

Managing a Portfolio of Interdependent New Product Candidates in the Pharmaceutical Industry

Gary E. Blau, Joseph F. Pekny, Vishal A. Varma, and Paul R. Bunch*

Highly regulated industries such as pharmaceuticals and agrochemicals face the challenge of maintaining a continuous stream of new products. This is difficult because of low probabilities of technical success, high development costs, uncertain market impact, a scarcity of good new product ideas, and limited human and capital resources available to develop them. The problem of evaluating and selecting which new products to develop and then of sequencing or of scheduling them is complicated further by the presence of dependencies between products both in the market place and in the development process itself.

This study proposes a portfolio management approach that selects a sequence of projects, which maximizes the expected economic returns at an acceptable level of risk for a given level of resources in a new product development pipeline. A probabilistic network model of distinct activities is used to capture all the activities and resources required in the "process" of developing a new drug. A prioritization scheme suggesting sequences for developing new independent drug candidates with unlimited resources is generated with a conventional bubble chart approach. These sequences initiate a genetic algorithm (GA)-based search for the optimal sequence in the presence of product dependencies and limited resources. By statistically evaluating the sequences generated during the GA search using a discrete event simulation model, it is possible to construct an economic reward-risk frontier that illustrates the trade-offs between expected rewards and risks. The model ideally is suited to answer various "what if" questions relative to changes in the resource level on pipeline performance.

The methodology is illustrated with an industrially motivated case study, involving nine interdependent new product candidates targeting three diseases. The dramatic results yield a candidate sequence with an expected return 28 percent higher than the sequence suggested by the bubble chart approach at almost the same level of risk. The synergism among the candidate dependencies, pipeline resources, and economic and technical uncertainties demonstrates the necessity of a computationally intensive approach if the best development strategy is to be realized.

Address correspondence to Gary Blau, Department of Chemical Engineering, Purdue University, West Lafayette, IN 47908. Tel: (765) 494-9472. Email: blau@purdue.edu.

*The authors wish to thank Karthik Rajan for his work on this project and Dow Agrosciences, LLC for their financial and technical support.

Introduction

Research shows that only around 60 percent of new products launched are successful commercially (Griffin, 1997). In addition, for every seven concepts that enter the new product development (NPD) process, on average only one is

successful commercially. Consequently, almost half of the resources that U.S. industry devotes to NPD are spent on products that fail or are cancelled. In the pharmaceutical industry this problem particularly is challenging because of long development times, low success rates, high capital requirements for building a manufacturing facility, and broad uncertainty in sales estimates (Kahn, 2002). This article will provide a computational decision support system for managing a new product pharmaceutical portfolio. The goal is to maximize expected financial returns at an acceptable level of risk for a given level of corporate resources.

In the spirit of Adler et al. (1995) and Pisano (1997), this study's approach uses models similar to manufacturing to characterize the highly regulated product development process in the U.S. pharmaceutical and agrochemical industries. To gain regulatory approval from government agencies, all new products follow the same route through the development process. The plan for developing each drug mirrors a

house-building plan with definitive precedence relationships and task durations. A major departure from the house analogy is that during, or at the end, of some of these tasks, the development of the new drug may be terminated. These termination points may occur because of unwanted side effects, marginal efficacy, or competition from in-house or competitor candidates. Once estimates of the termination probabilities and task duration distributions are supplied, it is possible to simulate the movement of new product candidates through the pipeline by using discrete event simulation techniques (Adler, 1995; Taylor and Moore, 1980). One additional complication is the need to consider the interdependent nature of the various drug candidates that arises from competition from both in-house and competitor products (Roberts, 1999), synergism between the duration of tasks of related products in the pipeline, and finally dependencies of technical success probabilities of other drug candidates preceding them in the pipeline. Incorporating these dependencies into the model of the pipeline represents a significant advance over the existing literature.

BIOGRAPHICAL SKETCHES

Gary E. Blau currently holds the position of visiting industrial professor of chemical engineering at Purdue University. He joined the faculty in 1998 after 30 years' experience with Dow Chemical and Dow AgroSciences. He has an active research program in developing methodologies to facilitate the adaptation of quantitative decision-making tools for managing risk in research and development (R&D) project selection and portfolio optimization. He is also president of The Modeling Connection, LLC, a consulting firm providing mathematical modeling expertise to industry. Dr. Blau has a B.A.Sc. degree from the University of Waterloo and a M.S. and Ph.D. from Stanford University, all in Chemical Engineering.

Joseph F. Pekny joined the Purdue University faculty in 1990 and is professor of chemical engineering with research interests in supply chain management, planning and scheduling systems, pharmaceutical pipeline management, model-based and data driven management, and real-time decision systems. His research group is active in risk management, decision support systems, algorithm engineering, combinatorial optimization, large-scale software systems, sensitivity analysis, and simulation based optimization architectures. He received his Ph.D. from Carnegie Mellon University and his B.S. from Princeton University, both in chemical engineering.

Vishal A. Varma is a graduate student in the Department of Chemical Engineering at Purdue University, pursuing a Ph.D. in chemical engineering. His major area of research includes stochastic modeling and optimization of bio-pharmaceutical new product development pipelines and other business systems related to the process industry.

Paul R. Bunch is currently manager of capacity planning and management at Lilly Research Laboratories, located in Indianapolis, Indiana. Dr. Bunch develops tools and methods for portfolio and capacity management of the drug discovery and development process at Eli Lilly and Company. He has held positions in manufacturing, engineering, project management, and R&D strategy. Dr. Bunch received both his B.S. and Ph.D. in chemical engineering from Purdue University.

Evolution of New Products in the Pharmaceutical Industry

There are three major stages in the life cycle of a pharmaceutical product:

Discovery → Development → Launch (Commercial)

Discovery Stage

Literally thousands of molecules are applied to targets developed to simulate various disease groups. Once an active molecule—one that has an effect on the target—is discovered, various permutations of the structure of the molecule are tested to improve various properties that predict whether the molecule eventually will be successful as a drug. The best molecule from these structure-activity relationships is tested for toxicological results in whole-body systems. If no particularly worrisome toxic effects are observed, the molecule is promoted to the status of "lead" molecule and becomes a candidate for development.

Development Stage

Enormous sums of money and resources are committed to the lead molecule to observe its behavior in

healthy volunteers first, then in patients with the disease, and finally in large-scale clinical studies conducted in concert with the Food and Drug Administration (FDA). Coincident with these studies, process research and formulations work is conducted to supply the drug for testing purposes as well as to design and to construct a commercial plant, if the product is launched.

Commercial Stage

If the drug is effective in clinical studies, has no unacceptable side effects, and is approved by the FDA, it may be marketed. Target markets are identified for a staged launch or “ramp-up” of the new compound. After a few years, a mature sales level usually is reached and is maintained until patent coverage on the molecule expires and/or until competition from generics is realized (Roberts, 1999). Once generics are available, an attempt usually is made to get approval of the drug for alternative markets and perhaps in different dosage forms. In any case, sales diminish after a patent expires.

This study is concerned with the decision-making process at the start of the development phase. New drug candidates are available from the discovery process, and management needs to decide which, if any, of these to select. Once selected, they need to decide how to sequence them with respect to each other and other drug candidates that can be supplied from other drug discovery companies through some type of in-licensing arrangement. The word *sequencing* means that when conflicts arise, due to resource constraints, for example, the earlier appearing candidate is processed first.

Details of the Drug Development/ Commercialization Process

Once a new candidate molecule is promoted to the status of lead molecule, it follows the steps shown in the flow diagram of Figure 1. All activities or tasks are described by rectangles and the precedence relationships between tasks by arrows. Key termination points are presented as diamonds and are not truly activities but rather are the consequence of the preceding activities. The entire flow diagram is referred to as a probabilistic network model to distinguish it from the traditional critical path method (CPM) or project evaluation and review technique (PERT) representation (Elmaghraby, 1977).

First Human Dose Preparation (FHD)

This is a planning activity in preparation for administration to healthy volunteers. This includes various pharmaco-kinetic studies including adsorption, distribution, metabolism, and excretion from the body as well as suitable dose levels.

Phase I

This includes the first clinical trials in which the drug is administered to healthy human volunteers. Acute/ chronic and reproductive studies in animals (mice/ rats) also are conducted simultaneously. Positive results indicate acceptable absorption, distribution, or elimination patterns. Unacceptable behavior in human and animal studies can terminate the study.

