16th European Symposium on Computer Aided Process Engineering
and 9th International Symposium on Process Systems Engineering
W. Marquardt, C. Pantelides (Editors)

A framework for capturing the impact of resource allocation policies in the selection of a new product portfolio

Juan Camilo Zapata[a], Vishal A. Varma[b], and G. V. Reklaitis[a]

[a] *School of Chemical Engineering, Purdue University, West Lafayette, IN 47907, USA*
[b] *Air Products and Chemicals, Allentown, PA 18195, USA*

Abstract

The high development costs, low probability of success and intensive competition faced by pharmaceutical companies make management of their product pipelines a high risk undertaking. The strategic decision involving the selection of the particular set of drugs to be developed has implications that affect the behavior of the pipeline for years. While recently reported research has captured the stochastic character of the pipeline, to date no methodology has explicitly included the impact of operational policies in the selection process. In this work, a multi-level Sim-Opt strategy is used to assess the effect of resource allocation on risk and rewards. Product sequences generated by a GA are statistically evaluated using a probabilistic network model. The model includes all the tasks that have to be accomplished in order to release a new drug into the market. The resources assigned to each drug in each task are rebalanced by an optimal policy every time a project fails and at the end of each year. Based on the results a reward-risk frontier is constructed and compared to the one generated when no reactive allocation is considered. Results show that the inclusion of this additional degree of freedom in the decision process causes a significant change in the portfolio mix.

Keywords: Strategic and tactical decisions; Sim-Opt; efficient frontier.

1. Introduction

Portfolio management is a dynamic decision process, that provides the framework in which senior management operationalizes its business strategy. The future directionality of such a strategy is translated into the R&D portfolio. However, the presence of uncertainty, multiple objectives and decision makers, project interdependencies and a constantly changing environment makes that translation a very difficult process. A wide variety of methodologies has been proposed to facilitate this process (Cooper et al., 1999). In spite of the differences, all of them share one characteristic: the implicit use of a decomposition philosophy. The problem is broken down into different independent hierarchical levels, where each level uses a set of data and models whose degree of aggregation depends of the scope of the corresponding level. Intuition suggests that such an approach is valid when the project execution process tends to be deterministic. However in the highly dynamic and constrained environment of R&D (projects fail, uncertainties are reshaped by internal changes and the surrounding environment, resources are not easy to replace, etc.) the answer is not so clear. Under these conditions, optimal strategic decisions may require the integrated consideration of key aspects of the tactical and strategic levels. Sensitivity analysis is sometimes implemented to assess the impact of using aggregated data. However, the results obtained by this strategy present two major pitfalls. First, the interdependencies between the projects are not captured. And second, completely different portfolio realization paths at the tactical level, with their corresponding rewards and risk levels, can be

obtained for the same set of aggregated values. This study explores the implications of the choices at the tactical level on the selection and prioritization of new products in the R&D portfolio of a pharmaceutical company.

2. Tactical vs. Strategic Decisions

The development of decision support strategies and systems for managing an R&D portfolio date back to the 60s. Since then it has become evident that R&D portfolio management is about minimizing risk while maximizing an objective or a set of objectives in the presence of constraints (Baker, 1975). In order to accomplish this goal many decisions have to be made at different levels in the organization. Depending on their scope they can be classified in two groups: strategic or tactical. Both of them imply the allocation of resources; the only difference is that the first group determines the objective, while the second one leads to it. At the strategic level some techniques were developed to support the selection of projects and their priority in the portfolio (Blau et al., 2004; Lin and Hsieh, 2004; Raynor and Leroux, 2004; Rogers et al., 2002). Others concentrated on the selection of one project from a group of candidates (Calantone et al., 1999; Loch and Bode-Greuel, 2001). At the tactical level, the focus has been on scheduling and allocating resources to activities within the projects (Maravelias and Grossmann, 2004; Subramanian et al., 2003). Regardless of the scope, all these methodologies are based on one of two paradigms, quantitative or qualitative. Real options, decision trees, discrete event simulation, mathematical programming, etc are at the core of quantitative decision support systems (Loch and Bode-Greuel, 2001; Maravelias and Grossmann, 2004; Raynor and Leroux, 2004; Subramanian et al., 2003). On the qualitative side the focus has been the direct translation of the decision makers' knowledge into portfolios and priorities. For that purpose, a wide spectrum of techniques that range from checklists to fuzzy theory have been used (Cooper et al., 1999; Lin and Hsieh, 2004). In spite of all the methodologies developed, quantitative and qualitative, none of them consider both decisions levels at the same time nor provide evidence to support the validity of the decomposition strategy.

