiments quickly enough to incorporate it into the next round. In such cases, the result is waste, confusion, and frustration. In other words, without careful and thorough planning, a new technology might not only fail to deliver on its promise of lower cost, increased speed, and greater innovation, it could actually decrease the overall performance of an R&D organization, or at a minimum disrupt its operations.

Misaligned objectives are another common problem. Specifically, some managers do not fully appreciate the trade-off

between response time and resource utilization. Consider what happens when companies establish central departments to oversee computing resources for performing simulations. Clearly, testing ideas and concepts virtually can provide developers with the rapid feedback they need to shape new products. At the same time, computers are costly, so people managing them as cost centers are evaluated by how much those resources are being used.


The busier a central computer is, however, the longer it takes for devel-

opers to get the feedback they need. In fact, the relationship between waiting time and utilization is not linear—queuing theory has shown that the waiting time typically increases gradually until a resource is utilized around 70%, and then the length of the delays surge. (See the exhibit “Waiting for a Resource.”) An organization trying to shave costs may become a victim of its own myopic objective. That is, an annual savings of perhaps a few hundred thousand dollars achieved through increasing utilization from 70% to 90% may lead to very long delays for dozens of development engineers waiting for critical feedback from their tests.

A huge negative consequence is that the excessive delays not only affect development schedules but also discourage people from experimenting, thus squelching their ability to innovate. So in the long term, running additional computer equipment at a lower utilization level might well be worth the investment. An alternative solution is to move those resources away from cost centers and under the control of developers, who have strong incentives for fast feedback.

Waiting for a Resource

According to queuing theory, the waiting time for a resource such as a central mainframe computer increases gradually as more of the resource is used. But when the utilization passes 70%, delays increase dramatically.



Because all the knowledge required about safety, design, simulation, and testing resided within a small group, the team was able to iterate experiments and develop solutions rapidly. After each round of simulated crashes, the team analyzed the results and developed new ideas for the next round of experiments. As expected, the team benefited greatly from the rapid feedback: it took them only a few days to accept, refine, or reject new design solutions—something that had once taken months.

As the trials accrued, the group members greatly increased their knowledge of the underlying mechanics, which enabled them to design previously unimaginable experiments. In fact, one test completely changed their knowledge about the complex relationship between material strength and safety. Specifically, BMW's engineers had assumed that the stronger the area next to the bottom of a car's pillars (the structures that connect the roof of an auto to its chassis), the better the vehicle would be able to withstand crashes. But one member of the development team insisted on verifying this assumption through an inexpensive computer simulation.

The results shocked the team: strengthening a particular area below one of the pillars substantially *decreased* the vehicle's crashworthiness. After more experiments and careful analysis, the engineers discovered that strengthening the lower part of the center pillar would make the pillar prone to folding higher up, above the strengthened area. Thus, the passenger compartment would be more penetrable at the part of the car closer to the midsection, chest, and head of passengers. The solution was to weaken, not strengthen, the lower area. This counterintuitive knowledge—that purposely weakening a part of a car's structure could increase the vehicle's safety—has led BMW to reevaluate all the reinforced areas of its vehicles.

In summary, this small team increased the side-impact crash safety by about 30%. It is worth noting that two crash tests of physical prototypes at the end of the project confirmed the simulation results. It should also be noted that the physical prototypes cost a total of about \$300,000, which was more than the cost of all 91 virtual crashes combined. Furthermore, the physical prototypes took longer to build, prepare, and test than the entire series of virtual crashes.

But to obtain the full benefits of simulation technologies, BMW had to undertake sweeping changes in process, organization, and attitude—changes that took several years to accomplish. Not only did the company have to reorganize the way different groups worked together; it also had to change habits that had worked so well in the old sequential development process.

