R&D Project Portfolio Analysis for the Semiconductor Industry

Banu Gemici
Jeffrey T. Linderoth
S. David Wu
Lehigh University

Jeffry Moore
Agere Systems

Report No. 06T-005

R&D Project Portfolio Analysis for the Semiconductor Industry

Banu Gemici*, Jeffrey T. Linderoth*,
S. David Wu*, Jeffry Moore[1]

October, 2006

* Industrial and Systems Engineering Department, Lehigh University, Bethlehem, PA
bag4@lehigh.edu, jtl3@lehigh.edu, sdw1@lehigh.edu
[1] Agere Systems, Allentown, PA, jeffrymoore@agere.com

ABSTRACT

We introduce a decision support framework for the Research and Development (R&D) portfolio selection problem faced by a major U.S. semiconductor manufacturer. R&D portfolio selection is of critical importance to high-tech operations such as semiconductor and pharmaceutical, as it determines the blend of technological development the firm must invest in its R&D resources. This R&D investment leads to differentiating technologies that drive the firm's market position. We developed a general, three-phase decision support structure for the R&D portfolio selection problem. First is the *scenario generation phase* where we transform qualitative assessment and market foresight from senior executives and market analysts into quantitative data. This is combined with the company's financial data (e.g., revenue projections) to generate scenarios of potential project revenue outcomes. This is followed by the *optimization phase* where a multistage stochastic program (SP) is solved to maximize expected operating income (OI) subject to risk, product interdependency, capacity, and resource allocation constraints. To measure portfolio risk, the SP model uses a modified mean-Gini risk coefficient that estimates product risk based on the variability of profits. The optimization procedure generates an efficient frontier of portfolios at different OI (return) and risk levels. The *refinement phase* offers managerial insights through a variety analysis tools that utilize the optimization results. For example, the robustness of the optimal portfolio with respect to the risk level, the variability of a portfolio's OI, and the resource level usage as a function of the optimal portfolio can be analyzed and compared to any qualitatively suggested portfolio of projects. The decision support structure is implemented, tested, and validated with various real world cases and managerial recommendations. We discuss our implementation experience using a case example, and explain how the system is incorporated into the corporate R&D investment decisions.

1 Introduction

Effective management of the Research and Development (R&D) portfolio is critical to effective market positioning in high-tech industries such as computers, semiconductors, and pharmaceutical. In these industries, a firm's market position is tied directly to its portfolio of intellectual property (IP), which must be developed, acquired, or licensed. Studies of these industries (c.f., [4], [2]) point to the important conclusion that a firm's ability to attain significant market share in any technology area depends on its ownership of the essential IP and its ability to leverage essential IP from other firms.

In this paper, we describe specific issues of R&D portfolio selection and management from the perspective of a specialty semiconductor firm. We developed a decision support structure based on risk modeling and scenario analysis. The idea is to combine qualitative and quantitative insights concerning market competition, customer performance, technology interdependence, and the status of the firm's IP development and R&D resources in a cohesive framework. This, in turn, allows the management to systematically evaluate and refine their R&D portfolio, trading off potential returns against risk exposure. Another important dimension of the decision support structure is the integration of strategic and operational planning decisions. Key strategic decisions in this context include customer, market, and product selections as well as timing of the product release. Operational decisions support strategic decisions through resource adjustments such as capacity, personnel, and budget allocations.

The most notable examples of using portfolio selection and management concepts in strategic planning can be found in pharmaceutical companies. In this industry, the drug development portfolio drives the R&D budget (which can constitute over 40% of the total operating cost), determines the firm's market positioning, and impacts the firm's brand image. Major firms try to balance their portfolio of drugs such that high-risk, high-margin specialty drug development projects are mixed with low-risk, high-volume developments. When facing R&D portfolio decisions, firms are not only subject to risks due to market uncertainties, but also other risks such as prolonged FDA approval cycles and potential legal liabilities. Pharmaceutical product development is one example of an emerging trend in the high-tech arena—many products have a relatively short life cycle, product-development is extremely capital extensive, and the production can require a long lead-time. While our focus in this paper is on the specialty semiconductor industry, the methodologies developed are generalizable and transferable to other industries such as pharmaceutical.

