



Modes of experimentation: an innovation process—and competitive—variable

Stefan Thomke ^a, Eric von Hippel ^{b,*}, Roland Franke ^c

^a *Harvard Business School, Boston, MA 02163, USA*

^b *MIT Sloan School of Management, Cambridge, MA 02139, USA*

^c *Natural Pharmaceuticals Inc., Cambridge, MA, USA*

Received 5 February 1997; revised 24 October 1997; accepted 23 April 1998

Abstract

The outputs of R&D, such as new research findings and new products and services, are generated with the aid of specialized problem-solving processes. These processes are somewhat arcane and have been largely ignored in studies of technical change. However, their improvement can significantly affect the kinds of research problems that can be addressed, the efficiency and speed with which R&D can be performed, and the competitive positions of firms employing them. In this paper, we first describe the general nature of the trial-and-error problem-solving processes and strategies for experimentation used in the development of new products and services. We next discuss the rapid advances being made in problem-solving methods, and the impact such advances can have on the competitive position of adopting firms. Finally, we offer a detailed case study of the impact one novel experimental method, combinatorial chemistry, is having on the economics of the drug discovery process. © 1998 Elsevier Science B.V. All rights reserved.

Keywords: Problem-solving; Experimentation; Technological innovation; R&D efficiency; Sticky information; Pharmaceutical drug development; Product development performance

1. Introduction

The impact of the outputs of the R&D process on firms and industries has long been acknowledged. For example, the major consequences of the development and continued improvement of semiconductors computerized manufacturing and have been noted by many. However, the outputs of R&D are themselves ‘manufactured’ with the aid of specialized problem-solving processes. These underlying processes have been largely ignored in studies of technical change.

The application of problem-solving processes, on the other hand, represent an increasing proportion of economic activity (Carter, 1995), and the processes themselves are improving rapidly both in terms of the kinds and efficiencies of outputs producible. These changes, in turn, are having and will increasingly have an impact on the competitive position of adopting firms.

In this paper, we will explore the general nature of the problem-solving processes used in R&D, and the potential impact of novel problem-solving methods on firm R&D performance and competitive standing. We begin by describing the problem-solv-

* Corresponding author.

ing process used in R&D in general terms, and showing how different experimentation strategies can influence R&D efficiency (Section 2). We next illustrate the rapid rate of advance affecting experimental methods, and then observe that early developers or adopters can have a competitive advantage over rivals, because new methods are often difficult to acquire and use (Section 3). We then present a field study of combinatorial chemistry—a new method now being introduced into the drug discovery process that promises to make drug discovery both faster and less costly (Section 4). This case is especially relevant to our topic because more effective drug discovery methods can convey a great competitive advantage to pharmaceutical firms: present drug discovery processes are currently very lengthy and expensive, and the commercial advantages to being first with a significant new product can be very large.

2. The problem-solving process

Research into the nature of problem-solving shows that it consists of trial and error, directed by some amount of insight as to the direction in which a solution might lie (Barron, 1988). This general finding is supported by empirical studies of problem-solving in the specific arena of product and process development (Allen, 1966; Alexander, 1964; Clark and Fujimoto, 1991; Iansiti, 1997; Marples, 1961; Smith and Eppinger, 1997; Thomke, 1997; von Hippel and Tyre, 1995; Wheelwright and Clark, 1992). Such studies show trial-and-error learning conducted via a process of conscious experimentation as a prominent feature. In this section, we begin by discussing the general nature of trial-and-error problem-solving via experimentation. Then, we discuss the creation of strategies for solving a given problem via a number of related experiments.

2.1. *Problem-solving via experimentation*

Experimentation using trial-and-error problem-solving begins with the selection or creation of one or more possible solutions. The alternatives selected may or may not include the best possible solutions—one has no way of knowing. These are then tested

against an array of requirements and constraints (Duncker, 1945; Marples, 1961; Simon, 1969). The new information provided by a trial-and-error experiment to an experimenter are those aspects of the outcome that he or she did not (was not able to) know or foresee or predict in advance—the ‘error’. Test outcomes are used to revise and refine the solutions under development, and generally, progress is made in this way towards an acceptable result.

One may view the experimental trial-and-error process as cycles that repeatedly ‘generate and test’ design alternatives (Simon, 1969). For example, one might conceive of, design, and build a prototype of a new, more rapidly deploying airbag for a car (generate alternative) and run an experiment to evaluate its actual deployment speed (test alternative). If the results of the first experiment are satisfactory, one stops. However, if the analysis shows that the results of the initial experiment are not satisfactory, as is usually the case, one may elect to modify one’s experiment and ‘iterate’—try again. Modifications may involve the experimental design, the experimental conditions, or even the nature of the desired solution. For example, a researcher may design an experiment with the goal of identifying a new cardiovascular drug. However, experimental results obtained on a given compound might suggest a different therapeutic use, and cause researchers to change their view of an acceptable solution accordingly.

Experimentation is often carried out using simplified versions (models) of the eventually-intended test object and/or test environment. For example, aircraft designers often conduct experiments on possible aircraft designs by testing a scale model of that design in a ‘wind tunnel’—an apparatus that creates high wind velocities that partially simulate the aircraft’s intended operating environment. The value of using models in experimentation is two-fold: to reduce investment in the aspects of reality that are irrelevant for the experiment, and to ‘control out’ some aspects of reality that would affect the experiment in order to simplify analysis of the results. Thus, models of aircraft being subjected to wind tunnel experiments generally include no internal design details such as the layout of the cabins—these are both costly to model and typically irrelevant to the outcome of wind tunnel tests, which are focused

on the interaction between rapidly moving air and the model's exterior surface.

Models used in experimentation can be physical in nature, as in the example just given, or they can be represented in other forms. Computer simulation, for example, involves representing experimental objects and experimental environments in digital form, and then simulating their interaction within a computer in a type of virtual experiment (Thomke, 1998). Thus, one might model an automobile and a crash barrier inside a computer, perform the computations needed to simulate the crash of the model car into the model barrier, and then calculate the effects of that crash on the structure of the car via finite element analysis. One could then assess the results of this virtual experiment by viewing a visual display of the 'crashed' car on a video display, and/or by looking at detailed calculations of the forces and accelerations generated during the simulated crash and the effects of these on the car's structure.

Designers will sometimes test a real experimental object in a real experimental context only after experimenting with several generations of models that isolate different aspects of the real and/or that encompass increasing amounts of the complexity of the real. Developers of pharmaceuticals, for example, might begin by testing a candidate drug molecule against just the purified enzyme or receptor it is intended to affect, and then test it again and again against successively more complex models of the human organism (e.g., tissue extracts, tissue culture, animal models, etc.) before finally seeking to test its effect on real human patients during clinical trials.

Models do not represent reality completely (if they did, they would be the reality they are to represent). This is, in part, by design and for the purposes mentioned earlier. The representation is also, in part, incomplete because one does not know and/or cannot economically capture all the attributes of the real situation, and so could not transfer them into a model even if one wanted to. The incompleteness of a model is a source of unexpected errors when a given model being used in testing is replaced by a different model or by the real context or object for the first time (Tyre and von Hippel, 1997). As an illustration, consider the airbag inflation example given earlier. If the gas used to inflate the airbag had been toxic, and the various experimental apparatus

used to test the airbag had not been capable of detecting this factor, the problem would have been detected as an unexpected error only when real airbags were deployed in the real use environment.