3. Portfolio Optimization

The pharmaceutical industry provides one of the most challenging areas in terms of R&D portfolio management. Regardless of the type of product, small molecule chemical compound or complex protein, the industry faces long development times, low success rates, very high investments and considerable uncertainty in sales revenue estimates (Blau et al., 2004; Loch and Bode-Greuel, 2001). It is thus a perfect testing ground for any stochastic decision support methodology.

There are three major stages in the lifecycle of a new drug: discovery, development and commercialization. The discovery stage is highly unpredictable and case specific, while the other two generally follow a well defined activity path. Also, the typical situation in the pharmaceutical industry is that there are seldom enough renewable and nonrenewable resources available to develop all the lead compounds in the pipeline at the same time. Therefore, all the attention from a portfolio management perspective is given to the development and commercialization stages. They are divided in the following sequential activities: First human dose preparation, clinical trials I, II and III, first submission for approval, prelaunch, ramp us sales, and mature sales. Paralleling these activities are all the engineering and marketing related tasks. For a thorough explanation of each of the activities the reader is refered to Blau et al (2004).

3.1. The portfolio management problem

In this study a multi-level optimization version of the SIM-OPT architecture developed by Subramanian et al (2003) is used to solve the portfolio management problem under constrained renewable resources. SIM_OPT combines discrete event simulation and optimization (Fig. 1). The inner loop contains a model of the process and an optimizer that is activated every time an event (i.e. project failure) takes place during the simulation. The outer loop optimizer makes higher order decisions based on the information collected from multiple runs of the inner loop. In the realization of the SIM-OPT architecture employed in Blau et al (2004) only the outer loop is included. A genetic algorithm (GA) is used to optimize the selection of drug candidates, while the discrete event simulation model, which is a complete representation of the probabilistic pipeline network, is used to evaluate the candidate selection and sequencing alternatives generated by the GA. In their work Blau et al (2004), which we call the base case, used the project sequence generated by the GA to determine the order in which projects are started in the development pipeline, as well as the priorities of the activities within projects and the resources allocated to them, regardless of the specific realizations of the uncertainties in the system. It follows the aggregation concept characteristic of the reported decision support methodologies. In our study Blau's work is extended by including an inner loop to reallocate the renewable resources every time a project fails, and refining the GA. In the extended model the portfolio and the priorities generated by the GA are used as starting point to allocate resources and schedule activities. However, those decisions are dynamically updated by an optimizer according to the resolution of uncertainties at the key termination points in the activity network. The optimizer is a decision support system developed by Varma (2005) that maximizes the expected economic return by collecting the most up to date system data and processing it through a series of control policies learned by running the model multiple times using a short time window 3 mode (upper, most likely and lower) resource allocation MILP.

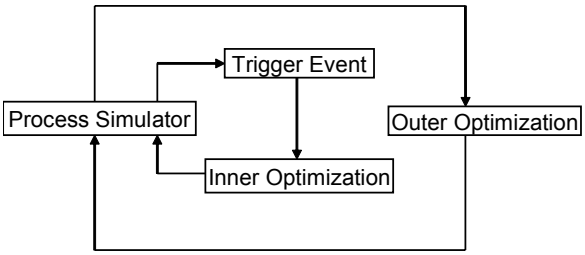

Fig 1. Sim-Opt Architecture

3.2. Outer loop (GA)

The same GA strategy is used for the base case and the extended model trial runs. Following Blau et al (2004) the portfolios are encoded in such a way that each gene contains the number of a drug candidate (with 0 indicating that a project was not selected), and its position in the chromosome represents the priority given to the project. The fitness function, Z_k, is given by:

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

Where $EPNPV_{min}$ and $EPNPV_{max}$ are the minimum and maximum expected positive net present values, respectively, in the current population; $Risk_{min}$ and $Risk_{max}$ are the maximum and minimum risk levels in the current population, measured as the probability of losing money; γ is a small positive number that prevents division by zero, and α weights the present value vs. the level of risk in a convex linear combination. The $\lambda + \mu$ selection strategy for new generations is also used. However, it was found that Blau's reproduction strategy within the GA algorithm caused the optimization to be trapped in a specific area of the search space. Therefore, an additional mutation operator was included to overcome this problem. The operator randomly changes some of the genes in the following way: If the gene holds a value, the corresponding project is abandoned, and if it is empty a project is randomly reentered into the portfolio. In addition to correcting the myopic behavior of the original GA, this adjustment also makes the algorithm very robust to changes in the initial population. Therefore, the risk of biasing the algorithm in the wrong direction from the start, due to the inadequate selection of the initial population, is considerably reduced.