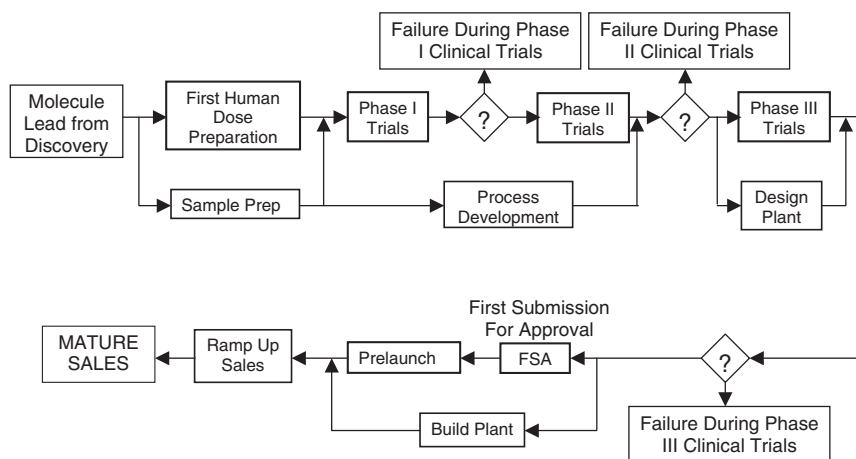


Figure 1. Flow Diagram of the Activities Involved in the Development and Commercialization of a New Drug Candidate

Phase II

The compound is administered to human patients with the disease by using the results of dosing studies from Phase I. Coincident with these studies are long-term oncogenic toxicological studies in animals and market research to obtain sales estimates. If the compound fails to treat the disease or is inferior to competitive products, it is destaged or is returned to the discovery phase for modification.

Phase III

This phase includes large-scale clinical studies on humans with the disease. The FDA is involved and indicates benchmarks for giving their approval. In addition to confirming the efficacy, these studies identify drug–drug interactions, human demographics, and so forth. This most expensive phase of the development process requires extensive global coordination and cooperation. The results should confirm what was learned in Phase II but on a much larger scale; otherwise, the compound may be terminated.

First Submission for Approval (FSA)

All information (e.g., efficacy, toxicology, process, drug–drug interactions, side effects) is combined and is submitted to the FDA. Simultaneously, the marketing strategy is evolving, price negotiations are being conducted with suppliers/distributors, and promotional materials are being developed. The construction of a commercial plant is in progress. Approval for selling the new drug is the anticipated outcome.

Prelaunch Activities

This is the final stage before launch: Approval has been received from the FDA; a global penetration strategy has been completed; the commercial plant has been built and has been started up; and a promotional campaign has been launched. This phase ends when the new drug is distributed.

Launch Activities

The compound is launched over a period of years in various global markets until mature sales levels are

realized—the ramp-up period. Mature sales are maintained until patents expire or until competition is realized either from competitors or from planned cannibalization.

Product Supply Chain Activities

Sample preparation, process research, process development, process design, and plant construction occur simultaneously with other development activities (Figure 1). Initial focus is on preparing sufficient sample material for various animal/human studies. Once the launch prospects appear promising, the emphasis changes to developing a process for commercialization. This includes a pilot plant that provides data for plant design as well as the larger quantities of material needed for Phase III clinical trials. During Phase III clinical trials, the new plant is designed, or other arrangements for product manufacture are made. Once phase III trials are successful, plans are launched to build a new plant, to modify existing facilities, or to find a suitable source of the materials needed for launch.

The Portfolio Management Problem

Managing the new product pipeline is a series of trade-offs among maximizing expected economic returns, minimizing risk, and maintaining diversity in the product mix for a given level of renewable and nonrenewable corporate resources. Uncertainty and resource limitations make this a challenging problem. For example, there is a very high attrition rate as candidates move through the pipeline. Twenty percent of the candidates drop out after Phase I, while 80 percent of the remaining candidates drop out after Phase II human efficacy testing (Rodriguez, 1998). As the drug moves down the pipeline from Phase I to III, the financial costs and resources required for testing dramatically increase. Investments close to \$500MM may be required just for the opportunity to launch the drug.

The portfolio must be selected in such a way that the competition among drug candidates for limited resources does not result in unusually long average product development times and hence late commercialization (Gupta and Wilemon, 1996). Further, the portfolio must be large enough to compensate for product failures yet must not be too large to overextend resources. The problem is complicated further by the presence of product dependencies. Much like the

exploitation of negative correlations between individual asset returns in managing a portfolio of financial assets (Markowitz, 1991), the structure of correlations and dependencies among development activity durations, capital investments, resource requirements, clinical success probabilities, and marketplace performance can be used in managing a portfolio of new drugs. However, unlike financial portfolio management, exploiting product dependencies requires prudent selection of drug candidates within a product portfolio and the sequence in which the drugs are developed.

We propose a rational and computationally feasible approach that incorporates these complicating uncertainties, product dependencies, and combinatorial decisions into a single framework. This study's overall framework is built upon a unique combination of discrete event simulation (Law and Kelton, 2000) and a specially designed genetic algorithm (GA) (Holland, 1975). The essence of this approach lies in capturing the probabilistic behavior of the drug product pipeline within the discrete event simulation model, under selection and sequencing decisions imposed by the GA. The GA makes various feasible drug selection and sequencing decisions that are evaluated by using the discrete event simulator. The probability distributions generated by the discrete event simulator then are evaluated to determine an optimized decision.

Review of Literature on Pharmaceutical Portfolio Management

The earliest portfolio management techniques applied in the pharmaceutical industry were based on economic analysis (Chapman and Ward, 1996). The discounted cash flow (DCF) method remains the most commonly used technique (Krishnan and Ulrich, 2001). However, it is based on expected values of uncertain parameters and is unable to generate quantitative details about the risk associated with a given drug candidate (Poh, Ang, and Bai, 2001). Economic analysis methods also have been criticized for their rigid focus on single criteria decision-making versus more realistic multiple criteria decision-making (Linton, Walsh, and Morabito, 2002).

These concerns have been addressed by decision theoretic methods (Morgan and Henrion, 1990). Decision theory formalizes the key concepts of risk and return by defining the decision-maker's utility function (Markowitz, 1991). Using this formalism, decision theory provides comprehensive portfolio management

methods such as decision trees that allow management to undertake complex resource allocation decisions among competing drug candidates with full consideration to the possibilities of drug failures (Sharpe and Keelin, 1994). The decision tree method also has addressed portfolio management issues such as how many projects to pursue and how many projects to terminate (Ding and Eliashberg, 2002). Unfortunately, decision tree methods have come under criticism for occurrence of unmanageably large decision trees, due to significantly rapid increase in the number of selection and sequencing decisions with the size of the portfolio (Copeland and Antikarov, 2001).

Real option valuation (ROV) has been promoted as a substitute to the decision tree method (Amram and Kulatilaka, 1999). While there is a large literature on pharmaceutical portfolio management using ROV (e.g., Loch and Bode-Greuel, 2001), in practice the method has been used effectively only to evaluate single projects (Copeland and Antikarov, 2001). Other methods such as the stage-gate™ process (O'Connor, 1994) appear mainly focused on tactical decisions such as regulating the flow of work in the pipeline rather than on strategic decisions such as project selection and sequencing.

In contrast to this, Monte Carlo/discrete event simulation methods (Law and Kelton, 2000) readily accommodate uncertainties and alternative performance criteria in a computationally feasible manner (Blau and Bunch, 2002). Simulation-based methods have been employed by Adler et al. (1995) to analyze a relatively complex engineering design process, by Blau et al. (2000) in simulating an industrial-scale pharmaceutical NPD pipeline, by Repenning (2001) in modeling the control and dynamics of a two-stage NPD system, and by Subramanian, Pekny, and Reklaitis (2003) in studying the effect of activity rescheduling on portfolio performance. However, to date no simulation-based method optimizes project selection and sequencing decisions jointly. This study will address project selection and sequencing decisions for the product development pipeline against product launch decisions (Hultink et al., 1997), which emphasize such issues as market introduction and advertising campaigns.

Modeling Uncertainties in Pharmaceutical Research and Development

Any attempts at detailed product portfolio analysis must begin by characterizing all uncertainties

Table 1. Probability Distributions for the NPD Pipeline Activities for Figure 1

Activity	Duration (days)			Cost (\$MM)			Total Available Resource (\$MM)
	Min	ML	Max	Min	ML	Max	
FHD Prep	300	400	500	72	80	88	275
Phase I	225	300	375	70	80	90	350
Phase II	375	500	625	75	80	85	175
Phase III	575	775	975	150	200	250	250
FSA	275	375	475	18	20	22	100
Prelaunch	75	100	125	45	50	55	550
Ramp Up 1 (RU1)	250	350	450	9	12	15	25
Ramp Up 2 (RU2)	250	350	450	19	22	25	50
Ramp Up 3 (RU3)	250	350	450	35	40	45	100
Mature Sales (Mat Sales)	250	350	450	46	53	60	150
Sample Prep	300	400	500	1.8	2	2.2	10
Process Development 1 (ProcDev1)	600	800	1000	7	10	13	16
Process Development 2 (ProcDev2)	600	800	1000	7	10	13	16
Design Plant (Des Plant)	550	750	950	8	10	12	12
Build Plant	600	750	900	52	62	72	120

associated with parameters of all development activities, manufacturing costs, and marketplace returns with probability distributions. For the flow diagram shown in Figure 1, historical data at a major pharmaceutical company were used to represent the parameters with a triangular distribution. The data are presented in Table 1. For example, the time required for Phase I testing ranges from minimum (min) of 225 days to a maximum (max) of 375 days with a most likely (ML) value of 300 days. Costs are not distributed between manpower and equipment/clinical costs. This type of detailed physical resource-based data is available and could be used for future resource planning but is beyond the scope of this article.