2.2. *Parallel and serial strategies for experimentation*

Researchers engaging in problem-solving via experimentation generally do not expect to solve a problem via a single experiment, and so often plan a series of experiments intended to bring them a solution to their problem in an efficient manner. The strategy they choose is in part a function of the information they have regarding the topography of the 'value landscape' which they plan to explore when seeking a solution for their problem (Alchian, 1950; Baldwin and Clark, 1997a).¹

A value landscape can be visualized as a flat plain with one or more hills rising from it. The total landscape represents the arena that the experimenters plan to search as to identify an acceptable solution to their problem. The probability of finding a solution increases as one ascends the 'hills' in the landscape, and so the experimenters' goal is to devise a series of experiments that will enable them to identify and explore those hills in an efficient manner. Real-world experimenters may not have much information regarding the value landscape they plan to explore when they begin their work—and may even abandon one landscape and switch to another as their work proceeds. (Explorations of the specification of 'well-structured' problems have shown that problem-solvers often may pursue a solution to a problem across a range of value landscapes rather than simply seeking to search a given landscape in

¹ The concept of a 'value landscape' is related to the study of evolutionary biology which regards fitness landscapes as the distribution of fitness values across a space of entities (Kauffman and Levin, 1987; Wright, 1932). More recently, fitness landscapes have been used in the study of organizational structure and strategy in the context of changing environments (Bruderer and Singh, 1996; Tushman and O'Reilly, 1996; Levinthal, 1997). In order to distinguish between biological evolution and the design and experimentation process, we instead use the term 'value landscape' for the remainder of the paper (for a good explanation of fitness and value landscapes and their respective differences, see Baldwin and Clark, 1997a).

an efficient manner.²) Nonetheless, experimenters' expectations regarding the topography of the value landscape(s) they have chosen are central to their construction of efficient experimental strategies.

As an illustration, consider the choice between a strategy of serial experimentation vs. parallel experimentation. When identification of a satisfactory solution to a problem involves more than one trial-and-error experiment, the information gained from a previous experiment(s) *may* serve as an important input to the design of the next one. Experiments which incorporate learning derived from other experiments in a set are considered to have been conducted in series. Experiments that are conducted according to an established plan that is not modified as a result of the finding from other experiments are considered to have been conducted in parallel. For example, one might carry out a pre-planned 'array' of experiments, analyze the results of the entire array, and then carry out one or more additional verification experiments (Montgomery, 1991). The experiments in the initial array are viewed as being carried out in parallel, while those in the second round are carried out in series with respect to that initial array.

Suppose that the problem at issue is to deduce the correct combination for a combination lock. Good locks may have 10^6 or more possible combinations, of which only one is correct. They are also designed so as to give an 'experimenter' (in this case, a robber) no indication as to how close he or she may be to the correct combination. That is, they are

designed to display a value landscape that is absolutely flat for all combinations except the correct one, which can be visualized as rising up from the landscape like a narrow tower with vertical sides. In a value landscape with this topography, a parallel experimentation strategy would be the fastest, although not necessarily the most efficient choice (see Table 1, strategies (a) and (c) and related discussion below). This is because in this landscape configuration, each failed trial provides very little information that would be of use in a serial experimentation strategy—only the information that 'the combination you just tried is not the correct one.'³

In contrast, suppose that the value landscape is a hill with only a single peak and sides that extend to all edges of the landscape. (This is the shape, for example, of the value landscape in the children's game in which a child is guided to a particular spot via a feedback from other children who say 'warmer' each time a step is taken towards that spot.) In such a case, a strategy of serial experimentation may be the most efficient choice, because the information gained from each step taken is so useful in guiding the direction of the next trial step that the correct solution is often found after only a few trials.

The relative efficiency of experimentation strategies can be estimated using what is known about the topography of the solution space, and what is known about the time and money costs associated with generating and testing alternatives in the solution space. Consider the following very simple search model in which the topography of the value landscape is known to consist of n points and to have the configuration described in the lock example dis-

² Well-structured problems have value landscapes for which one can precisely specify a process of trial and error that will lead to a desired solution in a 'practical' amount of time (Reitman, 1965; Simon, 1973; Pople, 1982). For example, a 'traveling salesman' problem of 'a size amenable to practical computation' is well-structured, because one can precisely specify a generator of alternative solutions and a solution-testing procedure that are guaranteed to eventually identify the best solution. (A traveling salesman problem involves determining the most efficient itinerary for a salesman who must physically visit each of a given list of cities.) A real-world problem-solver facing a traveling salesman problem may solve this problem as given or may decide to modify it—thereby creating a new value landscape(s) to explore. For example, the problem-solver might modify the original problem by deciding to consider the option of contacting customers in the specified list of cities by using the Internet rather than by arranging physical visits by a single salesman.

³ Simon (1969), (p. 206) uses a similar example in explaining problem-solving as natural selection, noting that the example was originally supplied by W. Ross Ashby. "Suppose that the task is to open a safe whose lock has 10 dials, each with 100 possible settings, numbered from 0 to 99. How long will it take to open the safe by a blind trial-and-error search for the correct setting? Since there are 100^{10} possible settings, we may expect to examine about one-half of these, on the average, before finding the correct one [...]. Suppose, however, that the safe is defective, so that a click can be heard when the dial is turned to the correct setting. Each dial can now be adjusted independently and does not need to be touched again while the others are being set. The total number of settings that has to be tried is only 10×50 , or 500."

Table 1

A comparison of different experimentation strategies (each trial consists of a generate and test step)

Massively parallel experimentation is particularly beneficial if the value of time is highly relative to the cost of an experiment. Real-world experimentation usually involves parallel and serial strategies, where the optimal combination is driven by many factors such as the cost and time to generate and test an alternative, the topography of the solution space, and prior knowledge of the experiment

Experi- mentation Strategy	(a) Parallel Experimentation	(b) Serial Experimentation (Rapid Learning)	(c) Serial Experimentation (Minimal Learning)
Learning between periods	None	50% reduction of search space after each period	Eliminate only un- successful alternative after each period
Expected # of periods	1	$\log_2 n$	$\frac{n+1}{2}$
# of trials per period	n	1	1
Expected # of trials	n	$\log_2 n$	$\frac{n+1}{2}$
Attributes	Highest cost Fastest speed	Lowest cost Medium speed	Medium cost Slowest speed

cussed above—flat except for a single point representing the correct solution.⁴

A parallel experimentation strategy (strategy (a) in Table 1) would require all experiments and their

⁴ Much more sophisticated models of search have been applied to the study of the R&D and design process. See, for example, Nelson (1961), Abernathy and Rosenbloom (1968), Evenson and Kislev (1976), Weitzman (1979), Roberts and Weitzman (1981), Nelson and Winter (1982), Nelson (1982), and Baldwin and Clark (1997b). The purpose of our simpler model is to help the reader in understanding the cost and time trade-offs between parallel and serial experimentation strategies.

tests to be done at the same time. Thus, one would not be able to incorporate what one has learned from 1 trial and apply it to the next trial. While this approach results in a very high number of experiments (n), it also reduces the total development time significantly as all experimental trials are done in parallel. Thus, in the case of this example, massively parallel experimentation would be the costliest but also the fastest strategy.

In contrast, a serial strategy applied to this sample problem would allow one to learn from each experimental trial and—equipped with this new knowledge—carefully select the next one. As shown in Table 1

(c), a strategy even with minimal learning (i.e., not repeating a trial that has failed) can halve the total number of experiments required on average, but would dramatically increase total development time relative to the purely parallel approach.⁵

Of course, if there is the opportunity for greater learning from each trial, the number of trials in the series likely to be required to reach the solution (and therefore the total elapsed time) is further reduced. For example, consider a very favorable learning scenario where the n trials are arranged on a linear scale (e.g., n different pressure settings) and that after each trial, one could learn whether to move up or down on that scale. Thus one would effectively reduce the search space by 50% after each experimental cycle and rapidly progress towards an optimal solution. An experimenter would start with $n/2$ (the midpoint) and move to either $n/4$ or $3n/4$, depending on the outcome of the first experiment, and continue in the same fashion until the solution is found. A real-world example for such a search can be found in the practice of system problem identification: very experienced electronic technicians tend to start in the middle of a system, find the bad half, and continue to subdivide their search until the problem is found. One can easily see that the expected number of trials until success using such a serial strategy (with the kind of learning described) can be reduced to $\log_2 n$ —a dramatic reduction in cost. However, total development time would exceed that of the purely parallel strategy by the same factor (see Table 1, strategy (b)).⁶

Real-world experimentation strategies can be

⁵ Assume that the set of possible experimental trials is of size n . After an alternative is generated, a screen tests if it is the solution (there exists only one solution in n). If the experimental trial results in a solution, the experimenter stops. If the experimental trial is unsuccessful, the experimenter randomly generates another alternative and continues. The experimenter only learns which trials have failed and thus should be avoided going forward—i.e., the experimental learning is minimal.