Our work is based on a two-year project with a major U.S. semiconductor manufacturer. The company has an annual R&D budgeting process where senior executives go through an extensive review to determine in which technology development projects to invest and the amount of resources (e.g., capital, personnel, and prototyping capacity) to be allocated. Due to the large amount of market information that must be digested and analyzed and the multiple sources of uncertainty constantly present in the process, it has long been recognized that a more methodical, data-driven approach is needed to assist in this process. The decision results in the allocation of tens of millions of dollars of capital expenditure.

Portfolio management for R&D projects involves a great deal of uncertainty. In order to assess revenue impacts of a specific project, it is necessary to consider the full life cycle of the product, from the initiation of the product development, to product launch, all the way to the end of the product life cycle. In early stages of product development, the product specifications are not well defined, thus the level of capital expenditures for the R&D project are not precisely known. It is typically the manufacturer's responsibility to make initial investments and build prototype models (e.g., microelectronic components of a more complex device) upon customers' request. If the prototypes are not adopted by the customer, then the project terminates and the costs must be absorbed by the manufacturer. If the prototypes are adopted by the customer, market demand for the customer's product (which drives the manufacturer's demand) remains unknown. Market demands may be driven by a variety of factors such as product quality, timing of the product release, and intensity of competition in the market. Adding to the market demand uncertainty, rapid changes in technology decrease the duration of a product's life cycle, eroding the revenue streams. A microelectronic chip that faces high demand today might be obsolete in a few months, and its revenues might fall well below expectation. It is due to the highly volatile nature of this industry that the portfolio management issues become so critical. Moreover, as each portfolio selection decision could lead to exceedingly different outcomes, the *robustness* of the portfolio decision is of significant importance, i.e., it is more desirable to have portfolios that perform well over a wide variety of scenarios than ones that perform exceptionally, but only under restrictive conditions.

Although the technology market is both complex and volatile, there are some clear traces of technological trend driving this market. Moreover, senior decision makers at the firm share a wealth of insights on broader market trends, their technological ramifications, and potential impacts to customer demands. We discovered early in our project that it is important for the portfolio selection model to look beyond quantitative data and takes advantage of qualitative knowledge that is distributed among top decision makers. We designed the first phase of our Decision Support System (DSS) to collect relevant qualitative information, which is in turn converted into comparable measures and used to provide valuation for the products under consideration. The valuation of a particular product is constructed under a wide range of scenarios, which enables us to incorporate uncertainty effects and the variability of revenue streams.

While it is important that the DSS includes all critical factors influencing a portfolio's value, if the system demands excessive detail or uses logic to which the decision makers cannot relate, then management will not expend the necessary time to accurately fulfil the model requirements, and the DSS will be of little practical value. On the other hand, management wants assurance that the tools include all practically relevant details and produce rational results that they can decipher. This is a legitimate demand as when a large amount of quantitative data is processed without proper contextual information, misguided solutions often result (e.g., a particular project may not look very attractive on its own merits, but it is essential in securing the market or the confidence of an important customer). For these reasons, during the design of the DSS, we established close interactions with the top management, evaluated their concerns, and set the data requirement such that data is either easy to extract from human knowledge or exists in the company database. The former, or the qualitative data requirement, is closely related to each project's life cycle. In the next section, we will first describe the business context of the specialty semiconductor industry and a typical R&D project's life cycle.

2 Semiconductor Product Development

Semiconductor manufacturers undertake the wafer fabrication, packaging, assembly, and test responsibilities of microelectronic chips. In the specialty semiconductor industry, a majority of the chips are custom-designed to handle special functionality of an electronic device, be it a cell phone, an iPod, or complex telecommunication equipment. The electronic device manufacturers are major-brand carrying corporations who contract their chip production to the specialty semiconductor manufacturer. Depending on the contract type, the brand-carrying customer and the semiconductor manufacturer may share the responsibilities in new product development (NPD). The semiconductor manufacturer needs to choose the set of contract proposals from customers that best enhances the company's profitability. These contract proposals might be a renewals of existing projects or the development of a new product. A typical NPD project includes the following steps:

1. ***Evaluation of project proposals from customers:*** At the beginning of this stage, the customer and the semiconductor manufacturer aim to characterize the product. They discuss and identify the type of technology to be used, technical specifications, expected resource requirements, time to market, and the demand potential. If the project proposal is for a renewal of a previous project, then most of this information already exists. However, if the project proposal calls for NPD, typically very limited information is available. Typically for NPD, the semiconductor manufacturer must develop the product before entering production. At this stage, sources of uncertainty include resource requirements during product development, time to market, and demand potential. After evaluating the project proposals, the manufacturer must decide which proposals to accept.