There is a relationship between activity times and costs in Table 1 for specific drug candidates. For example, new drugs from a class of chemistries or biotech products would require activity resource levels closer to the maximum of the triangular time and cost distributions than those familiar to a company. This

relationship is captured with a simple parameter called the degree of difficulty (DoD). Subjective estimates of DoD can be obtained from the various principal investigators, although the values may be different between work processes. However, since the focus of this article is on project selection and sequencing rather than resource planning, the analysis can be simplified by using a single value of DoD ranging from 1 (very easy) to 10 (very difficult). Table 2 lists DoD for a set of new product candidates. The reported DoD values are used as follows: (1) The minimum and maximum of the triangular distribution remain the same as the values shown in Table 1 for all the drug candidates; while (2) the most likely value of the distribution is proportional to DoD. If DoD is one, for example, the most likely value is set equal to the minimum of the triangular distribution while the maximum remains the same. Conversely, if DoD is 10 the most likely value is set to the maximum while the minimum remains the same.

Table 2. Example Data for Nine New Drug Candidates

Drug Name	Disease Type	Success Probabilities			Capital Cost (MMS)			Mature Sales (MMS)			Degree of Difficulty
		Phase1	Phase2	Phase3	Min	ML	Max	Min	ML	Max	
D1	III	0.9	0.3	0.9	40	50	60	800	900	1000	5
D2	I	0.85	0.2	0.85	20	30	40	400	500	900	2
D3	I	0.95	0.35	0.95	30	45	60	600	750	850	8
D4	II	0.87	0.22	0.80	28	34	40	700	1000	1200	9
D5	II	0.97	0.36	0.99	25	40	75	1000	1200	1600	3
D6	I	0.83	0.18	0.86	50	60	70	500	650	800	7
D7	I	0.94	0.4	0.94	65	75	90	400	475	800	1
D8	II	0.86	0.2	0.88	60	65	90	1000	1500	1700	4
D9	II	0.98	0.34	0.92	52	62	72	1000	1200	1300	10

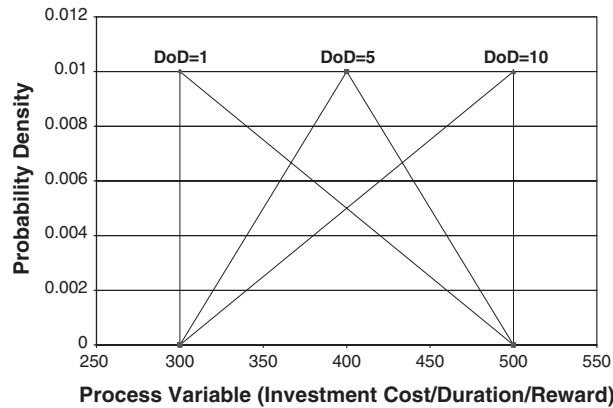


Figure 2. Impact of Degree of Difficulty (DoD) on Triangular Distributions

For intermediate values of DoD, the most likely values are linearly scaled by their values above or below DoD equal to five, the value used to generate the most likely values in Table 1. This scaling procedure is shown in Figure 2. More sophisticated scaling is possible but probably is not justified because of the extent of the other uncertainties involved.

Although there is considerable uncertainty in specifying the resource requirements for individual activities, there is little or no uncertainty in specifying the “total” level of resources available to the decision-maker at any point in time in the pipeline. The resource availabilities for each activity are given in the last column of Table 1. The types of complex interaction between skill sets of the various teams, as considered by Adler et al. (1995), are less important in the drug development environment where the process follows a well-defined, predetermined path.

The remainder of the uncertainties in the NPD process is associated with the individual drug candidates themselves. These include the success probabilities in the three phases of clinical trials, the capital required to manufacture the drug, and the sales at maturity. Although the drug candidates never have undergone any actual testing beyond the discovery stage, it often is possible to extract subjective probability estimates of their anticipated performance from the personnel involved (Morgan and Henrion, 1990; Nutt, 1998). Clinical directors, for example, generally have an opinion on anticipated behavior of the drug in clinical trials based on its performance in discovery research and their experience with similar drugs. Engineers are trained to estimate the capital required for manufacturing a candidate drug from the structure of the molecule and the chemical or biological process

used to manufacture discovery samples. Finally, market research studies and forecasting practices (Cooper, Edgett, and Kleinschmidt, 1999; Kahn, 2002) provide price and sales estimates for the product at some launch date in the future. In all three of these situations, the uncertainty in the estimates is quite large, ranging from 50 to 100 percent of the most likely values.

Types of Product Dependencies

Individual product returns frequently are influenced by other products in the pipeline and by competitor products (Roberts, 1999). In some instances, it may be even advantageous to prioritize and to schedule a candidate drug for early development despite unattractive financials and low technical success probabilities, because it will provide a knowledge base to forecast the success better for dependent drug candidates later in the product sequence.

In this article, four frequently occurring types of dependencies are considered: (1) resource dependencies (Verma and Sinha, 2002); (2) manufacturing cost dependencies; (3) financial return dependencies; and (4) technical success dependencies. Learning curve effects frequently lead to *resource dependencies*. A common example occurs when the development times are reduced for the trailing candidate of two functionally similar drug types. *Cost dependencies* occur when the combined cost of a development activity for two drug candidates is less than the sum of the individual costs because of resource sharing. For example, it may be possible to use the same production facilities for two chemically or biologically similar drug candidates. *Financial return dependencies* occur when there is synergism or antagonism in the marketplace (Roberts, 1999). For example, cannibalization can occur when two drug candidates are aimed at developing products that compete with one other in the marketplace. Alternatively, several candidates can complement one another such that sales of one actually may enhance the sales of another. *Technical dependencies* occur when the technical success or failure of a drug candidate affects the probability of technical success of an as-yet-untested trailing drug candidate. For example, two drug candidates might be developed to release an active ingredient in a controlled fashion. If the precedent candidate is successful, the probability of success of the as-yet-untested second candidate will be increased.

A Pipeline Portfolio Case Study

Blau and Sinclair (2001) introduced an example problem for managing the portfolio of nine new drug candidates. Building on this example to accommodate interdependencies, the problem case study shown in Tables 2 and 3 was selected. The problem is large enough to ensure that at least one drug will be launched yet is small enough to illustrate that the impact of the various dependencies. The uncertainty data associated with the drug candidates, simply identified as D1 through D9 and targeted to treat three different diseases, is shown in Table 2. The technical success probabilities in each of the three clinical trial

stages are given as simple point estimates, while the manufacturing costs of designing and building the manufacturing facilities are represented by triangular probability distributions. Although it is impossible to estimate the sales ramp-up five or six years into the future, market research can supply triangular probability distributions for sales estimates at maturity (Table 2). Additionally, the right-most column of Table 1 shows the deterministic total resource availability for each activity of Figure 1. In effect, the nine drug candidates must compete for this limited resource pool.

This example may appear restrictive in the sense that all candidates enter the pipeline at the same place:

Table 3. Dependency Data for Nine-Drug Example of Table 2

Disease	Drugs	Type of Dependency	Explanation of the Dependency
Disease I	D2, D3, D6, D7	Financial Dependency	Here sales are dependent on development sequence. If two drugs for Disease I are successfully developed then sales for each candidate are 0.85 of the independent sales of individual value shown in Table 2, if three drugs for Disease I make it through the pipeline successfully, then sales are 0.75 of the independent sales, if all four drugs make it through successfully, then sales of each candidate is 0.6 times the independent sales figures from Table 2.
		Technical Dependency	If the first drug in the sequence of drugs targeted for Disease I fails, the probability of technical success for all succeeding drugs decreases by 50%. On the other hand, if the first in the sequence for testing Disease I succeeds, the probability of technical success for all succeeding drugs for Disease I increases by 10%.
		Manufacturing Cost Dependency	For any sequence of drugs for Disease I, the 1 st drug uses full capital shown in Table 2, the 2 nd drug in the sequence uses $\frac{1}{2}$ of its individual capital, the 3 rd drug uses $\frac{1}{3}$ of its capital cost, while the 4 th drug uses $\frac{1}{4}$ of its capital shown in Table 2.
		Resource Dependency	The time reduction by virtue of learning curve experience is translated into a degree of difficulty reduction of 20% of the values shown in Table 2 for every drug in the sequence for Disease I.
Disease II	D4, D5, D8, D9	Financial Dependency	Total market for the drugs of Disease II is fixed at 9000 million dollars.
		Technical Dependency	If the first drug in the sequence targeted for Disease II fails, the probability of technical success for all succeeding drugs for Disease II decreases by 50% from the values shown in Table 2. On the other hand, if the first drug in the sequence targeted for Disease II succeeds, the probability of technical success of all succeeding drugs for Disease II increases by 10% of the values shown in Table 2.
		Manufacturing Cost Dependency	For any sequence of drugs for Disease II, the 1 st drug uses full capital listed in Table 2, the 2 nd drug in sequence uses 1/2 of its individual capital, the 3 rd drug uses 1/3 of its given in Table 2.
		Resource Dependency	The time reduction by virtue of learning curve experience is translated into a degree of difficulty reduction of 20% of the value shown in Table 2 for every drug in sequence targeted for Disease II.
Disease III	D1	No Dependencies	