⁶ A related strategy is the renowned Newton's method that is widely used in numerical analysis and was first introduced in Newton's *Principia Mathematica* to solve a cubic polynomial. The method's iterative and sequential search, known for its rapid convergence, is guided by knowledge of the underlying function and its gradient to quickly find an accurate estimate of the numerical value being sought (Gerald and Wheatley, 1984).

much more complex than our simple model, and will often contain a combination of serial and parallel approaches (see, for example, Ward et al., 1995). As we will see in the discussion of drug discovery in Section 4, pharmaceutical firms typically employ both serial and parallel experimentation in the search for promising candidate drug molecules. Factors such as cost and time of generating and testing alternatives, and knowledge of the topography of the solution space are affecting the degree to which parallelism is employed in the drug discovery process.

3. Advances in experimentation and problem-solving and their implications for firms

The methods and tools available to help solve many types of problems are rapidly changing and improving. These advances are affecting all of the elements of the experimentation process that have been described in Section 2. That is, they are rapidly reducing the cost and time involved in designing and executing and analyzing many types of experiments, and are also affecting the type of experimentation strategies that may be most effective for an experimenter. In this section, we first illustrate the rapid advances being made in experimental methods by noting the rapid evolution⁷ of experimentation via computer simulation. We will next discuss the implications that such advances can have for the competitive position of firms.

3.1. Advances in methods for experimentation and problem-solving

There is no general index that documents the rate of advance in problem-solving methods and tools. However, those with a professional interest in these matters generally judge that the rate of change today is very rapid. Advances in some fields, such as computer simulation, are applicable to a wide variety of subject matter. Others, such as the scanning tunneling electron microscope, are germane to only a

⁷ We regard the evolution of technologies and methods as a decentralized, adaptive process that can be characterized by interplay between users and developers. Because both derive economic value from advances in areas such as simulation, they tend to reinforce each other and thus accelerate the overall evolution of such novel technologies and methods.

narrow range of applications—although the range of application seen for a given technique often broadens significantly over time (Rosenberg, 1982). The reader may find a brief overview of the rapid evolution of computer simulation techniques to be a useful way to gain a feeling for what we mean by ‘rapid advances in methods and tools’ in the case of a generally applicable tool. Later, in Section 4, we will provide a detailed description of the nature of and impact of an advance with a narrower range of application—combinatorial chemistry.

As was noted earlier, experimentation via computer simulation involves representing experimental objects and experimental environments in digital form, rather than in the form of physical objects tested within physical environments. Then, their interaction is within a computer in a type of virtual experiment. The advantages of substituting virtual experimentation via a computer for experimentation with real physical objects can be very significant. For example, studying automobile structures via real car crashes is clearly quite expensive and time-consuming—a crash prototype can cost in excess of US\$100,000 and may take up a year to build and test. In contrast, once the proper digital models have been created, a virtual car crash can be run again and again within a computer under varying conditions at very little additional cost per run. Furthermore, consider that a real car crash experiment happens very quickly—so quickly that the experimenter’s ability to observe details is typically impaired, even given high-speed cameras and well-instrumented cars and crash dummies. In contrast, one can instruct a computer to enact a virtual car crash as slowly as one likes, and can zoom in on any structural element of the car (or minute section of a structural element) that is of interest and observe the forces acting on it and its response to those forces ‘during’ the crash. Thus, computer simulation may not only decrease the cost and time of an experimental cycle but can also increase the depth and quality of analysis, leading to improved learning and ultimately products of higher quality (Thomke, 1998).

The steady (and really quite spectacular) improvement in the capabilities of digital computers over the past few decades has made it possible and desirable to carry out more and more experiments via computer simulation, rather than via physical experimen-

tation. Computer simulation is today being used as a substitute for or supplement to physical experimentation in fields ranging from the design of drugs (e.g., rational drug design) to the design of mechanical products (e.g., finite element analysis), to the design of electronic products (e.g., simulation of digital circuitry), and to the analysis of financial positions (e.g., simulation of novel financial instruments). The ability to usefully substitute a simulation for a ‘real’ experiment requires, of course, more than the development of advanced computer equipment. It also requires the development of simulation models that are accurate from the point of view of a given experimental purpose. Often, a simulation model will not be fully accurate in ways that later turn out to matter. When this is recognized, virtual and physical experiments may be conducted in some combination in order to combat this source of error. For example, auto designers will supplement data gathered from virtual car crash experiments with data from real crash experiments using real cars, in order to assure themselves that the results of the virtual experiments also hold in the real world.

At the same time, of course, methods for conducting physical experiments are also advancing. For example, significant advances are being made in reducing the costs and time of building the various types of prototypes. Complex three-dimensional objects used to require days or weeks of work in a machine shop to fabricate. Many such shapes can now be made rapidly—in very few hours—by using computer-controlled machining equipment and/or equipment for creating objects via ‘three-dimensional printing’ (Sachs et al., 1992). This also applies to physical prototypes of complex electrical circuitry—custom integrated circuits—used to take months to create via ‘full custom’ methods, and weeks to create via ‘Application-Specific Integrated Circuits’ (ASIC) technology. Now, designers can create customized circuits in minutes at their desks or lab benches using the so-called ‘Field Programmable Gate Arrays’ (FPGAs) (Thomke, 1997; Villasenor and Mangione-Smith, 1997).

3.2. Impact of changes in experimental methods on firm competitiveness

The adoption of more effective experimental methods for problem-solving and the development of

new products and services, such as those just described, can lead to significant competitive advantages for adopting firms relative to rivals *if* novel techniques that offer such advantages are not rapidly picked up by rivals as well. Or, as Barney (1986) and Wernerfelt (1984) put it with respect to core competencies: a core competence can be a source of long-term competitive advantage for a firm if it is difficult or impossible to buy or sell in the available factor markets, and if it is difficult to replicate.

We argue that the new and more effective experimental methods and techniques that are rapidly emerging are indeed often difficult to buy and sell, and difficult to replicate as well, and can therefore serve as a significant source of long-term competitive advantage for innovators and early adopters. The reason is that new methods require (1) the transfer of significant amounts of new information to the adopting firm, including new skills, and (2) some reorganization of a firm's R&D activities as well.

The requirement that new information must be transferred to a firm adopting a new experimental technique is in itself a barrier to adoption in many instances, because information is often costly to transfer to a new site in a form usable by a given information seeker, that is, sticky information. Transfer costs are affected by attributes of the information itself (e.g., how the information is encoded), and also by attributes of and choices made by information seekers and information providers (Arora and Gambardella, 1994; Cohen and Levinthal, 1990; Griliches, 1957; Mansfield, 1968; Nelson, 1982; Pavitt, 1987; Rosenberg, 1982; Teece, 1977; von Hippel, 1994).

Thus, consider that only *some* of the information associated with the ability to execute new experimental methods may be embodied in equipment that can be purchased and installed by an adopting firm—a relatively easy form of transfer. For example, a firm can buy computers and computer programs that can be used to do experiments via computer simulation, but new equipment and new software provide only a portion of the information a firm needs to actually become competent at performing a new experimental method. Typically, new skills and expertise are also needed, and, as Polanyi (1958) has pointed out, skill and expertise are often encoded within an expert's mind as tacit information that is

difficult to transfer to another. For example, in a study of biology lab practising an experimental method known as cell fusion, Barley and Bechky (1994) (pp. 98–99) reported that "...experienced research support specialists and technicians [carrying out the cell fusion work] made use of signs that could not be found in textbooks, and that were difficult to define except ostensively. Partially for this reason, practices successful in one lab often failed in another unless technicians from the first trained technicians from the second..."