Previously, for example, engineers were often loath to release less-than-perfect data. To some extent, it was in each group's interest to hold back and monitor the output from other groups. After all, the group that submitted its

information to a central database first would quite likely have to make the most changes because it would have gotten the least feedback from other areas. So, for instance, the door development team at BMW was accustomed to—and rewarded for—releasing nearly flawless data (details about the material strength of a proposed door, for example), which could take many months to generate. The idea of releasing rough information very early, an integral part of a rapid and parallel experimentation process, was unthinkable—and not built into the incentive system. Yet a six-month delay while data were being perfected could derail a development program predicated on rapid iterations.

Thus, to encourage the early sharing of information, BMW's managers had to ensure that each group understood and appreciated the needs of other teams. The crash simulation group, for example, needed to make the door designers aware of the information it required in order to build rough models for early-stage crash simulations. That transfer of knowledge had a ripple effect, changing how the door designers worked because some of the requested information demanded that they pay close attention to the needs of other groups as well. They started to understand that withholding information as long as possible was counterproductive. By making these kinds of organizational changes, BMW in Germany significantly slashed development time and costs and boosted innovation.

2. Fail early and often, but avoid mistakes. Experimenting with many diverse—and sometimes seemingly absurd—ideas is crucial to innovation. When a novel concept fails in an experiment, the failure can expose important gaps in knowledge. Such experiments are particularly desirable when they are performed early on so that unfavorable options can be eliminated quickly and people can refocus their efforts on more promising alternatives. Building the capacity for rapid experimentation in early development means rethinking the role of failure in organizations. Positive failure requires having a thick skin, says David Kelley, founder of IDEO, a leading design firm in Palo Alto, California.

IDEO encourages its designers “to fail often to succeed sooner,” and the company understands that more radical experiments frequently lead to more spectacular failures. Indeed, IDEO has developed numerous prototypes that have bordered on the ridiculous (and were later rejected), such as shoes with toy figurines on the shoelaces. At the same time, IDEO's approach has led to a host of best-sellers, such as the Palm V handheld computer, which has made the company the subject of intense media interest, including a *Nightline* segment with Ted Koppel and coverage in *Serious Play*, a book by Michael Schrage, a co-director of the e-markets initiative at the MIT Media Lab, that describes the crucial importance of allowing innovators to play with prototypes.

Removing the stigma of failure, though, usually requires overcoming ingrained attitudes. People who fail in experiments are often viewed as incompetent, and that attitude can lead to counterproductive behavior. As Kelley points out, developers who are afraid of failing and looking bad to management will sometimes build expensive, sleek prototypes that they become committed to before they know any of the answers. In other words, the sleek prototype might look impressive, but it presents the false impression that the product is farther along than it really is, and that perception subtly discourages people from changing the design even though better alternatives might exist. That's why IDEO advocates the development of cheap, rough prototypes that people are invited to criticize—a process that eventually leads to better products. “You have to have the guts to create a straw man,” asserts Kelley.

To foster a culture in which people aren't afraid of failing, IDEO has created a playroomlike atmosphere. On Mondays, the different branches hold show-and-tells in which employees display and talk about their latest ideas and products. IDEO also maintains a giant “tech box” of hundreds of gadgets and curiosities that designers routinely rummage through, seeking inspiration among the switches, buttons, and various odd materials and objects. And brainstorming sessions, in which wild ideas are encouraged and participants defer judgment to avoid damping the discussion, are a staple of the different project groups.

3M is another company with a healthy attitude toward failure. 3M's product groups often have skunk-works teams that investigate the opportunities (or difficulties) that a potential product might pose. The teams, consisting primarily of technical people, including manufacturing engineers, face little repercussion if an idea flops—indeed, sometimes a failure is cause for celebration. When a team discovers that a potential product doesn't work, the group quickly disbands and its members move on to other projects.

Failures, however, should not be confused with mistakes. Mistakes produce little new or useful information and are therefore without value. A poorly planned or badly conducted experiment, for instance, might result in ambiguous data, forcing researchers

to repeat the experiment. Another common mistake is repeating a prior failure or being unable to learn from that experience. Unfortunately, even the best organizations often lack the management systems necessary to carefully distinguish between failures and mistakes.