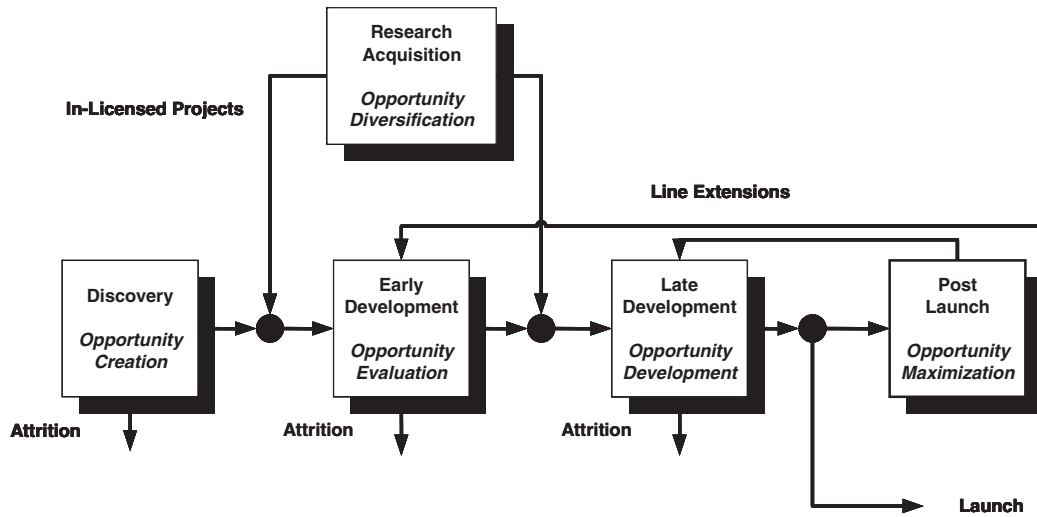


Figure 3. A Typical Pharmaceutical Business Process (Blau and Bunch, 2002)

the start of development. In reality, candidates can enter at different points along the pipeline to accommodate in-licensing and outsourcing options as well as existing product line extensions. Figure 3, taken from Blau and Bunch (2002), shows these various options. Such an analysis is beyond the scope of the current article, although it is accommodated readily with the methodology that will be described.

For the nine-drug-candidate example introduced in Table 2, a realistic set of dependencies is shown in Table 3. The financial return dependencies are different for two of the three targeted diseases; however, the technical, manufacturing cost, and resource dependencies are the same. In both cases the financial return from developing more than one candidate is less than the sum of developing the drugs alone. However, for candidates targeting disease II the total market is fixed, while for disease I the total market increases to accommodate the additional candidates, although at a slower rate. The probabilities of technical success, resource, and cost dependencies all depend on the development sequence. Consequently, until this sequence has been defined, it is not possible to calculate the required resources, returns, or the success probabilities.

In summary, the portfolio-planning problem associated with this case study is to decide which of the nine new drugs to select for development and, once selected, the sequence to move them through the NPD pipeline to maximize the return at the least possible risk for the given resource capacities. The methodology for solving this problem will be described next.

Discrete Event Simulation, Bubble Charts, and Valuation of New Drug Candidates

The network of activities shown in Figure 1 and the probability-based data shown in Table 1 together constitute a probabilistic network model. Such a network is an extension of the PERT network formulation (Elmaghraby, 1977) to accommodate the conditional nature of the product path following an activity. For example, consider the simple network shown in Figure 4, consisting of two activities followed by a decision point. Once activities A and B are completed, a decision must be made either to advance to activity C, to return to a previous activity, or terminate the development effort. Since the decision depends on the uncertain outputs of A and B, the network is said to be probabilistic.

A discrete event simulation program, CSIM (2000), was used to simulate the pharmaceutical pipeline shown in Figure 1. Given a set of drugs and the sequence by which they enter the pipeline, the program simulates the behavior of the drug pipeline for randomly selected values drawn from the probability




Figure 4. Probabilistic Decision-Making




Figure 5. NPV Distribution for D1

distributions of Tables 1 and 2. As these drugs move through the pipeline they accumulate development costs, including costs of building a manufacturing facility, until they enter the market place where the product is ramped up until mature sales are realized. The total resource levels specified in Table 1 are treated as inviolate; consequently, queues of drugs at the various activities may occur. These queues slow the development process, and, consequently, the time to launch the various candidate drugs will be lengthened. If a drug is terminated during development, no returns are obtained, and all development costs are lost. By selecting a sufficient number of different “samples of values” at random from the probability distributions (Monte Carlo trials), performing the simulations, and gathering the results, it is possible to determine probability distributions for various economic returns and risk indicators.

Before considering a portfolio of products, it is informative to examine the behavior of each individual drug candidate. Using net present value (NPV) with an internal rate of return of 15 percent as the economic criterion, the behavior of each drug can be simulated by using the discrete event simulation program. The NPV distribution for D1 is shown in Figure 5 following a sufficiently large number of Monte Carlo trials. This two-peaked distribution is typical of a new drug candidate in the pharmaceutical industry. The first peak corresponds to the loss of money in those instances when D1 fails to pass all the clinical trials. That is, D1 technically is unsuccessful. The second distribution corresponds to the returns following a successful product launch. For D1, the mean “loss” for those trials in which the product failed to advance to the launch state is \$214MM. The primary variation around this value is due to the stage

at which the product failed. The mean “reward” following launches is \$3887MM. The mean value or expected net present value (ENPV) of the overall distribution, which is \$1425MM, lies somewhere between the mean loss and mean reward and truly is meaningless since it never can occur. This is a profound result. It is remarkable that anyone would use a single, deterministic, meaningless ENPV value to describe the anticipated return from a drug. Yet this is frequently the case in pharmaceutical and other NPD companies in order to simplify calculations or to avoid the reluctance of dealing with uncertainties. To overcome the limitations of the ENPV measure, the mean “reward” is described in this study in terms of expected positive net present value (EPNPV), which is defined as the expected value over the positive axis of the NPV distribution. The information about the negative part of the distribution will be conveyed by using a risk measure called probability of losing money—the area under the negative axis of the NPV distribution.

Blau et al. (2000) and Blau and Sinclair (2002) introduced the concept of a reward/risk ratio, obtained by dividing the mean reward by the mean loss, to measure the attractiveness of new product candidates. Intuitively, drug candidates with high reward/risk ratios would be more attractive than those with low reward/risk ratios. For D1 this ratio is 18.2, the expected reward (or return) for every dollar expected to be lost in the NPD. Here risk is defined as the probability of losing money.¹ The risk of losing money for D1 is 24 percent, which is the area under the curve in Figure 5 from an NPV of $-\infty$ to zero. Since

¹ i.e., $P(\text{NPV} < 0)$

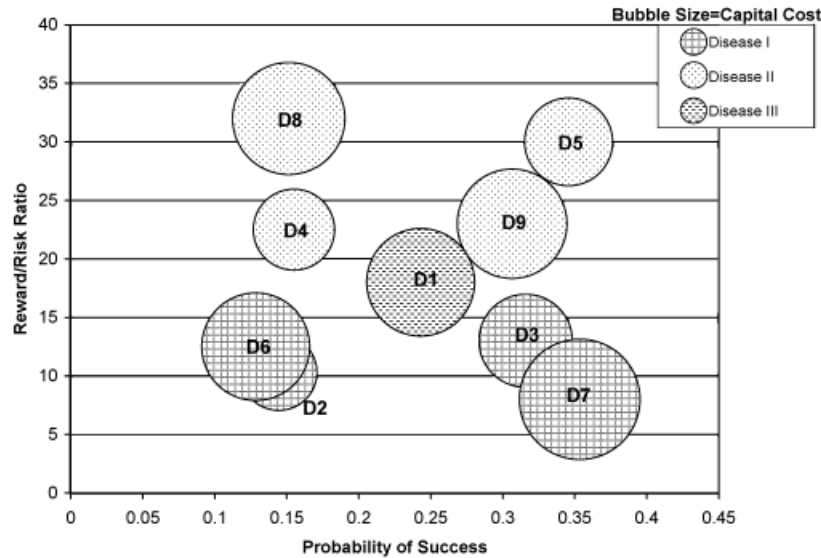


Figure 6. Bubble Diagram for Nine New Candidate Drugs of Table 2

losses only can occur because of failure of the clinical trials, risk is simply the complement of the product of the probabilities of the drug successfully passing the three clinical trials, or 76 percent for D1.