2. ***Development Phase:*** If the semiconductor manufacturer commits to the project proposal, then the development phase starts. Throughout this phase, the two parties' design teams work together to identify the functional needs of the chips. The duration of the development phase depends on the complexity of the design, the skill level of the R&D staff, and the availability of relevant prior technology. If prior technology needs to be extended, or new technology needs to be invented, then investments must be made to develop the essential intellectual property

(IP). In some cases, IP can be licensed or acquired from a third party. A completed design leads to the development of prototypes. If the prototypes meet the specifications, then the development phase continues as scheduled; otherwise the date of the product release must be delayed, and more investment is needed to develop the IP. Once the product is successfully prototyped and tested, the customer chooses to accept or reject the product design. In case the customer rejects the design, the project is terminated, and majority of the development costs are borne by the semiconductor company. This is known as a *loss of design win* and is a major source of uncertainty for the semiconductor manufacturer.

3. ***Production Ramp:*** When the customer accepts the product design, the production phase starts and resources are allocated for the production. Now, market and capacity-related risks dominate. Market risks are influenced by both external factors, such as the customer's marketing efforts and the nature of market competition, and internal factors, such as the product quality. Capacity risks arise from the uncertain nature of available resources. These factors collectively affect the expected revenue streams. If early sales or other market signals suggest that the revenues of a certain product is unlikely to recuperate the production and the development costs, the product might be pulled from the market so that resources can be allocated to more profitable projects.

Figure 2 illustrates a typical product life cycle for a semiconductor manufacturer. Each product to be considered for the R&D portfolio follows a similar life cycle. The duration of each phase in the life cycle and the quantity produced during each phase depend on various (unknown) event realizations, but decisions on a particular project proposal must be made at the beginning of the product life cycle, in anticipation of the uncertainties.

While uncertainties in the product life cycle represent a major challenge to the portfolio decisions, product relationships and dependencies also have to be taken into account. The relationship between products might be in different forms. For example, a new project proposal might make a set of existing products obsolete, or competing customers might offer similar proposals. In either case, these are indications that certain projects cannot co-exist in the same portfolio. Later, we refer to these as *mutually exclusive projects*. Projects may also have *prerequisite relations*. For example,