3. Anticipate and exploit early information. When important projects fail late in the game, the consequences can be devastating. In the pharmaceutical industry, for example, more than 80% of drug candidates are discontinued during the clinical development phases, where more than half of total project expenses can be incurred. Yet although companies are often forced to spend millions of dollars to correct problems in the later stages of product development, they generally underestimate the cost savings of early problem solving. Studies of software development, for instance, have shown that late-stage problems are more than 100 times as costly as early-stage ones. For other environments that involve large capital investments in production equipment, the increase in cost can be orders of magnitude higher.

The Benefits of Front-Loaded Development

In the 1990s, Toyota made a major push to accelerate its product development cycle. The objective was to shorten the time from the approval of a body style to the first retail sales, thereby increasing the likelihood that Toyota kept up with the rapidly changing tastes of consumers.

Toyota made a concerted effort to identify and solve design-related problems earlier in product development—a concept known as *front-loading*. To accomplish that, the company implemented a number of initiatives, such as involving more manufacturing engineers during the product-engineering stage, increasing the transfer of knowledge between projects, investing substantially in computer-aided design and engineering tools, and developing rapid-prototyping capabilities.

To measure the benefits of these initiatives—and to monitor the company's evolving capabilities for early problem solving—Toyota tracked problems over multiple development projects. (See the exhibit “Solving Problems Earlier.”) The knowledge that a higher percentage of problems were being solved at earlier stages reassured Toyota's managers that they could aggressively reduce both development time and cost without risking product quality. In particular, between the first and third front-loading initiatives, Toyota slashed the cost (including the number of full physical prototypes needed) and time of development by between 30% and 40%.

It should be noted that in the early 1990s Toyota substantially reorganized its development activities, resulting in more effective communication and coordination between the different groups. This change most likely accounted for some of the performance improvements observed, particularly during the first front-loading initiatives.


In addition to financial costs, companies need to consider the value of time when those late-stage problems are on a project's critical path—as they often are. In pharmaceuticals, shaving six months off drug development means effectively extending patent protection when the product hits the market. Similarly, an electronics company might easily find that six months account for a quarter of a product's life cycle and a third of all profits.

New technologies, then, can provide some of their greatest leverage by identifying and solving problems upstream—best described as *front-loaded development*. In the automotive industry, for example, “quick-and-dirty” crash simulations on a computer can help companies avoid potential safety problems downstream. Such simulations may not be as complete or as perfect as late-stage prototypes will be, but they can force organizational problem solving and communication at a time when many downstream groups are not participating directly in development. (See the sidebar “The Benefits of Front-Loaded Development.”)

Several years ago, Chrysler (now DaimlerChrysler) discovered the power of three-dimensional computer models, known internally as digital mock-ups, for identifying certain problems in early development stages. When Chrysler developed the 1993 Concorde and Dodge Intrepid models, the process of decking—placing the power train and related components like the exhaust and suspension in the prototype automobile—took more than three weeks and required many attempts before the powertrain could be inserted successfully. By contrast, the early use of digital mock-ups in the 1998 Concorde and Intrepid models allowed the company to simulate decking to identify (and solve) numerous interference problems before the physical decking took place. Instead of taking weeks, decking was completed in 15 minutes because all obstruction problems had been resolved earlier—when it was relatively inexpensive and fast to do so.

Of course, it is neither pragmatic nor economically feasible for companies to obtain all the early information they would like. So IDEO follows the principle of three

Solving Problems Earlier



Source: Stefan Thomke and Takahiro Fujimoto, “The Effect of ‘Front-Loading’ Problem-Solving on Product Development Performance,” *The Journal of Product Innovation Management*, Vol. 17, No. 2, March 2000

As Toyota intensified its front-loading efforts, it was able to identify and solve problems much earlier in the development process.

In the early 1990s (see top graph), the first initiatives for front-loading began. Formal, systematic efforts to improve face-to-face communication and joint problem solving between the prototype shops and production engineers resulted in a higher relative percentage of problems found with the aid of first prototypes. Communication between different engineering sections (for instance, between body, engine, and electrical) also improved.