Individual reward/risk ratios can be determined for all the drugs in Table 2. By plotting these values versus an independent factor, the probability of successfully navigating the pipeline (i.e., probability of technical success), the bubble diagram in Figure 6 is obtained. The size of the bubble in Figure 6 is the expected capital cost of building a manufacturing facility, and the shading of the bubble corresponds to the target disease being treated. Intuitively, the most attractive drug candidates are in the upper-right quadrant where both the reward/risk ratios and probability of technical success are high. Within this quadrant smaller bubbles are preferred over larger ones since the capital requirements are less. In the same way, large bubbles for drug candidates in the lower-left quadrant are the least attractive since they have a low reward/risk ratio and low probability of technical success and are expensive to manufacture. Using such intuitive arguments, D5 and D9 should be given preferential treatment whereas D2 and D6 are unattractive. It is less clear what to do with the drug candidates from the other two quadrants. They involve tradeoffs among reward/risk ratios, the probabilities of technical success, and the size of the bubble. The disease being targeted frequently can help in this selection process. For example, in Figure 6 the four drugs targeted to treat disease type II have higher re-

ward/risk ratios than the four drugs targeted to treat disease type I. If another criterion were a desire to maintain diversity in the portfolio, it would make sense to develop D3 and D7 ahead of D4 and D8. By using this type of reasoning for the four criteria (reward/risk ratio, probability of technical success, manufacturing cost, and a diverse product mix), an intuitively appealing relative ranking for the nine candidates would be D5, D9, D3, D7, D1, D4, D8, D2, and D6.

In practice, such an intuitive prioritized approach is too simplistic. It includes uncertainties and deals with multiple criteria but does not handle dependencies between drug candidates, nor does it consider resource conflicts. However, it serves as a good heuristic and starting point for the more comprehensive selection and sequencing process described in the next section.

Optimizing a Portfolio of Interdependent New Drug Candidates

A formal statement of the portfolio optimization problem is as follows: Select a set of new drug candidates, and sequence them for the development process in such a way that the economic return expressed as the expected positive new present value (EPNPV) is maximized for a given level of risk measured as the probability of losing money. Stated as a mathematical program the portfolio optimization problem is to

Table 4. Initial Population for the Genetic Algorithm

Ranking Basis	Drug1	Drug2	Drug3	Drug4	Drug5	Drug6	Drug7	Drug8	Drug9	Average NPV (million \$)	Probability of Losing Money (%)
Bubble Chart	6	0	4	7	1	0	5	3	2	1131.43	33.23
Bubble Chart	3	9	4	7	1	8	5	6	2	1265.66	25.44
Bottleneck	5	2	7	8	3	6	1	4	9	894.228	31.34
Random Order	4	0	2	0	6	5	1	0	3	1045.63	20.84
Random Order	4	1	0	0	0	2	3	5	6	473.222	33.37
Random Order	4	2	1	3	5	0	0	0	0	819.52	25.73
Random Order	5	0	4	0	2	3	6	1	0	1060.48	32.82
Random Order	2	3	0	7	1	6	0	5	4	1179.92	26.7
Random Order	1	0	0	0	4	5	0	3	2	1350.28	29.38
Random Order	1	0	3	0	4	0	6	2	5	1353.62	19.82

Maximize $EPNPV$ (over all available drug selections and sequences)

Subject to

$P(NPV < 0) < \beta$ (Risk Constraint)

where β , the risk of losing money, may be set at 5 percent, 10 percent, or any other value coincident with the risk tolerance level of management for their NPD efforts. Decreasing β below a certain level (i.e., being excessively risk averse) will reduce the $EPNPV$ significantly, while high levels of β will introduce excessive risk into a portfolio.

In order to find the “optimal” selection and sequence of drugs, this conceptual mathematical programming model needs to be transformed into a quantitative model. It is important to note that a highly combinatorial and statistically complex decision problem such as research and development (R&D) portfolio management cannot be formulated by using traditional differential equation or mathematical programming models without making unrealistic assumptions. As an alternative, a “computational model” is presented that is built out of the aforementioned bub-

ble charts, a discrete event simulator, and a genetic algorithm. The term *sequence* will encompass not only information on the order in which drugs are developed—their priorities—but also on which drugs are included in the portfolio: their selection. The methodology can be summarized as follows:

- (1) An initial list of 10 sequences of drug candidates is generated, some from the bubble chart using individual drug analysis and others at random (Table 4). In Table 4 the “0” indicates that a candidate was not selected, and a positive integer indicates the position of the drug in the sequence.
- (2) For every sequence, the probability distributions associated with the activities for each of its selected drug candidates are modified or are “pre-processed” to account for the dependencies in Table 3.
- (3) The behavior of each sequence is simulated by using a discrete event simulator.
- (4) The results from these simulations are used by a genetic algorithm to search for improved drug sequences.

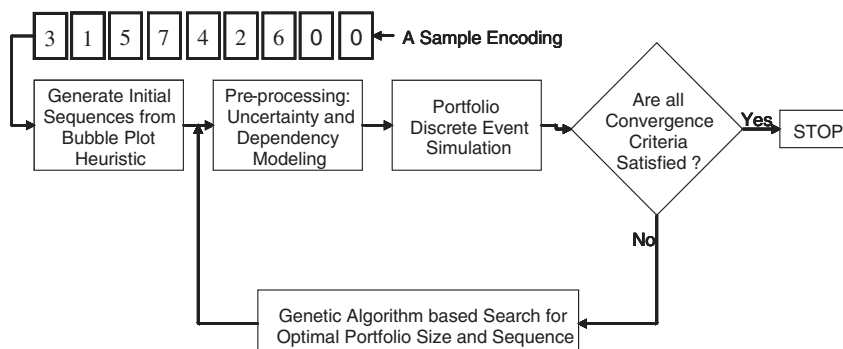


Figure 7. Computational Framework for Determining Optimal Portfolio

The various steps in the algorithm are presented in Figure 7. Since the genetic algorithm is the key to finding improved sequences, its implementation requires some explanation. Genetic algorithms are biased stochastic search algorithms (Gen and Cheng, 1997; Goldberg, 1989; Holland, 1975) based on analogies to Darwinian principles of natural selection. As with evolutionary processes, genetic algorithms are iterative, gradually evolving from some initial state to an optimized one. Here, the initial state is the initial list of drug sequences, with successive iterations improving on previous sequences until a suitably optimized one is found. The GA “encodes” sequence selection decisions as a rank ordering of drug candidates. In Figure 7, the sample encoding (3, 1, 5, 7, 4, 2, 6, 0, 0) implies that D1 is developed third in priority, D2 is developed first, D3 is developed fifth, and so on, with D8 and D9 not developed. A collection of such encoded sequences is called a *population*. In our implementation, 10 sequences are maintained per population. Thus, 10 sequences generated from the bubble chart-based procedure and from random selections were used to initialize the GA procedure. It is important to select a good set of sequences from the bubble chart procedure in order to keep the number of evolutions (i.e., iterations) to realize the optimal sequence at a computationally reasonable level.

For each of these candidate sequences included in the initial or subsequent populations, a so-called “fitness” function is assigned, which measures how closely the sequence not only maximizes economic performance but also minimizes the probability of losing money. A fitness function, Z_k , is calculated for each of the $k=1,2,\dots,10$ candidate sequences in the current population by normalizing the EPNPV and risk as follows:

$$Z_k = \alpha \left(\frac{EPNPV_k - EPNPV_{\min}}{EPNPV_{\max} - EPNPV_{\min} + \gamma} \right) + (1 - \alpha) \left(\frac{Risk_{\max} - Risk_k}{Risk_{\max} - Risk_{\min} + \gamma} \right),$$

where $EPNPV_{\min}$ and $EPNPV_{\max}$ are the minimum and maximum expected positive NPVs, respectively, in the current population; $Risk_{\max}$ and $Risk_{\min}$ are the maximum and minimum risk probabilities in the current population; and γ is a small positive number that prevents division by zero. The nonnegative number α (between zero and one inclusive) is inversely proportional to the cost per unit violation of the risk constraint, written at a risk tolerance level of β . Further

quantitative details about α are provided in Appendix A (see <http://atom.ecn.purdue.edu/~varmav/Pharma-Suite/Appendix.doc>).