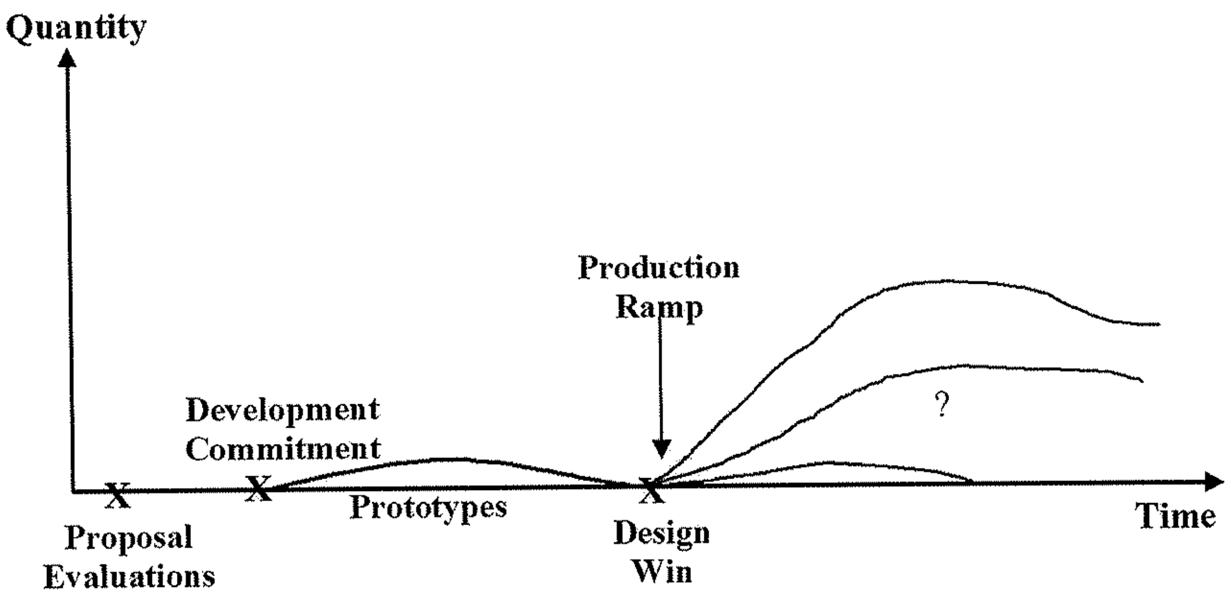

Product Life Cycle

projects may create IP that is necessary for the development of another set of projects. There are often other complex relationships between semiconductor R&D projects. For example, it is possible to have a prerequisite project of a key project that happens to be mutually exclusive with another key project in the portfolio. Portfolio decisions must be made considering all project dependencies, leading to complex trade-offs. Moreover, uncertainties associated with a project are often amplified when the project is considered together with its prerequisite projects. For example, the market release time of a project depends on the successful completion of its prerequisite projects. Thus, uncertainty in the completion time of prerequisite projects compounds the uncertainty associated with the completion of a dependent project. As a result, the valuation of a dependent project may be subject to a much higher level of variability when considered together with its prerequisite projects. This leads to the conclusion that the value and variability of the project portfolio *as* a whole must be considered by management. In the DSS designed for the company, the scenario analysis tool takes into account all inter-project relationships.

The rest of the paper is organized as follows. We summarize relevant literature for project/product portfolio management, followed by the problem statement and model formulation. We then introduce the three-stage decision support structure for project portfolio management, and we describe the implementation and use of this DSS at a particular semiconductor firm. We discuss the robustness of the tool, the comparison against current methods, and key feedbacks from the management. The paper finishes with conclusions and directions for further research.

3 Related Literature

The concept of Product Portfolio Management (PPM) emerged in late 1950's. Through 1970's, it became an established planning tool. In the 1980's and 1990's, the use of portfolio management was extended into the area of new product development(NPD) and R&D projects evaluation. The tools and methods for PPM have evolved through time, but the key objectives remained the same, as listed by [6]:

1. Maximize the value of the portfolio
2. Achieve the right balance of projects, and
3. Achieve a strategically aligned portfolio

Methods used for PPM range from qualitative such as unstructured peer review, to various quantitative techniques. The latter includes approaches ranging from mathematical programming and portfolio optimization, to economic models using (*internal rate of return (IRR), net present value (NPV), return on investment (ROI)*, decision analysis tools such as *multi attribute utility theory (MAUT), decision trees, risk analysis, analytical hierarchy process (AHP)*), and interactive methods such as emphDelphi, Q-sort, behavioral decision aids (BDA), decentralized hierarchical models (DHM). Specific examples of these approaches are too numerous to name.

The success of each method depends on the business environment, the information requirements, the users' understanding of the techniques, and the buy-in from senior management. Each method has its own pros and cons. For example, mathematical programming methods rely heavily on data, and capturing that data might not be feasible in the business environment; scoring techniques might not be sufficient to capturing the complexity of product dependencies. In the last decade, use of *Decision Support Systems* (DSS) became popular in PPM. DSS are interactive and computer-based systems to help decision makers in their decision process, generally by combining multiple methods into an integrated system. Hereafter, we focus on main components of project selection process and different applications of these components in PPM literature.