In the mid-1990s (see middle graph), the second front-loading initiatives called for three-dimensional computer-aided design, resulting in a significant increase of problem identification and solving prior to stage 3 (first prototypes).

In the ongoing third front-loading initiatives (see bottom graph), Toyota is using computer-aided engineering to identify functional problems earlier in the development process, and the company is transferring problem and solution information from previous projects to the front end of new projects. As a result, Toyota expects to solve at least 80% of all problems by stage 2—that is, before the first prototypes are made. And because the second-generation prototypes (stage 5) are now less important to overall problem solving, Toyota will be able to eliminate parts of that process, thereby further reducing time and cost without affecting product quality.

R's: rough, rapid, and right. The final R recognizes that early prototypes may be incomplete but can still get specific aspects of a product right. For example, to design a telephone receiver, an IDEO team carved dozens of pieces of foam and cradled them between their heads and shoulders to find the best possible shape for a handset. While incomplete as a telephone, the model focused on getting 100% of the shape right. Perhaps the main advantage of this approach is that it forces people to decide judiciously which factors can initially be rough and which must be right. With its three R's, IDEO has established a process that generates important information when it is most valuable: the early stages of development.

In addition to saving time and money, exploiting early information helps product developers keep up with customer preferences that might evolve over the course of a project. As many companies can attest, customers will often say about a finished product: "This is exactly what I asked you to develop, but it is not what I want." Leading software businesses typically show incomplete prototypes to customers in so-called beta tests, and through that process they often discover changes and problems when they are still fairly inexpensive to handle.

4. Combine new and traditional technologies. New technologies that are used in the innovation process itself are designed to help solve problems as part of an experimentation system. A company must therefore understand how to use and manage new and traditional technologies together so that they complement each other. In fact, research by Marco Iansiti of Harvard Business School has found that, in many industries, the ability to integrate technologies is crucial to developing superior products.

A new technology often reaches the same general performance of its traditional counterpart much more quickly and at a lower cost. But the new technology usually performs at only 70% to 80% of the established technology. For example, a new chemical synthesis process might be able to obtain a purity level that is just three-quarters that of a mature technique. Thus, by combining new and established technologies, organizations can avoid the performance gap while also enjoying the benefits of cheaper and faster experimentation. (See the exhibit "Combining the New with the Traditional.")

Indeed, the true potential of new technologies lies in a company's ability to reconfigure its processes and organization to use them in concert with traditional technologies. Eventually, a new technology can replace its traditional counterpart, but it then


might be challenged by a newer technology that must be integrated. To understand this complex evolution, consider what has happened in the pharmaceutical industry.

In the late nineteenth century and for much of the twentieth century, drug development occurred through a process of systematic trial-and-error experiments. Scientists would start with little or no knowledge about a particular disease and try out numerous molecules, many from their company's chemical libraries, until they found one that happened to work. Drugs can be likened to keys that need to fit the locks of targets, such as the specific nerve cell receptors associated with central nervous diseases. Metaphorically, then, chemists were once blind, or at least semiblind, locksmiths who have had to make up thousands of different keys to find the one that matched. Doing so entailed synthesizing compounds, one at a time, each of which usually required several days at a cost from \$5,000 to \$10,000.

Typically, for each successful drug that makes it to market, a company investigates roughly 10,000 starting candidates. Of those, only 1,000 compounds make it to more extensive trials in vitro (that is, outside living organisms in settings such as test tubes), 20 of which are tested even more extensively in vivo (that is, in the body of a living or-

Combining the New with the Traditional

A new technology will reach perhaps just 70% to 80% of the performance of a traditional technology. A new computer model, for instance, might be able to represent real-world functionality that is just three-quarters that of an advanced prototype model. To avoid this performance gap – and potentially create new opportunities for innovation – companies can use the new and traditional technologies in concert. The optimal time for switching between the two occurs when the rates of improvement between the new and mature technologies are about the same – that is, when the slopes of the two curves are equal.



ganism such as a mouse), and ten of which make it to clinical trials with humans. The entire process represents a long and costly commitment.