Thus, any violation of the risk constraint is penalized at a “cost” of $(1 - \alpha)$, thereby progressively leading the GA to sequences feasible with respect to the risk constraint or at worst violate the risk constraint to as small a degree as possible. Typically, α is specified based on the decision maker’s risk aversion. For instance, a “speculative” pharmaceutical management may fix α at a high value, thereby downplaying the cost of risk. In this study, α was fixed at 0.8, thus rendering the return as the dominant criterion.

The GA proceeds to find sequences that improve the fitness function by a process called “reproduction,” i.e., creation of a population of new sequences or “offspring” by using sequences in the initial population as “parents.” This procedure is carried out by employing certain probabilistic transition rules called genetic operators that modify the encoded information of the parents to form offspring. These rules are discussed in Appendix B (available at <http://atom.ecn.purdue.edu/~varmav/Pharma-Suite/Appendix.doc>).

At the end of this reproduction process, there exists a population of 10 parent sequences and at most 10 offspring sequences. For each of these offspring sequences, the fitness function must be calculated by using the interdependency preprocessor and the discrete event simulator. The number of Monte Carlo trials used by the discrete event simulator affects the quality of the NPV distribution. A relatively small number of trials (e.g., 100) were adequate to estimate the EPNPV. However, the estimate of risk required many more trials (e.g., 10,000) to generate a reasonable degree of accuracy because it was calculated from the tail of the distribution. Once the fitness of each sequence is calculated, the GA selects the 10 fittest sequences from the expanded population of parent and offspring sequences, resulting in a new population of sequences. The remaining less fit sequences “die.” The sequences in the new population then form the parents for the generation of the next set of sequences, and the reproduction–selection procedure repeats itself until a suitable fit is obtained. The key feature that makes this search “intelligent” as compared to “random” is the selection process. The selection process favors the mating of the fittest sequences in the population resulting in more fit populations until the improvement from population to population is minimal.

Initiation of the GA Search Procedure

For the nine-drug candidate example, an initial population of 10 sequences was determined as follows:

- (1) Two possible sequences were suggested from the bubble chart of Figure 6: (D5, D9, D1, D3, D7, D8, D4) and (D5, D9, D1, D3, D7, D8, D4, D2, and D6). They differ in that the former does not include D2 and D6 because they both have very low probabilities of technical success and also a low reward/risk ratio. In both sequences D5, with its high probability of success and reward/risk ratio, is developed first followed by D9 and D1. D3 and D7 are preferred over D8 and D4 because D3 and D7 target disease I with a higher level of technical success.
- (2) By examining the resource availabilities for the activities described in Table 1 it is apparent that the Phase III clinical trial is the most limiting resource. Furthermore, drug candidates with high values of DoD (Degree of Difficulty) use more resources than others. Based on these observations, Phase III clinical trial resource requirements were calculated for all nine drugs from Table 1. The sequence D7, D2, D3, D8, D1, D6, D3, D4, D9 was generated by arranging these resource requirements in increasing order.
- (3) An additional seven sequences were generated at random, which leads to greater diversity of sequences in the initial population of the GA. This randomization increases the breadth of the search

space from the narrow region defined by deliberately chosen candidates.

Table 4 shows the initial population used for the genetic algorithm search. The GA was run for a maximum of 100 populations (or generations) on a 1 GHz Intel Pentium III, 256 MB RAM single processor machine. With 10 sequences per population, the search explored a total of 1,000 sequences. The total computer processing unit (CPU) time involved in the search was about 60 hours. Of course, using computers to perform the calculations simultaneously could reduce the time dramatically.

Results of the Search Procedure

The return as measured by the EPNPV and the probability of losing money (portfolio risk) are presented in Figure 8 for each of the 1,000 sequences generated by the GA. All points corresponding to the maximum EPNPV for a given level of risk are joined to form an approximate return-risk frontier. This frontier indicates that the maximum expected return from the portfolio first increases and then decreases as the risk increases, which somewhat is counterintuitive. It would seem more likely that returns would increase as the probability of losing money increases: The greater the risk, the greater the reward (Markowitz, 1991). Such an analogy to conventional portfolio analysis does not apply in managing a pharmaceutical portfolio. Basically, three interdependent factors




Figure 8. Return versus Risk Frontier for the Nine Drug Candidates

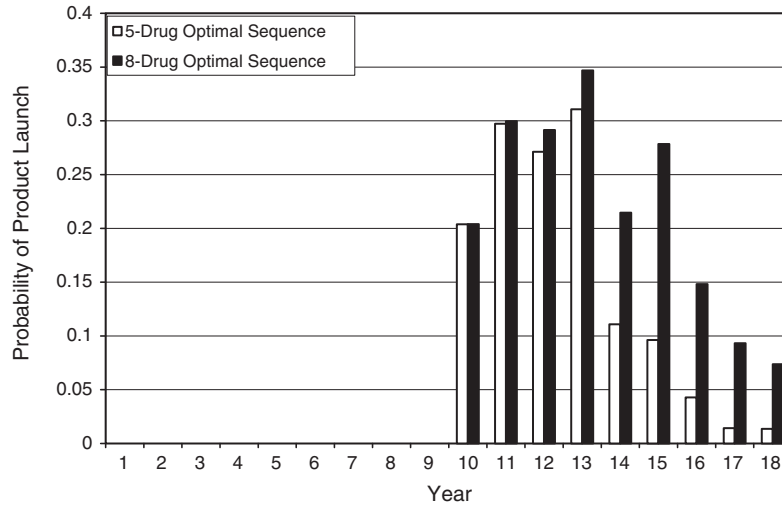


Figure 9. Probability of Product Launch by Year

are operating here: the length of the sequence, the available resources, and the impact of drug failures. Long sequences (eight or nine drugs) cushion the impact of drug failures but strain the available resources, resulting in long queues at bottleneck activities, which in turn extend the time before a successful candidate can be launched. On the other hand, short sequences result in short times-to-launch since adequate resources are available, but the impact of a drug failure dramatically enhances the probability of losing money. Even such gross observations of system behavior are masked by the impact of the dependencies from Table 3; consequently, computationally intensive calculations such as the GA are necessary to capture the sequence substructure. The highest point on the return/risk frontier shown in Figure 8, \$2369MM at a

risk of 22 percent, is projected for the five-drug sequence D5, D8, D9, D3, D1, with D2, D4, D6, and D7 not developed. This return is about 6 percent higher and the risk is about 4 percent less than the best eight-drug sequence D5, D8, D9, D1, D7, D4, D3, D2, with D6 not developed. Other sequences on the frontier, some of which are shown in Figure 8, consist of five, six, seven, or eight drugs.

Sequence length impact can be understood best by comparing the best five- and eight-drug sequences. Figure 9 shows a comparison of the probability of product launch by year for the best five- and eight-drug sequences. A significant proportion of the products launched for the eight-drug sequence will occur after 14 years with only a few years available to recoup development costs before the product loses

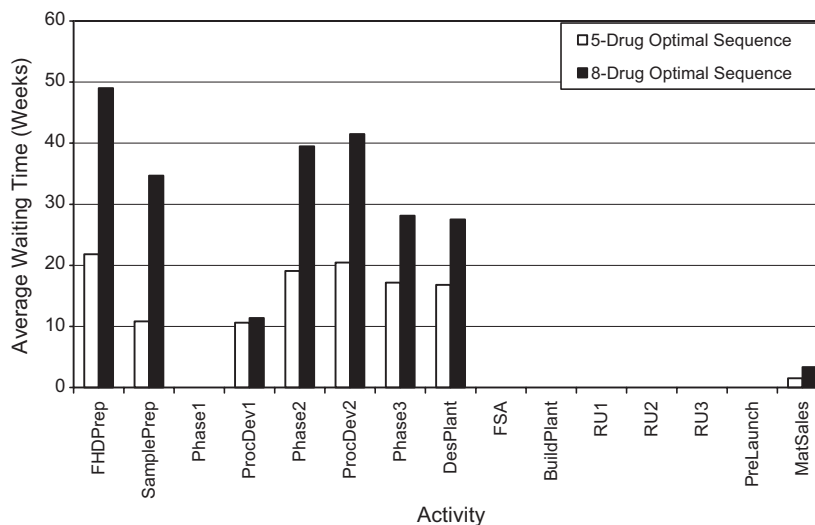


Figure 10. Average Waiting Times on Each Activity for Five- and Eight-Drug Optimal Portfolios