Adopting novel experimental methods may also require considerable change in the organizational arrangements prevailing in the adopting firm. As Morison (1966) and Schön (1967) have pointed out, organizations are often built up around and adapted to existing technologies. When this is so, changes in technologies may require changes to organizational structures and routines. As an illustration, suppose that a firm wishes to replace some physical experimentation methods being carried out in its labs with computer simulation methods. To do this, it must typically hire new kinds of people and also reorganize the relationships between the various specialists who jointly carry out the experiments. In its existing organizational arrangements designed for physical experimentation, for example, the firm might have routines in place that enable researchers to work with design engineers and modelmakers to design and build the experiments that they wished to run. Next, the procedures might dictate that the completed experimental apparatus be transferred to experts at specialized test facilities who would actually run the experiments, collect the resulting data, and then supply that data to the researcher for analysis. In contrast, experimentation via computer simulation would require quite different organizational routines. In some cases, these would enable the researcher to do the entire experimental cycle in his or her own lab. In other cases, they might facilitate collaborative arrangements between the researcher and various types of experts not previously employed by the firm who specialize in different aspects of computer simulation.

With respect to the difficulty of achieving such organizational change, we note that Holmstrom and Tirole (1991) have argued that organizational arrangements cannot serve as sources of enduring

competitive advantage because they can be easily replicated. However, much of the literature on organizational change suggests otherwise (see, for example, Milgrom and Roberts, 1990; Henderson and Cockburn, 1994). Thus, Henderson and Cockburn (1994), in a study of cardiovascular drug discovery, reported that organizational capabilities found associated with improved productivity at this type of research task are in fact often very difficult to transfer from firm to firm.⁸ Further, they note that such arrangements can have an important impact on research productivity. In their study, about 30% of the observed variation in the ‘productivity’ of firms in drug discovery (number of drugs discovered per R&D dollar invested) was due to unique organizational capabilities (represented by a variety of measures such as to the degree to which the firm actively manages the integration of knowledge across disciplinary and firm boundaries).

4. Field study: the impact of new drug discovery methods on pharmaceutical drug development

To this point, we have described the general nature of problem-solving via experimentation in R&D and related parallel and serial strategies, have observed that methods for accomplishing this task are evolving rapidly, and have argued that competence at problem-solving via experimentation can be

⁸ Henderson (1994) (pp. 624–626) illustrates difficulties associated with replicating organizational capabilities associated with better performance at drug discovery by presenting examples experienced by firms in their sample. Thus, there was a period when leading-edge drug discovery processes were shifting from simple mass screening of compounds for possible medicinal effects to a more precise form of research based on an understanding of a drug’s mechanism of action. This change was being driven by the academic research community. The drug firm ‘Alpha,’ which had long-term ties to the academic community and which employed leading-edge researchers who were accepted as peers in that community, had no difficulty in quickly adopting the new approach to drug discovery. In contrast, the firm ‘Beta,’ which had not had a practice of employing scientists known to and respected by the academic community, found it very difficult to make the change. For example, they found it difficult to hire ‘better’ people from academia who were experts in the new approach, because they did not have a reputation as a leading-edge place to work.

important with respect to the competitiveness of firms that perform R&D. In this section, we develop these points further via a field study of a recent improvement in experimental methods used in the drug discovery process—‘combinatorial chemistry’. We begin by describing the serious drug development problem currently facing pharmaceutical firms. Next, we describe the drug discovery process and then, we describe combinatorial chemistry. Finally, we describe a research project that clearly illustrates the impact that this new method can have on the drug discovery process—and with it, upon the competitiveness of pharmaceutical firms.

4.1. The product development problem facing pharmaceutical firms

If improvements in problem-solving methods are important to any firm, they should certainly be important to firms in the pharmaceutical industry. On the one hand, pharmaceutical firms face many potentially profitable opportunities to create new drugs to cure or ameliorate diseases ranging from cancer to heart disease, particularly if firms manage to receive patent protection and reach the market before their competitors do. Markets for new drugs typically involve US\$50–400 million in annual sales, and can sometimes be in billions as in the case of Zantac, a stomach acid inhibitor drug for ulcer treatment. On the other hand, the drug development process is currently one of the most time-consuming and costliest product development processes in any industry.

A widely cited study of pharmaceutical drugs developed between 1972 and 1987 found that the expected capitalized development cost per marketed drug was on the average US\$230.8 million (1987) (DiMassi et al., 1991), with total development times well above 10 years. Various other studies have shown a trend that has caused much concern in the pharmaceutical industry: the cost and time of new drug development has increased significantly over the last 30 years (e.g., DiMassi et al., 1994). Besides the impact of lower R&D productivity on firm cost and profitability, longer development times have also raised important public policy concerns. As the industry remains the dominant provider of life-saving and life-prolonging medicines, it is in the public

Table 2
Expected phase costs per New Chemical Entity (NCE) (in millions of 1987 dollars) (from DiMassi et al., 1991)^a

Testing phase ^b	Uncapitalized expected cost	Mean phase length (months)	Capitalized expected cost ^c
Preclinical	65.5	42.6	155.6
Long-term animal	5.3	33.6	8.2
Other animal	0.4	33.6	0.7
Phase I	9.3	15.5	17.8
Phase II	12.9	24.3	21.4
Phase III	20.2	36.0	27.1
Total	113.6		230.8

^bThe New Drug Approval (NDA) review period was estimated to last 30.3 months.

^cCosts were capitalized at a 9% discount rate.

^aAll costs were deflated using the GNP Implicit Price Deflator. A 23% clinical approval rate was utilized.

interest to have promising new drugs available to patients as quickly as possible (Savello, 1996).

The complete drug development and approval process involves three phases. It begins with a pre-clinical research phase devoted to the discovery and optimization of one or a few 'lead' chemical compounds that appear to hold sufficient promise as drugs to merit investment in clinical testing. Phase II, clinical development, consists typically of three clinical phases to determine and document the safety and efficacy of the proposed drugs. The final phase Phase III, involves regulatory New Drug Approval (NDA) review processes of the clinical trial outcome. The average cost and duration of preclinical and clinical development for drugs developed between 1972 and 1987 are provided in Table 2.

4.2. The drug discovery process

Drugs achieve their effect by binding with very specific molecular receptors or enzymes or biologically important molecules that are present in the human body or on/in disease-causing agents such as bacteria, fungi, and viruses. The goal of drug discovery or drug design is therefore to discover or create a molecule that will in fact bind to a particular, say, receptor with a required degree of tenacity (binding affinity), and that will, at the same time, not bind to other receptors that may be structurally similar but have different functions.

The drug discovery process can involve either or a combination of two basic approaches. (1) One can start with little or no knowledge about the structure of a disease target (receptor, enzyme, molecule) associated with a particular disease, and simply try out many candidate molecules until one finds one that happens to bind properly with the target receptor. (2) One can strive to determine the structure of the relevant receptor with biophysical methods, and then attempt to design or select a molecule that will bind to it.

Until the 1970s, methods of drug discovery necessarily relied on the first of these two approaches because the technical ability to determine the molecular structure of a protein receptor did not yet exist. Researchers at early pharmaceutical firms (often subsidiaries of chemical manufacturing firms) implemented this approach by setting up a systematic trial-and-error drug discovery system known as the 'mass screening' system, which is still used today.

The mass screening system begins with the selection or design of a 'screen'—e.g., a disease-causing bacterium or an isolated receptor that is known to be associated with the disease under study. 'Masses' of chemical compounds are then applied to this screen (one at a time), with the goal of identifying compounds that cause the screen to display a desired effect (e.g., killing of the disease-causing bacterium: evidence of binding to the receptor).

Traditionally, there have been two different sources of input materials to the mass screening process. The first source is proprietary archival libraries of known chemical compounds that have been collected by chemical and pharmaceutical firms over the years. A given major firm might have an archival library of perhaps half a million known compounds. The second source is extracts of plants, microorganisms, and animals, each of which may contain perhaps up to 100,000 unknown chemical compounds.