But in the last ten years, new technologies have significantly increased the efficiency and speed at which companies can generate and screen chemical compounds. Researchers no longer need to painstakingly create one compound at a time. Instead, they can use combinatorial chemistry, quickly generating numerous variations simultaneously around a few building blocks, just as today's locksmiths can make thousands of keys from a dozen basic shapes, thereby reducing the cost of a compound from thousands of dollars to a few dollars or less.

In practice, however, combinatorial chemistry has disrupted well-established routines in laboratories. For one thing, the rapid synthesis of drugs has led to a new problem: how to screen those compounds quickly. Traditionally, potential drugs were tested in live animals—an activity fraught with logistical difficulties, high expense, and considerable statistical variation.

So laboratories developed test-tube-based screening methodologies that could be automated. Called high-throughput screening, this technology requires significant innovations in equipment (such as high-speed precision robotics) and in the screening process itself to let researchers conduct a series of biological tests, or assays, on members of a chemical library virtually simultaneously.

The large pharmaceutical corporations and academic chemistry departments initially greeted such “combi-chem” technologies (combinatorial chemistry and high-throughput screening) with skepticism. Among the reasons cited was that the purity of compounds generated via combi-chem was relatively poor compared to traditional synthetic chemistry. As a result, many advances in the technology were made by small biotechnology companies.

But as the technology matured, it caught the interest of large corporations like Eli Lilly, which in 1994 acquired Sphinx Pharmaceuticals, one of the start-ups developing combi-chem. Eli Lilly took a few years to transfer the new technologies to its drug discovery division, which used traditional synthesis. To overcome the internal resistance, senior management implemented various mechanisms to control how the new technologies were being adopted. For example, it temporarily limited the in-house screening available to chemists, leaving them no choice but to use some of the high-throughput screening capabilities at the Sphinx subsidiary and interact with the staff there.

Until now, pharmaceutical giants like Eli Lilly have used combinatorial chemistry primarily to optimize promising new drug candidates that resulted from an exhaustive search through chemical libraries and other traditional sources. But as combinatorial chemistry itself advances and achieves levels of purity and diversity comparable to the compounds in a library, companies will in-

creasingly use it at the earlier phases of drug discovery. In fact, all major pharmaceutical companies have had to use combi-chem and traditional synthesis in concert, and the companies that are best able to manage the new and mature technologies together so that they fully complement each other will have the greatest opportunity to achieve the highest gains in productivity and innovation.

Enlightened Implications

New technologies reduce the cost and time of experimentation, allowing companies to be more innovative. Automotive companies, for example, are currently advancing the performance of sophisticated safety systems that measure a passenger's position, weight, and height to adjust the force and speed at which airbags deploy. The availability of fast and inexpensive simulation enables the massive and rapid experimentation necessary to develop such complex safety devices.

But it is important to note that the increased automation of routine experiments will not remove the human element in innovation. On the contrary, it will allow people to focus on areas where their value is greatest: generating novel ideas and concepts, learning from experiments, and ultimately making decisions that require judgment. For example, although Millennium's R&D facilities look more and more like factories, the value of knowledge workers has actually increased. Instead of carrying out routine laboratory experiments, they now focus on the early stages (determining which experiments to conduct, for instance) and making sense of the information generated by the experimentation.

The implications for industries are enormous. The electronic spreadsheet has already revolutionized financial problem solving by driving down the marginal cost of financial experimentation to nearly zero; even a small start-up can perform complex cash-flow analyses on an inexpensive PC. Similarly, computer simulation and other technologies have enabled small businesses and individuals to rapidly experiment with novel designs of customized integrated circuits. The result has been a massive wave of innovation, ranging from smart toys to electronic devices. Previously, the high cost of integrated-circuit customization made such experimentation economical to only the largest companies.

Perhaps, though, this era of enlightened experimentation is still in its bare infancy. Indeed, the ultimate technology for rapid experimentation might turn out to be the Internet, which is already turning countless users into fervent innovators. 

Reprint R0102D

To place an order, call 1-800-988-0886.