In PPM, decision makers' main task is to quantify the value of the projects, because the selection of a project at the beginning, or deciding to keep the project in the portfolio in the middle of the project's life cycle is generally based on the project's value and its rank. However, in most of the business environments, the projects are not independent. For example, some projects may share common resources, or there might be prerequisite relation between them. Thus, value of a project depends on its relation with the other projects, and depending on company strategies and the business environment, multiple criteria are considered to assess a project's value. AHP is one of the most frequently used techniques to represent these values. It is a systematic method for comparing list of objectives or alternatives. [8] use AHP to incorporate qualitative criteria and use priority rankings of AHP to represent a measure of value for each project. [18] use scoring techniques to rank the projects and identify the projects that are worthy of further evaluation. [7] use a dependency matrix to quantify the interdependencies between projects. Each element in the dependency matrix, d_{ij}, varies from zero to one indicating the level of dependence that project i has on project j for its financial success. They use d_{ij}'s to find how much value of the project is attributable to itself and how much is attributable to its interdependencies. After the seminal work of [11], versions of Mean-Variance (M-V) models became popular in stock portfolio analysis. In M-V analysis, the key component is covariance matrix of the stock returns. It measures the extent that a stock's return depends on the other stocks. Because of the difficulty in measuring these dependencies, M-V Analysis has not been widely used in PPM. [5] develop a mean variance model to determine the optimal toll in a built-operate-transfer roadway project under traffic demand uncertainty. [12] discuss the application of Markowitz's modern portfolio theory to upstream decision making in the oil and gas industries. In our case, the value of each project is determined by the events related to the project's life cycle and the effect of these events on the project's revenues and costs. By considering the project's life cycle, we include the relationship between the project and its dependencies. Each event is associated with a probability and an effect on the project's value. We calculate the value of a project under different scenarios, which are constructed from the combination of the project's life cycle events.

Incorporating risk in PPM decisions adds one more dimension to the characterization of the project's value. When risk is a part of the portfolio analysis, often the risk parameters are considered to be given. However, taking risk as given would be far from being realistic. The complication of including risk in decision process is that it is hard to obtain a reliable risk measure. [1] propose a set of risk-measure axioms and declared the risk measures that satisfy these axioms as "Coherent Measures of Risk". According to them, a coherent measure of risk should satisfy 1) Translation invariance, 2) Subadditivity, 3) Positive homogeneity, and 4) Monotonicity. M-V models aim to minimize the variance of the stock portfolio return while maintaining some target return level. However, as mentioned above, use of covariance matrix in PPM is not as practical as in stock portfolio analysis. Another point is that the indivisibility of the projects makes the M-V analysis in PPM different than M-V analysis in stock portfolio. Although, variance is a widely known measure of variability, it does not satisfy the subadditivity axiom of coherent risk measures, hence it is not a coherent measure of risk. [15] present a Mean Gini approach for analyzing risky prospects and construct optimum portfolios. They show that M-V analysis is a special case of mean Gini approach. [16] use Mean Gini Risk to measure the extent that each prospect's return moves with a given portfolio's return, and they identify dominant prospects based on Marginal Conditional Stochastic Dominance. Mean Gini Risk measure satisfies the four axioms of coherency. [13] extend the method of [16] to R&D portfolios. In our analysis, we use Mean Gini Risk measure as in [13] to calculate the risk of each project among all the candidate projects. However, they analyze a given portfolio in one period, while we consider multi period planning horizon. Thus, in our case, the risk of a project is different from period to period. Moreover, we construct the portfolio with a mathematical model while they use stochastic dominance criterion to determine dominant projects given an existing portfolio.