3.3. Inner loop (Scheduling and resource allocation)

Once the portfolio and the priorities of the projects have been identified, it is necessary to schedule the activities within a project and allocate resources to them. In the base case it is assumed that the priority sequence obtained for the projects applies to every activity, and the level of resources assigned is the most likely (ML) one according to the degree of difficulty of the project. The multi-level SIM-OPT strategy, refered as the extended case, on the other hand, follows an adaptive approach in which activities are scheduled based on the GA sequence, but resources assigned to a particular task are dynamically increased or decreased to speed up or slow down a project in response to events such as terminations or launches according to certain policies. The resource allocation control policies were obtained following the SIM-OPT framework conceived by Varma,(2005), which used an architecture in which there are two loops. The inner loop contains the discrete event simulation of the development and commercialization activity network and a resource allocation MILP. The outer loop includes an observer that learns the optimal static policy, while the inner loop is run hundreds of times with the MILP as the resource allocation decision maker. The state space for each drug in the policy is defined as: {DS(i), NLEV(i), NHEV(i)}, where DS(i) = Development Stage of Drug i, NLEV(i) = Number of drugs having Lower Expected Value than drug i in the same development stage, and NHEV = Number of drugs having Higher Expected value than drug i in the same development stage. The control space is a vector that has as many elements as there are projects in the portfolio. Each element can take only one of 3 values, which correspond to upper, ML and lower resource allocation; that are associated with three different levels of duration for each task. Finally, it is relevant to mention that Varma (2005) explored the use of a MRCPSP MILP in the inner loop to consider the impact of a reactive schedule on top of the dynamic allocation of resources. However, it was shown that the difference between this more computational expensive approach and the allocation only MILP was not significant.

3.4. Case study

The same 9 drug portfolio case study reported in Blau et al (2004) is used. Most of the statistical distributions and parameters for the model were replicated. The only changes were in the duration and cost distributions, which were fixed at their most likely values. This was necessary due to the lack of data for the same distributions at higher or lower

resource allocation levels. The interactions between projects, as well as the simulation philosophy were also retained. Finally, the same financial, technical, manufacturing cost and resource dependencies for drugs aimed to diseases I and II were included. It is important to highlight that in spite of the simplifications in the discrete event simulation, the results are not significantly affected. The influence of fixing the duration and cost distributions is strongly dominated by the uncertainties retained in the model.

The base and extended case results for this study were obtained using the improved version of the GA with multiple α weights (0, 0.5, and 0.8), and a population size of 10. However the values of some of the parameters of the mutation operators were individually adjusted to prevent the algorithm from converging prematurely.

The resource allocation levels for the inner loop in the extended case were identified from suggestions provided by managers in the industry. The upper and lower levels correspond to using ±15% resources than the ML, while inversely changing activity duration by ±7.5%. Those flexibility values were doubled in a second run to better understand the implications of decisions made at the tactical level within different managerial frameworks.