Mass screening proceeds differently depending on which type of input is used. In the case of archival libraries, the known compounds are tested against the disease target screen one by one, and the effect of each on the screen is observed. In the case of natural extracts, the entire extract is tested against the screen. If a desired effect is observed, the compound responsible for that effect must then be iso-

lated via a complex series of fractionations and retestings.

As an illustration for the natural compound process, consider the development of antibiotics based on ‘magainins.’ When researchers noticed that frogs living in bacteria-contaminated water did not appear to get skin infections, they suspected that a new and useful antibiotic compound in a frog’s skin might be involved (Zasloff, 1987). To identify it, they began by grinding up frog skin and subjecting the whole mixture—consisting of literally hundreds of thousands of different compounds—to mass screening tests for antibiotic activity. When these tests did indicate antibiotic activity, they next had to identify which compound(s) in the complex mixture was (were) the source of that activity. This was done by biochemical separation of the compounds found in frog skin into fractions, followed by a test of each fraction for the presence of antibiotic activity. The active fraction was then subjected to further cycles of fractionation and test until finally, the active compound was isolated.

When an active compound is finally identified via mass screening, it will generally not meet all of the criteria required to make it a ‘lead’ candidate for a new drug. For example, it may display the needed medical effect very powerfully, but may display unacceptable side effects at the same time, such as toxicity or mutagenic effects in animals, or may not become available in the bloodstream after ingestion or injection. Therefore, the lead optimization process in the drug discovery process is to create and test a number of variations (analogs) of the originally identified molecule in order to find one or more that appears to have all the attributes needed for a successful new drug. One lead compound is then advanced into the clinical development phase where its effects are tested on humans. Experimentation with analogs is carried out in series, with some elements of parallelism: a few molecules are created and tested during each round with the objective of learning as much as possible between rounds—a strategy discussed in Section 2.2.

At this point, we should note that the traditional process used to create analogs to a proposed drug in order to create a lead drug compound is typically a very costly and time-consuming matter. In order to create analogs to the original compound, medicinal

chemists (specialized organic chemists employed by pharmaceutical firms) maintain the basic structure of that compound, but add, exchange, or remove chemical groups from it. On the average, it takes 7 to 10 days and approximately US\$7500 to synthesize one such analog (Longman, 1993). According to the statistics of the Centre for Medicines Research, the average American pharmaceutical company synthesizes approximately 6100 chemical compounds for each successful drug that makes it to the market place (Halliday et al., 1992). This amounts to an average of US\$46 million for the creation of analogs alone, with a total time requirement of about 170 person/year.

The reason for the necessity to develop so many potential solutions to the receptor problem is because many drugs must be precisely tailored to sharply discriminate between very similar receptors. For example, researchers working to develop a drug for Alzheimer’s disease are targeting a particular muscarinic receptor located in the brain. However, five subtypes to this muscarinic receptor are known to exist in the gut and elsewhere, and the desired drug must *not* affect these. Compounds displaying the needed selectivity can be very difficult to find without extensive creation of analogs.

4.3. Rational drug design and combinatorial chemistry

In this century, the knowledge in chemistry, biology, and the molecular basis of disease has increased exponentially. In the beginning of the 1980s, advanced methods of protein structure determination and computer-supported molecular modeling became the focus of the pharmaceutical industry. This new technology was thought by many in the industry to be sufficiently advanced to allow the creation of a *rational drug design* methodology as an alternative to the traditional medicinal approach to drug discovery.

Pharmaceutical researchers using the rational drug design approach would use very advanced scientific methods such as X-ray crystallography and/or nuclear magnetic resonance (NMR) spectroscopy to determine the three-dimensional shape of a receptor or an enzyme that they wish to influence with a drug. They would then enter the structure of this receptor

into a computer software package containing information on the configuration and strengths of the chemical bonds that can form between atoms. This software would then allow them to use simulation to design drug molecules that bind properly to the target receptor. Real molecules would then be created by chemists in the laboratory as specified by the computer-modeling exercise, and these would be tested for the desired pharmaceutical effect. Thus, rational drug design is an example of a strategy that tries to maximize the amount of learning between trials and thereby achieve a total reduction in the number of laboratory chemical trials (as explained in Section 2.2).

However, the ‘rational’ approach to drug design has proven to be problematic for two reasons. First, the molecular modeling of a drug requires very accurate data on the structure of the target receptor, and the required degree of accuracy is often very difficult for researchers to attain using present-day methods. Second, it has been found that the shape of a target receptor can change dramatically when a drug is inserted into the receptor’s ‘binding pocket’ (see Fig. 1). The effect of such ‘induced fit’ shape changes is that a drug that has been designed to fit a receptor’s empty binding pocket may in fact not fit at all. Induced fits between receptor and drug are too complex to be modeled by present computer simulation tools. As a consequence, rational drug design has not been proven to be a full replacement for traditional drug design methods. Instead, computer-based molecular modeling exercises prescribed by the rational drug design procedure are still followed

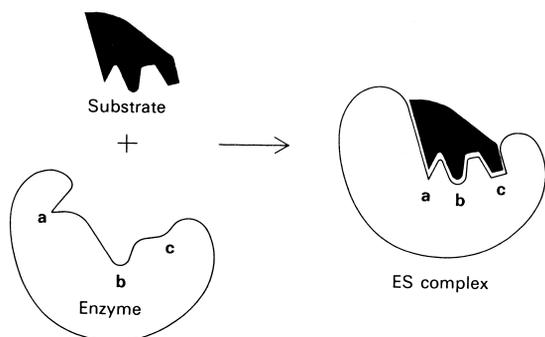


Fig. 1. Illustration of dramatic structural alterations to the binding pocket of a receptor resulting from the insertion of a drug molecule (from Stryer, 1995).

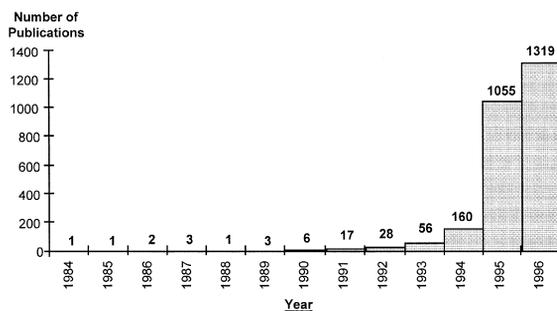


Fig. 2. The rapid increase of scientific publications on combinatorial chemistry (based on a search of the Science Citation Index, 1984–96, a database that covers 90% of the world’s significant scientific and technical literature).

up by medicinal chemists who create and test analogs to the rationally designed compound, just as was done in traditional drug development.

In the last few years, a new method called *combinatorial chemistry* has emerged very rapidly (Fig. 2), primarily due to its impact on the underlying experimentation economies (Plunkert and Ellman, 1997). Combinatorial chemistry makes the synthesis of proposed drug compounds and analogs radically faster and cheaper (the basic principle of combinatorial chemistry and the underlying process technologies are described in the Appendix A). For example, cost reductions from about US\$7500 per compound (traditional medicinal chemistry) to perhaps US\$1 to US\$10 per compound have been reported, with reductions in preparation time of comparable magnitude. While a skilled medicinal chemist required 7 to 10 days to create a single analog using traditional methods, a chemist can now use automated equipment and combinatorial chemistry techniques to create thousands of analog compounds—each precisely identified by an attached chemical ‘tag’—in a matter of days, (Franke, 1995). Thus, as we will see in Section 4.4, the dramatically different economies of combinatorial chemistry are inducing drug developers to shift to mixed experimentation strategies with a strong emphasis on parallelism, while at the same time, reducing total development time dramatically (similar to strategy (a) in Table 1).