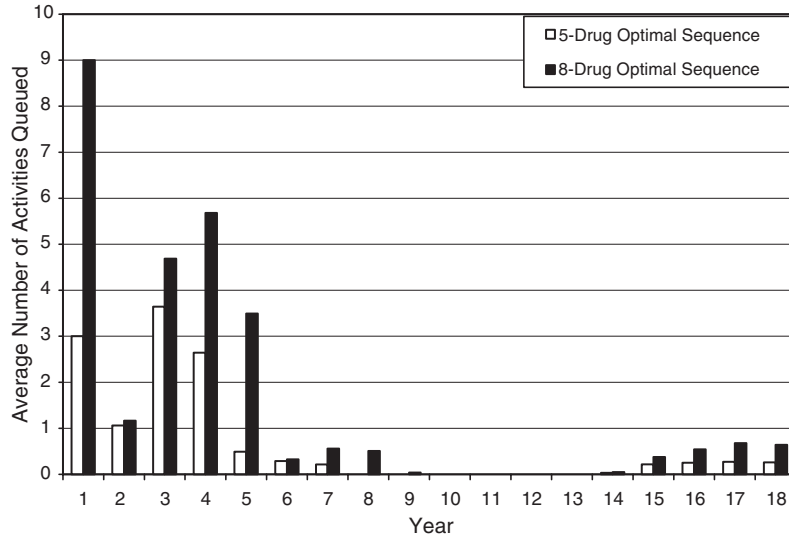


Figure 11. Profile of Number of Activities Queued in System

patent protection. Naturally, successful products could get to market more quickly and losers could be eliminated earlier if more resources were available. Figure 10 displays the average waiting times on resources for the various development activities. The same activities are bottlenecks for both sequences, but the waiting times are significantly longer for the eight-drug versus the five-drug sequences. Another way of looking at the impact of resources is to plot the average number of activities in a queue waiting for resources to be made available (Figure 11). Once again the pipeline has sufficient resources to accommodate the five-drug sequence more readily than the eight-drug sequence. The impact of the limiting resources is to delay the launch of a product by an average of one

year in the eight-drug sequence over the five-drug sequence and the time to launch the first product by a half year. All these factors combine to favor the five-drug sequence over the eight-drug sequence. As a further demonstration of the utility of the approach presented here, it is interesting to compare these GA-generated results with the results from the earlier bubble chart analysis. The nine-drug sequence of D5, D9, D3, D7, D1, D4, D8, D2, D6 is the optimal sequence obtained from inspection of the bubble chart (Figure 6). When this sequence is introduced into the pipeline and the dependencies are accommodated, the EPNPV for this sequence development is about \$1848MM. However, the maximum EPNPV from the best eight-and five-drug sequences is about 20

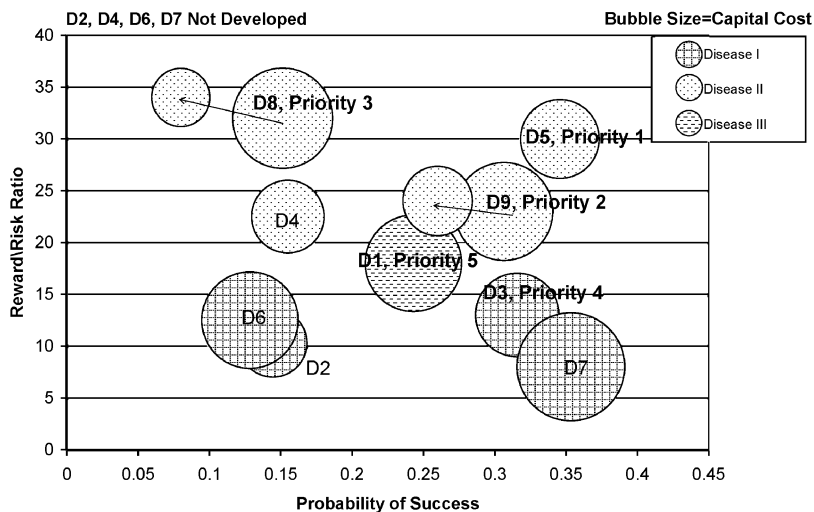


Figure 12. Effective Bubble Chart for Five-Drug Optimal Sequence (5,9,8,3,1)

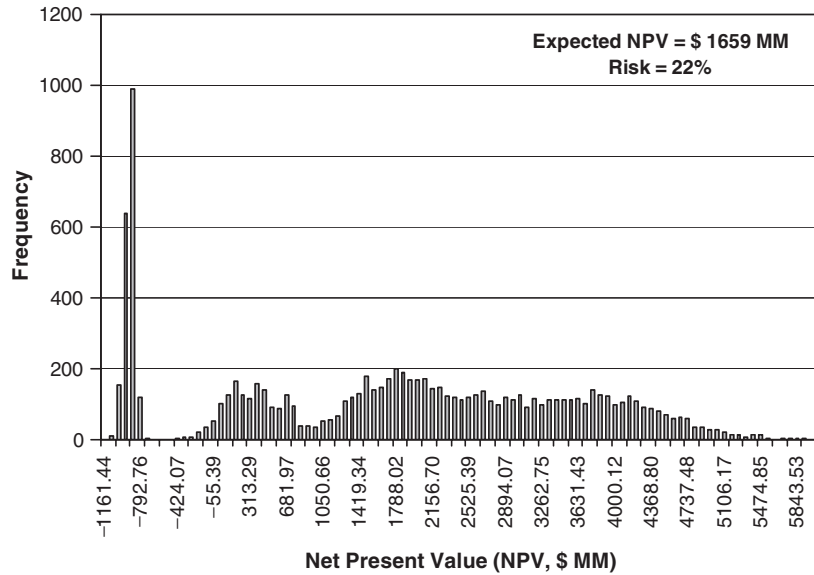


Figure 13. Net Present Value Distribution for the Best Five-Drug Portfolio

percent and 28 percent higher, respectively, than the nine-drug sequence.

Figure 12 is a superposition of the best five-drug sequence (D5, D8, D9, D3, D1) on the bubble chart of Figure 6—where all drugs were assumed independent and the development resources unlimited. The movement of the bubbles quantifies the impact of dependencies. For example, the probabilities of success decrease for all drugs except the first one targeted for a specific disease, because there is a 50-percent

reduction in the probability of success if the first drug fails, while there is only a 10-percent increase in the probability of success if the first drug is launched successfully. Further insights can be gained by examining the impact of the dependencies for the different disease targets.

For disease II, the best sequence from the GA for the best five-drug sequence is D5, D8, D9, which is different than this study’s initial heuristic sequence D5, D9, D8, and D4 obtained from inspection of the

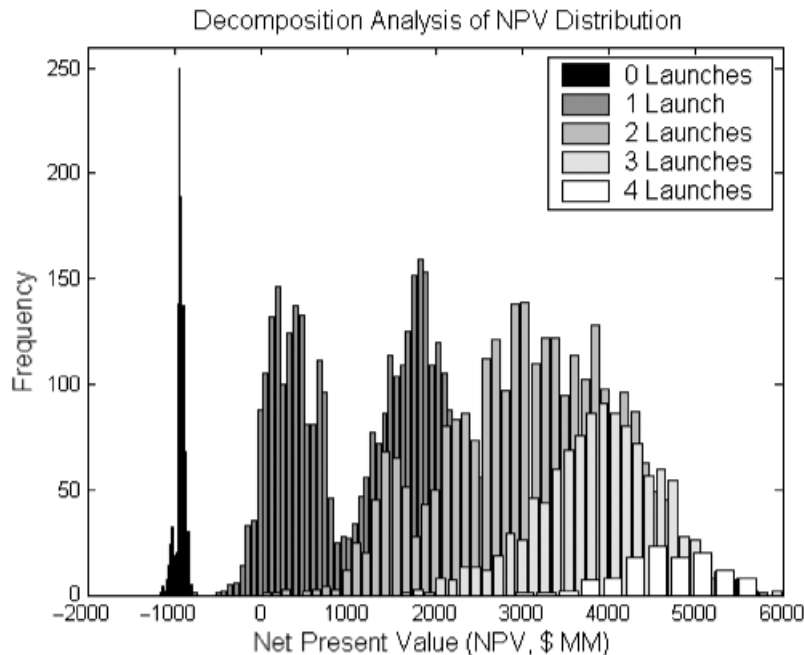


Figure 14. Decomposition Analysis of the NPV Distribution for the Best Five-Drug Portfolio

bubble diagram. D5, by virtue of its high reward/risk ratio and high probability of technical success, is developed first. However, D8 is selected ahead of D9 because the dependency relationship between them changes the probability of technical success and manufacturing capital requirements of D9.

Among the drugs targeting Disease I, D3 is developed for the best five-drug sequence, while D7, D6, and D2 are not developed. On the other hand, the best eight-drug sequence develops three D3, D7, and D2. Yet the best eight-drug sequence results in lower expected return and higher risk than the best five-drug sequence. This outcome is attributed to the combination of resource constraints and the sales dependency effects in which products cannibalize each other's sales.