4. Results

The return as measured by the EPNPV and the probability of losing money (portfolio risk) for the base case are presented in Fig. 2. All the points corresponding to the maximum EPNPV for a given level of risk are linked to form an approximate return-risk frontier. At first sight it looks like its shape reflects the general form found by Markowitz in financial portfolios (Luenberger, 1998), but a closer look reveals that the direct correlation between return and risk is violated in the middle section. The number of projects in the portfolios in that area is considerably higher than the number of those on the rest of the frontier. That demonstrates that the base case can not capture the trade off between the inclusion of more projects to reduce the level of risk due to failure, and the reduction in returns due to developmental delays, caused by exceeding the limits of available resources. Fig. 3 (only the dominating portfolios are plotted) shows that the addition of resource allocation flexibilities mitigates the effect of this trade off and pushes the frontier to higher returns for the same risk levels. This result can be explained from a conceptual point of view as follows. The lower or higher level of resource allocations constitute real options to delay or to expedite, which means that some of the flexibilities in the decision process where captured with a consequent rise in the value of the portfolios. However, the composition of the portfolios on the frontier in the base case and the extended one is remarkably different in the area where the depression is found. The extended case efficient portfolios can not be obtained by simply adding projects to the base case results. This demonstrates that it is not possible to decouple the strategic and tactical decision making processes at certain levels of risk. It was also found that the inclusion of flexibilities does not guarantee the improvement or sustainability of the performance of a specific portfolio. It was even found that the results of some interior portfolios from the base case completely dominate those observed in the extended cases. Most of the time the portfolio chosen at the strategic level is in the interior region, which means that its behavior is completely unpredictable based on aggregated quantitative or non quantitative methods. Finally, it is important to mention that the GA converges much faster in the extended cases than in the base case. Based on the progression of the algorithm we believe that such a behavior is due to the reduction in the search space. The inclusion of flexibilities decreases the importance of the position of the projects in the sequence. Therefore, what really matters in the initial

generations is the presence of the projects that in later generations will be reordered to obtain the optimal solution.

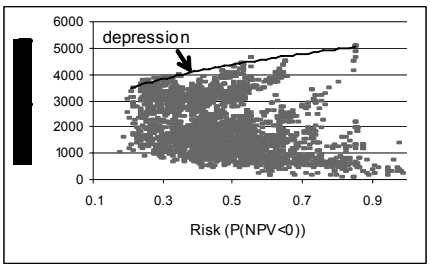

Fig 2. Efficient frontier base case

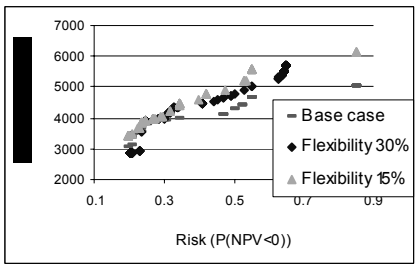

Fig 3. Efficient frontier extended cases

5. Conclusions

The study demonstrated the use of a multi-level Sim-Opt strategy to determine an optimal portfolio of new drug candidates. It has shown that the use of aggregated strategic decision models in the presence of uncertainties, project interdependencies and constraints can be misleading. It also has shown and quantified the importance of considering flexibilities in the valuation of projects. Although the approach is computationally demanding, it is an important tool to help management understand how tactical and strategic decisions are inter-related in order to maximize portfolio returns at specific levels of risk.

References

Baker, N.R., 1975. R + D Project Selection Models - Assessment. R & D Management, 5: 105-111.

Blau, G.E., Pekny, J.F., Varma, V.A. and Bunch, P.R., 2004. Managing a portfolio of interdependent new product candidates in the pharmaceutical industry. Journal of Product Innovation Management, 21(4): 227-245.

Calantone, R.J., Di Benedetto, C.A. and Schmidt, J.B., 1999. Using the analytic hierarchy process in new product screening. Journal of Product Innovation Management, 16(1): 65-76.

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.

Lin, C.H. and Hsieh, P.J., 2004. A fuzzy decision support system for strategic portfolio management. Decision Support Systems, 38(3): 383-398.

Loch, C.H. and Bode-Greuel, K., 2001. Evaluating growth options as sources of value for pharmaceutical research projects. R & D Management, 31(2): 231-248.

Luenberger, D.G., 1998. Investment science. Oxford University Press, New York, xiv, 494 pp.

Maravelias, C.T. and Grossmann, I.E., 2004. Optimal resource investment and scheduling of tests for new product development. Computers & Chemical Engineering, 28(6): 1021-1038.

Raynor, M.E. and Leroux, X., 2004. Strategic flexibility in R&D. Research-Technology Management, 47(3): 27-32.

Rogers, M.J., Gupta, A. and Maranas, C.D., 2002. Real options based analysis of optimal pharmaceutical research and development portfolios. Industrial and Engineering Chemistry Research, 41(25): 6607-6620.

Subramanian, D., Pekny, J.F., Reklaitis, G.V. and Blau, G.E., 2003. Simulation-optimization framework for Stochastic Optimization of R and D pipeline management. AIChE Journal, 49(1): 96-112.

Varma, V., 2005. Development of computational models for strategic and tactical management of pharmaceutical R&D pipelines. Ph.D Thesis, Purdue University.