The impact of this new capability on the drug discovery promises to be very significant. The amount of information that must be acquired about the structure of a receptor via computer modeling,

crystallographic studies, etc., can be greatly reduced with the application of combinatorial chemistry methods, because these can be used to create literally hundreds or thousands of compounds that might fit the receptor. The most promising (that is, the ones with the best desired influence on the target receptor) can then be identified via a mass-screening process. Next, one can create a new library of hundreds or thousands of analogs for each of these ‘round one winners’ within a few days or weeks. One can then repeat this screening, selection and analoging process until one gets the compounds displaying an excellent level of binding to only the target receptor. Because one has been able to create and test so many analogs, one can generate a ‘lead’ compound more quickly and more cheaply. Even more important, perhaps, is that one can identify a *better* lead compound to carry forward into the very expensive clinical development phase.

4.4. Field study: drug discovery at Pharmacoepia via combinatorial chemistry methods

We next illustrate the impact combinatorial methods can have on the economics and experimentation strategies of the drug discovery process by describing a research project carried out by Pharmacoepia of Princeton, NJ. Pharmacoepia is a well-known leader in the novel field of combinatorial chemistry using solid support libraries (see Appendix A). We compare the costs and outcomes that were achieved by using combinatorial chemistry in this case vs. estimated costs and outcomes that would have been achieved by using more traditional methods. We find that traditional methods would have been dramatically slower and costlier in this case—and would probably only have produced ‘lead’ drug candidates with little chance of clinical success. The high cost and time required to create and test compounds using traditional methods would have severely limited the number of compounds considered—and thus reduced the related search space—and would have focused the search to a region with the least promising molecules.

The drug discovery project we report upon deals with the identification of lead drug candidates to be used in the treatment of an important eye disease (glaucoma) that affects 1 in 100 adults. Glaucoma is

a wide-spread human disease responsible for impaired vision and eventual blindness. To document this project, we interviewed Pharmacoepia’s leading scientists and executives. These interviewees provided us with information about Pharmacoepia’s combinatorial chemistry technology in general and detailed information on a drug development case in particular. Personal interviews were followed with a detailed questionnaire that provided us with data on the efficiency of the drug discovery process used in this project (for detailed scientific description of the underlying chemistry used, see Burbaum et al., 1995).

Glaucoma is caused by a build-up of pressure within the human eye which in turn causes damage to optical nerve cells. Scientific research has shown that excessive pressure can be treated with the aid of what is known as the ‘diuretic’ effect (i.e., a reduction of liquid causes a decrease in pressure). It is also known that a certain group of drugs—known as carbonic anhydrase drugs—can precisely cause this diuretic effect, leading to stabilized pressure within the eye and long-term preservation of vision. The glaucoma project’s objective was to find sulfonamide compounds that ‘lock’ into and inhibit the function of the human carbonic anhydrase enzyme (hCAI) which regulates the production of liquid in

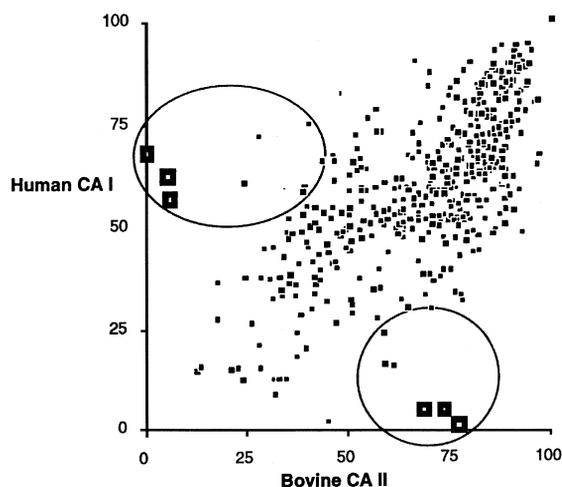


Fig. 3. Affinity of compounds for human and bovine isozyme. Each of the points represents a compound tested. The abscissa and ordinate indicate the affinity of the compounds for the human and bovine isozyme, respectively. Note that, from the large number of compounds screened, only three compounds per receptor (shown as large squares) displayed the desired discriminatory capability.

the human eye, thus leading to a reduction of both pressure and damage to optical nerve cells. As it is usually the case with pharmaceutical drugs, a promising lead compound had to discriminate against enzymes that have very similar structures in order to avoid unacceptable side effects.

In Pharmacopeia's case, a promising lead compound that might guide the way to an eventual new drug had to interact with hCAI, but discriminate against the bovine isozyme (bCAII)—two very similar receptors. (The bovine isozyme acts as a starting model for hCAII which is the human isozyme; once a drug that discriminates against bCAII is found, it acts as an excellent lead for discrimination against the human isozyme hCAII). The sulfonamide compounds identified served as leads for additional phases in the drug development process. Identifying suitable lead compounds was very difficult because, as was learned later, only 3 compounds out of thousands tested eventually displayed the searched-for selectivity (see Fig. 3; only compounds close to the abscissa and ordinate can discriminate against the respective enzymes). Failure to discriminate against

enzymes other than hCAI, however, was known to cause serious side-effects such as difficulties in breathing, convulsion, muscle cramps, and trembling.

4.5. R&D efficiency of combinatorial chemistry vs. traditional drug discovery methods

Lead compounds for the glaucoma project were identified using the combinatorial chemistry methods we described earlier. Data on development time, cost and experimentation strategies were collected for the actual mode used (combinatorial chemistry) and an estimated case using traditional medicinal chemistry was constructed (Table 3).

The estimated case was based on considerable experience with projects that were comparable in complexity and degree of difficulty but were developed using the traditional medicinal approach. A 'rational' drug design approach would have focused on reducing the number of compounds tested in the traditional approach by maximizing the learning between successive rounds—but at considerable addi-

Table 3

Comparing the parallel combinatorial chemistry approach with serial traditional medicinal chemistry in the discovery of promising lead compounds for the treatment of glaucoma^a

Project variable	Combinatorial approach	Traditional approach ^b
(1) Total development time	3.5 months	5 years ^c
(2) No. of chemists needed	4	15 ^c
(3) No. of compounds tested	~ 9000	~ 3750 ^d
(5) No. of (serial) rounds	1 ^e	100 (250 max.) ^f
(6) No. of compounds per round	~ 9000	~ 38
(7) Cost of screen per round	US\$10,000	US\$10,000
(8) Total cost (chemists only) ^g	US\$167,000	US\$18.75 million
(9) Total cost per compound	US\$19 ^h	US\$5000

^bBased on estimates from developers who are very experienced with medicinal drug development.

^cTypical time and resources planned for a project of the complexity and strategic importance equal to that of the combinatorial chemistry approach reported in this table.

^dA skilled chemist can prepare 50 compounds per year.

^eDuring a short second round, 220 compounds were prepared over a 2-week period. The compounds were a subset of the first round and did not contain new members.

^fWhile 250 rounds is theoretically possible, it does not allow sufficient time for learning and analysis between rounds. Thus, 100 rounds is a realistic number.

^gApproximately US\$250,000 per year is needed to pay a skilled chemist. In the combinatorial approach, chemists were only needed for 2 months.

^hSince the marginal cost of preparing additional compounds using combinatorial chemistry is negligible, a 10-fold increase in the number of compounds prepared would result in a per unit cost of approximately US\$2.

^aIn the case of both methods the project was followed by further refinements of the lead candidates in order to increase the probability of clinical success.

tional cost and time. Thus, it is unclear whether current ‘rational’ drug design would have improved the efficiency and output of the traditional approach at all. The data from Table 3 show not only that the combinatorial approach was more cost effective but the parallel nature of experimentation also led to a dramatically lower discovery lead time, allowing the firm to move to the clinical phase much earlier than the more serial traditional approach. It also identified lead compounds that showed a high degree of selectivity and thus much promise of success for the next development phases. In fact, interviewees strongly felt that with the cost and time required using traditional medicinal methods, (1) it would have been unlikely that the project had been pursued; and if pursued, (2) it was nearly impossible to identify a promising lead compound with the required selectivity.