Considering the best five-drug portfolio in its entirety, D5 is developed first followed by D8 and D9, all targeted for Disease II. Clearly, the cost, resource, and learning curve dependencies are realized to a significant degree for this five-drug sequence. Thus, resource dependency is likely to significantly reduce queuing for resources due to lower capital and resource requirements when drugs of the same disease are developed. The five-drug sequence includes D3, which targets disease I, and D1, which targets Disease III. Hence, the portfolio is sufficiently diverse, even in the presence of a limited number of drugs. It is remarkable that diversity and returns have been balanced without increasing the risk for the five-drug sequence.

Resource Planning

An important use of the methodology presented here is to determine the impact of changing resource levels on the performance of the NPD. Figure 10 can be used to identify which resources are limiting. However, the combinatorial nature of the problem makes it difficult to evolve an optimal strategy to increase these resources while simultaneously increasing the returns and decreasing risk. For example, consider the possibility of outsourcing all manufacturing activities: sample preparation, process development, plant design, and plant construction. From Figure 10 there are queues of products waiting to execute these activities. However, when the optimization procedure was used to determine the best sequence, there was little change in the expected returns or risk. The queues behind other resource-limited activities simply grew in length. However, increasing the resources for the FHD preparation activity in

conjunction with manufacturing outsourcing dramatically increased the returns and reduced risk by making it possible to develop an eight-drug sequence without excessive delays versus a five-drug sequence.

This type of analysis is very application specific but is essential to providing management with valuable insights into the interaction among the various selection criteria and the capacity of the pipeline.

The Broad Variance of the NPV Distribution

The NPV distribution for the five-drug optimal sequence with 10,000 Monte Carlo trials is shown in Figure 13. The return ranges from –\$1,200 million to \$5,700 million. The distribution does not show a strong central tendency and seems to be more representative of a uniform rather than the intuitively anticipated normal distribution. This variation is attributable mostly to the failure rates of the drug candidates. For this sequence six possible outcomes can occur: 0, 1, 2, 3, 4, or 5 drugs may be successfully launched. The returns from these six events are dispersed widely. When no drugs are launched, the returns are negative, whereas the returns are extremely large and positive when all five candidates are launched. Naturally, each of these cases has an associated distribution resulting from all the other uncertainties alluded to in this paper. In Figure 14, the distributions for the first five of the six different cases (different number of product launches) are distinguished. Because of this behavior it is important to deemphasize the importance of a central tendency measure such as the mean of the NPV distribution and instead focus on other characteristics of the distribution such as risk, capital, and diversity issues.

Conclusion

The paper has demonstrated a novel quantitative approach to determining an optimal portfolio for the development of new pharmaceutical compounds. The method provides a uniform treatment of both project uncertainties and dependencies, which are inherent to this industry. Combining discrete event stochastic simulation with genetic algorithms provides the computational synergism required to optimize the highly combinatorial portfolio management problems facing modern pharmaceutical businesses. This study has offered a discussion of simple bubble-chart-based

graphical aids that not only assist management in understanding the potential rewards and risk for various resource levels and mix of products but that also provide “good” starting solutions to the GA. Although the approach is computationally demanding, the rewards are significant in allowing management a vehicle for determining an optimal portfolio as well as appreciating the robustness of the portfolio to changing economic conditions, acceptable risk, and resource levels. Although illustrated for the pharmaceutical industry, this work can be extended readily to any industry such as aerospace, agrochemical, or biotechnology that is regulated highly and loses new product candidates during the development process.

References

- Adler, P.S., Mandelbaum, A., Nguyen, V. and Schwerer, E. (1995). From Project to Process Management: An Empirically Based Framework for Analyzing Product Development Time. *Management Science* 41(3):458–484.
- Amram, M. and Kulatilaka, N. (1999). *Real Options: Managing Strategic Investment in an Uncertain World*. Boston, MA: Harvard Business School Press.
- Blau, G.E. and Bunch, P. (2002). Process Modeling in New Product Development: PDMA New Product Development Toolbook. New York: John Wiley & Sons.
- Blau, G.E., Mehta, B., Bose, S., Pekny, J.F., Sinclair, G., Kuenker, K. and Bunch, P. (2000). Risk Management in the Development of New Products in Highly Regulated Industries. *Computers and Chemical Engineering* 24(7):659–664.
- Blau, G.E. and Sinclair, G. (2001). Dealing with Uncertainty in New Product Development. *Chemical Engineering Progress* 97(6):80–83.
- Chapman, C. and Ward, S. (1996). *Project Risk Management: Processes, Techniques and Insights*. Rexdale, Ontario: Wiley.
- Cooper, R.G., Edgett, S.J. and Kleinschmidt, E.J. (1999). New Product Portfolio Management: Practices and Performance. *Journal of Product Innovation Management* 16(4):333–351.
- Copeland, T.C. and Antikarov, A. (2001). *Real Options: A Practitioner's Guide*. New York: Texere Publishing Limited.
- CSIM (2000). The CSIM 18 User Manual. <http://www.mesquite.com>. Mesquite Software Inc.
- Ding, M. and Eliashberg, J. (2002). Structuring the New Product Development Pipeline. *Management Science* 48(3):343–363.
- Elmaghraby, S.E. (1997). *Activity Networks: Project Planning and Control by Network Models*. New York: Wiley Science.
- Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. New York: Wiley.
- Gen, M. and Cheng, R. (1997). *Genetic Algorithms and Engineering Design*. New York: Wiley.
- Griffin, A. (1997). PDMA Research on New Product Development Practices: Updating Trends and Benchmarking Best Practices. *Journal of Product Innovation Management* 14(6):429–458.
- Gupta, A.K. and Wilemon, D. (1996). Changing Patterns in Industrial R&D Management. *Journal of Product Innovation Management* 13(6):497–511.
- Holland, J.H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- Hultink, E.J., Griffin, A., Hart, S. and Robben, H.S. (1997). Industrial New Product Launch Strategies and Product Development Performance. *Journal of Product Innovation Management* 14(4):243–257.
- Kahn, K.B. (2002). An Exploratory Investigation of New Product Forecasting Practices. *Journal of Product Innovation Management* 19(2):133–143.
- Krishnan, V. and Ulrich, K.T. (2001). Product Development Decisions: A Review of the Literature. *Management Science* 47(1):1–21.
- Law, A.M. and Kelton, D.W. (2000). *Simulation Modeling and Analysis*. Boston, MA: McGraw-Hill & Co.
- Linton, J.D., Walsh, S.T. and Morabito, J. (2002). Analysis, Ranking, and Selection of R&D Projects in a Portfolio. *R&D Management* 32(2):139–148.
- Loch, C.H. and Bode-Greuel, K. (2001). Evaluating Growth Options as Sources of Value for Pharmaceutical Research Projects. *R&D Management* 31(2):231–243.
- Markowitz, H.M. (1991). *Portfolio Selection*, 2d ed. Cambridge, MA: Blackwell Publishers.
- Morgan, G. and Henrion, M. (1990). *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge, UK: Cambridge University Press.
- Nutt, P.C. (1998). How Decision-Makers Evaluate Alternatives and the Influence of Complexity. *Management Science* 44(8):1148–1166.
- O'Connor, P. (1994). Implementing a Stage-Gate Process: A Multi-Company Perspective. *Journal of Product Innovation Management* 11(3):183–200.
- Pisano, G.P. (1997). *The Development Factory: Unlocking the Potential of Process Innovation*. Cambridge, MA: Harvard Business School Press.
- Poh, K.L., Ang, B.W. and Bai, F. (2001). A Comparative Evaluation of R&D and Project Evaluation Methods. *R&D Management* 31(1):63–75.
- Repenning, N.P. (2001). Understanding Fire Fighting in New Product Development. *Journal of Product Innovation Management* 18(5):285–300.
- Roberts, W.R. (1999). Product Innovation, Product-Market Competition, and Persistent Profitability in the U.S. Pharmaceutical Industry. *Strategic Management Journal* 20(7):655–670.
- Rodriguez, D. (1998). Decisions of Pharmaceutical Firms for New Product Development. MIT Industrial Performance Working Paper. Also available at <http://web.mit.edu/ipc/www/Decisions.pdf>.
- Sharpe, T. and Keelin, T. (1998). How Smithkline-Beecham Makes Better Resource Allocation Decisions. *Harvard Business Review* (March–April):45–57.
- Subramanian, D., Pekny, J.F. and Reklaitis, G.V. (2003). Simulation–Optimization Framework for Stochastic Optimization of R&D Pipeline Management. *AIChE Journal* 49(1):96–112.
- Taylor, B.W. and Moore, L.J. (1980). R&D Project Planning with Q-GERT Network Modeling and Simulation. *Management Science* 26(1):44–59.
- Varma, V. (2004). Appendix. <http://atom.ecn.purdue.edu/~varmav/pharma-suite/Appendix.doc>.
- Verma, D. and Sinha, K.K. (2002). Toward a Theory of Project Interdependencies in High-Tech R&D Environments. *Journal of Operations Management* 20(5):451–468.