Of course, combinatorial chemistry does not offer the same advantage for all projects, and is currently not applicable at all to some kinds of molecules. Thus, the method is not now very effective for the kinds of complex molecules often dealt with in studies of natural compounds. However, combinatorial chemistry’s area of applicability is rapidly expanding and many companies are working on the conversion of classical organic chemical reactions to combinatorial systems. Expert interviewees contacted during our case study estimate that the advantages offered by combinatorial methods over traditional experimental methods, for projects where combinatorial chemistry is applicable, range from a 10% to an 80% reduction in the cost and time devoted to lead optimization—and, as was noted earlier, the development of better quality lead compounds than is customarily accomplished by traditional medicinal chemistry techniques.

5. Conclusion

In this paper, we have argued that the economics of problem-solving and the related R&D efficiency are being radically affected by the use of new and greatly improved versions of methods such as computer simulation and combinatorial chemistry. We explained how the introduction of novel methods could affect both the experimentation strategies

adopted by firms (e.g., serial vs. parallel experimentation) and the efficiency with which those strategies can be executed. Via a field study of the impact of combinatorial chemistry techniques on pharmaceutical drug discovery, we then illustrated the dramatic economic changes that can result from the adoption of novel experimental methods.

We also noted that novel experimental methods can importantly affect the relative competitive position of firms if techniques that offer such advantages are difficult or impossible to buy or sell in the available factor markets, and difficult to replicate as well. Many novel methods require novel skills and/or organizational arrangements to implement, and are likely to meet these criteria. (Certainly, as we explained, it is likely that the methods discussed in our case, combinatorial chemistry and rational drug design, meet the criteria. Equipment required by both are available on the market—but both also require novel skills and organizational arrangements that are not easily acquired by firms seeking to adopt them.)

In sum, we propose that strategies and modes of experimentation can be an important factor in the effectiveness of a firm’s innovation processes and its relative competitive position. We therefore propose that further studies on this topic may be of interest to both innovation researchers and innovation practitioners. For example, it would be useful to explore whether and how differences in the relative effectiveness of firm innovation processes can be traced to differences in the experimental methods employed, and the skill with which those methods are used. It would also be useful to explore the attributes of experimental methods that convey the greatest competitive advantage to firms using them. For example, it is likely that the methods that are the hardest to transfer to new users will be the ones that offer the greatest competitive advantage to method *users*—while method *sellers* are likely to appropriate the most benefit from methods that are easily transferred.

Acknowledgements

We would like to thank John Baldwin, John Chabala, Joseph Mollica and Nolan Sigal from Pharma-

copeia, Michael Pavia from Millennium Pharmaceuticals, our colleagues Carliss Baldwin, David Bell, Steven Wheelwright and two anonymous reviewers for their help and comments on earlier drafts of this paper. The financial support of the Harvard Business School Division of Research and the Research on Science and Technology Program at the National Science Foundation are gratefully acknowledged.

Appendix A. Background on combinatorial chemistry

Combinatorial chemistry is a very novel experimentation methodology and has evolved over the last decade. It consists of several new chemical synthesis strategies for the efficient generation of a large number of chemical compounds. This large number of chemical compounds, also called ‘compound libraries,’ is subsequently used in pharmaceutical drug screening projects. The term ‘combinatorial’ originates from chemical synthesis methods applied to most of these libraries.

The following is a brief description of the main process technologies that enable a large number of parallel experiments to be generated quickly and at low cost.

- *Biochip libraries.* Photolithographic synthesis methods are used for the creation of compound libraries on the surface of a silicon chip (Fodor et al., 1991). Up to 10,000 individual compounds can be synthesized on a silicon chip with a little more than 1 cm² surface area.

- *Solid support libraries.* Compounds are synthesized on the surface of polymer beads. This method

allows the chemist to attach a certain type of molecule to glass beads and split the pool of glass beads to continue with different synthesis steps (see example below).

- *Solution libraries.* Mixtures of compounds react chemically in a carefully designed system to form solution libraries with tens to thousands of different compounds within a few hours.

- *Rapid parallel synthesis libraries.* Robotic equipment is custom tailored to dispense chemicals into individual reaction chambers, carry out many individual chemical reactions in parallel, and extract and purify the reaction products automatically. Although this process is significantly slower than the other three technologies, it results to individually purified compounds at quantities sufficiently large for elaborate second round screening. (The other methods require chemical resynthesis which may cause a small but significant time delay.)

To illustrate how such a combinatorial chemistry works, consider the process of building solid support libraries. In the first synthesis phase, polymer beads are reacted in three different reaction vessels with chemical A in vessel 1, chemical B in vessel 2 and chemical C in vessel 3. After the reactions are completed, all the polymer beads are pooled and mixed. The mixture is now split into 3 equal portions and placed in vessels 1, 2 and 3. *Each* vessel now contains three mixtures: polymer beads covered with A, B and C.

In the second phase, Vessel 1 is reacted with chemical D, vessel 2 is reacted with chemical E and vessel 3 is reacted with chemical F. The result of this reaction is as follows:

-vessel 1 contains polymer beads carrying A–D, B–D, and C–D,

Table 4
Building solid support libraries

Phase (<i>k</i>)	Vessel 1	Vessel 2	Vessel 3	<i>N</i>
1	A	B	C	3
2	A–D, B–D, C–D	A–E, B–E, C–E	A–F, B–F, C–F	9
3	A–D–G, B–D–G, C–D–G	A–D–H, B–D–H, C–D–H	A–D–I, B–D–I, C–D–I	27
	A–E–G, B–E–G, C–E–G	A–E–H, B–E–H, C–E–H	A–E–I, B–E–I, C–E–I	
	A–F–G, B–F–G, C–F–G	A–F–H, B–F–H, C–F–H	A–F–I, B–F–I, C–F–I	
<i>k</i>	–	–	–	3 ^{<i>k</i>}

At the beginning of each round, a new reagent is introduced to each vessel. For example, A is added to vessel 1 in round 1, D is added to vessel 1 in round 2, G is added to vessel 1 in round 3, etc.

-vessel 2 contains polymer beads carrying A–E, B–E, and C–E, and

-vessel 3 contains polymer beads carrying A–F, B–F, and C–F.

The output of the second synthesis phase is a diversity of $3^2 = 9$ compounds. Without the splitting and mixing of the polymer beads after the first synthesis round and the combining of the individual pools, the second synthesis round would have yielded only 3 compounds. Combinatorial chemistry can increase the number of compounds by the power of the synthesis phases, resulting in a large chemical diversity very quickly. For instance the third round of the example given above would result in $3^3 = 27$ compounds (see Table 4).

Chip technology and solution libraries are following different chemistries but have similar exponential increases in compounds generated with each synthesis round. Rapid parallel synthesis, however, achieves its efficiency gain through robot technology and parallel execution of many individual reactions. However, the efficiency gain is not exponential with the synthesis phase.

References

- Abernathy, W., Rosenbloom, R., 1968. Parallel and sequential R&D strategies: application of a simple model. *IEEE Trans. Eng. Manage.* EM 15 (1).
- Alchian, A., 1950. Uncertainty, evolution and economic theory. *J. Pol. Econ.* 58 (3).
- Alexander, C., 1964. *Notes on the Synthesis of Form*. Harvard Univ. Press, Cambridge, MA.
- Allen, T.J., 1966. Studies of the problem-solving process in engineering design. *IEEE Trans. Eng. Manage.* EM 13 (2), 72–83.
- Arora, A., Gambardella, A., 1994. The changing technology of technological change: general and abstract knowledge and the division of innovative labor. *Res. Policy* 23 (55), 523–532.
- Baldwin, C., Clark, K., 1997a. Design options and design evolution, Working Paper #97-038. Harvard Business School.
- Baldwin, C., Clark, K., 1997b. The value of modularity: splitting and substitution, Working Paper #97-039. Harvard Business School.
- Barley, S., Bechky, B., 1994. In the backrooms of science: the work of technicians in science labs. *Work and Occupations* 21 (1), 85–126.
- Barney, J., 1986. Strategy factor market: expectation, luck and business strategy. *Manage. Sci.* 32 (10), 1231–1241.
- Barron, J., 1988. *Thinking and Deciding*. Cambridge Univ. Press.
- Bruderer, E., Singh, J., 1996. Organizational evolution, learning, and selection: a genetic-algorithm-based model. *Acad. Manage. J.* 39, 1322–1349.
- Burbaum, J. et al., 1995. A paradigm for drug discovery employing encoded combinatorial chemistry. *Proc. Natl. Acad. Sci. USA* 92, 6027–6031.
- Carter, A., 1995. Measuring the performance of the knowledge-based economy, Working Paper. Brandeis University.
- Clark, K., Fujimoto, T., 1991. *Product Development Performance*. Harvard Business School Press, Boston.
- Cohen, W., Levinthal, D., 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Sci. Q.* 35 (1), 128–152.
- DiMassi, J., Hansen, R., Grabowsky, H., Lasagna, L., 1991. Cost of innovation in the pharmaceutical industry. *J. Health Econ.* 10, 107–142.
- DiMassi, J., Seibring, M., Lasagna, L., 1994. New drug development in the United States from 1963–1994. *Clin. Pharmacol. Ther.* 55, 609–622.
- Duncker, K., 1945. On problem solving. *Psychol. Monographs* 58 (5), translated by Lynne S. Lees.
- Evenson, R., Kislav, Y., 1976. A stochastic model of applied research. *J. Pol. Econ.* 84 (2).
- Fodor, S., Read, J., Pirrung, M., Stryer, L., Lu, A., Solas, D., Solas, A., 1991. Light-directed, spatially-addressable, parallel chemical synthesis. *Science* 251, 767–773.
- Franke, R., 1995. The impact of combinatorial chemistry on the pharmaceutical drug discovery process, Unpublished Master's Thesis. MIT Sloan School of Management, Cambridge, MA.
- Gerald, C., Wheatley, P., 1984. *Applied Numerical Analysis*. Addison-Wesley, Reading, MA.
- Griliches, Z., 1957. Hybrid corn: an exploration in the economics of technical change. *Econometrica* 25 (4), 501–522.
- Henderson, R., 1994. The evolution of integrative capability: innovation in cardiovascular drug discovery. *Industrial and Corporate Change* 3 (3).
- Henderson, R., Cockburn, I., 1994. Measuring competence: exploring firm effects in pharmaceutical research. *Strategic Manage. J.* 15, 63–84.
- Halliday, R., Walker, S., Lumley, C., 1992. R&D philosophy and management in the world's leading pharmaceutical companies. *J. Pharm. Med.* 2, 139–154.
- Holmstrom, B., Tirole, J., 1991. The theory of the firm. In: Schmalensee, R., Willig, R. (Eds.), *Handbook of Industrial Organization*. North Holland, Amsterdam.
- Iansiti, M., 1997. *Technology Integration: Making Critical Choices in a Turbulent World*. Harvard Business School Press, Boston.
- Kauffman, S., Levin, S., 1987. Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* 128, 11–45.
- Levinthal, D., 1997. Adaptation on rugged landscapes. *Manage. Sci.* 43 (7).
- Longman, R., 1993. *The Promise of Combinatorial Chemistry*. In: Windhover Information, *Pharmaceutical Strategic Alliances*, 5th edn., January 1993–June 1994.
- Mansfield, E., 1968. *Industrial Research and Technological Innovation: An Econometric Analysis*. W.W. Norton, New York.
- Marples, D., 1961. The decisions of engineering design, *IRE Trans. Eng. Manage.*, June 55–71.

- Milgrom, P., Roberts, J., 1990. The economics of modern manufacturing: technology, strategy and organization, *American Econ. Rev.*, pp. 511–528.
- Montgomery, D., 1991. *Design and Analysis of Experiments*. Wiley, New York.
- Morison, E., 1966. *Men, Machines, and Modern Times*. MIT Press, Cambridge, MA.
- Nelson, R., 1961. Uncertainty, learning, and the economics of parallel research and development efforts. *Rev. Econ. Stat.* 43, 351–364.
- Nelson, R., 1982. The role of knowledge in R&D efficiency. *Q. J. Econ.*, August 1982.
- Nelson, R., Winter, S., 1982. *An Evolutionary Theory of Economic Change*. The Belknap Press of Harvard University Press, Cambridge, MA.
- Pavitt, K., 1987. The objectives of technology policy. *Sci. Public Policy* 14 (4), 182–188.
- Plunkert, M., Ellman, J., 1997. Combinatorial chemistry and new drugs. *Scientific American*, April 1997.
- Polanyi, M., 1958. *Personal Knowledge: Towards a Post-Critical Philosophy*. Univ. of Chicago Press, Chicago, IL.
- Pople, H., 1982. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In: Szolovits, P. (Ed.), *Artif. Intell. Med.*, Chap. 5. Westview Press, Boulder, CO.
- Reitman, W.R., 1965. *Cognition and Thought*. Wiley, New York.
- Roberts, K., Weitzman, M., 1981. Funding criteria for research, development, and exploration projects. *Econometrica* 49 (5), 1261–1288.
- Rosenberg, N., 1982. *Inside the Black Box: Technology and Economics*. Cambridge Univ. Press, New York.
- Sachs, E. et al., 1992. CAD-casting: direct fabrication of ceramic shells and cores by three-dimensional printing. *Manufacturing Rev.* 5 (2), 117–126.
- Savello, D., 1996. Speeding new medicines to patients. *Drug Discovery Today* 1 (2), 45.
- Schön, D., 1967. *Technology and Change*. Delacorte Press, New York.
- Simon, H.A., [1969] 1981. *The Sciences of the Artificial*, 2nd edn. MIT Press, Cambridge, MA.
- Simon, H.A., 1973. The structure of ill-structured problems. *Artif. Intell.* 4, 181–201.
- Smith, R., Eppinger, S., 1997. A prediction model of sequential iteration in engineering design. *Manage. Sci.* 43 (8).
- Stryer, L., 1995. *Biochemistry*, 4th edn. Freeman, New York.
- Teece, D., 1977. Technology transfer by multinational firms: the resource cost of transferring technological know-how. *Econ. J.* 87 (346), 242–261.
- Thomke, S., 1997. Managing experimentation in the design of new products. Working Paper #96-037, Harvard Business School, *Manage. Sci.*, in press.
- Thomke, S., 1998. Simulation, learning and R&D performance: evidence from automotive development, Working Paper #97-089, Harvard Business School, *Res. Policy*, in press.
- Tushman, M., O'Reilly, C. III, 1996. The ambidextrous organization: managing evolutionary and revolutionary change. *California Manage. Rev.* 38 (4).
- Tyre, M.J., von Hippel, E., 1997. The situated nature of adaptive learning in organizations. *Organization Sci.* 8 (1), 71–83.
- Villasenor, J., Mangione-Smith, W., 1997. Configurable computing. *Scientific American*, June, pp. 66–71.
- von Hippel, E., 1994. Sticky information and the locus of problem-solving: implications for innovation. *Manage. Sci.* 40 (4), 429–439.
- von Hippel, E., Tyre, M.J., 1995. How learning by doing is done: problem identification in novel process equipment, *Res. Policy* 24, 1–12.
- Ward, A., Liker, J., Cristiano, J., Sobek, D., 1995. The second Toyota paradox: how delaying decisions can make better cars faster. *Sloan Manage. Rev.*, Spring 1995.
- Weitzman, M., 1979. Optimal search for the best alternative. *Econometrica* 47 (3), 641–654.
- Wernerfelt, B., 1984. A resource-based view of the firm. *Strategic Manage. J.* 5, 171–180.
- Wheelwright, S., Clark, K., 1992. *Revolutionizing Product Development*. The Free Press, New York.
- Wright, S., 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proc. 11th Int. Congr. Genet.* 1, 356–366.
- Zaslloff, M., 1987. *Proc. Natl. Acad. Sci. USA* 84, 5443–5453.