The portfolio selection is made such that the portfolio increases total value of the company most. Depending on the business structure, the priority on company objectives might be different. Often there is more than one objective and the number of candidate projects might be too large. In this situation, generally mathematical models are used in DSS to select the best portfolio. For example, [8] use 0-1 integer model to construct the portfolio. They maximize the portfolio value with a budget constraint. The portfolio value is the sum of selected projects' priority rankings that are determined by AHP. They also extend their model with additional balancing constraints. [7] solve a nonlinear integer model to maximize the Net Present Value (NPV) of the portfolio with budget and strategic alignment constraints, and constraint that sets limit on the number of selected projects. [18] solve multi-objective integer linear model to find pareto efficient portfolios. Then decision makers select the best fit portfolio among the pareto efficient portfolios by setting aspiration levels and upper or lower bounds for certain objectives. In our DSS we solve multistage SP to maximize the expected value of the portfolio. The value of the portfolio is calculated by taking the weighted sum of selected projects' values. We also consider operational implication of the portfolio by adjusting resource levels, and reflect the cost of these operational decisions in the objective function. In addition to these, the selection of projects are limited by total return variability, and there are strategic constraints that mandate some projects to be selected or not selected depending on decision makers' choice. SP allows us to consider uncertainty in decision making. While the decisions are made in the presence of uncertainty, the decision makers are able to adjust their decisions after the realization of uncertain events. In general, SP models are hard to solve because of their sizes. However, with the availability of inexpensive computer power and sophisticated solvers, SP methods are increasingly popular. SP is used in the optimization component of our DSS, and the decision makers do not need to solve the problem by themselves. Our user friendly program produces the results in terms of charts, visual aids and tables, after getting initial directions from the decision makers. Although, we designed our DSS to address the needs of the semiconductor manufacturer, after some adjustments, our DSS can be used for any kind of business environment that suffers from high level of uncertainty.

4 Model Formulation

4.1 Justification

Prior to our study, the management had employed two different techniques to aid their PPM decisions. The first approach was to rank the projects based on NPV of the cash flows. A budget constraint limited the total investment, and heuristic strategic alignment criteria were used to narrow the number of allowable projects. This methodology was able to select projects with strategic considerations in mind, but the assessment and incorporation of risk into the decision process was qualitative and subjective.

In their second approach, the management worked on a M-V model. The interdependencies between projects were quantified using several factors, and selection of these factors and subsequent determination of the covariance matrix were the critical keystones of this approach. In practice, direct application of this method was impeded by not having the depth of data available in a classic stock portfolio problem. The required data for this approach was not captured in the routine transaction data of normal business operations such as orders and invoice data, bill of materials, and other supply chain transactions. Moreover, new projects had little or no historical data on which to base and compare the risk statistics. A final critique of the M-V model was that it was not dynamic enough to account for the fast rate of changes of the semiconductor industry.

Based on these previous attempts, we had the following design requirements for our PPM decision support tool:

1. use data available from routine business transactions,

2. dynamically account for changes in the business environment,

3. incorporate a risk measure for projects in the portfolio, and

4. include long term (strategic) and short term (operational) considerations.

In the light of these design criteria, we decided to use a SP approach to model the problem. First, the SP model can easily be built from scenario data generated from routine business transactions. Second, in SP models, decision makers are able to change their decisions as they learn new information in a dynamic environment. Third, a risk measure can easily be incorporated into the model with he inclusion of explicit scenarios. Last, strategic and operational aspects would be easily included by logical constraints into the mathematical model.

4.2 Data and Uncertainty

As part of the renewed project management effort, the company keeps basic financial figures, customer information, and technical product information in their database. At any point in time, the company is considering a set of projects P that will consume resources from a set R. Each project $p \in P$ is associated with a set of prerequisite projects $Q_p \subset P$ and with a set of mutually exclusive projects $E_p \subset P$. Project p can be undertaken if and only if all projects in Q_p are undertaken and none of the projects E_p are done. Forecast cost and resource information is kept at a project level for quarterly intervals. For each project $p \in P$ the following data is available for each period t in a specified time horizon $\mathcal{T} = \{1, 2, \ldots, T\}$.

- Forecast gross margins of the project p at period t: GM_{pt},
- Forecast fixed costs of the project p in period t: FC_{pt},
- Forecast usage of resource r in period t by project p: W_{rpt}, and
- Periodic unit cost of each resource r: c_r.

As new information is available forecast figures are updated. The management is aware of the need for accurate information, and they believe the existence of good information is necessary for making good portfolio choices in a systematic and sustained fashion.

Forecast figures reflect the current status of the business environment, but these figures change as new information is available or as the uncertainties resolve. From the analysis of the project life cycle and through discussions with management, we identified that a new business state can be modeled by adjusting GM_{pt}, FC_{pt} and W_{rpt} based on random events that occur during the project's life cycle. That is, these quantities are *functions* of some random variables. To fix notation, let $\Omega_t, t \in \mathcal{T}$ be the set of all possible random events that occur during period t. The set of all sequences of events is then $\Omega \stackrel{\text{def}}{=} \Omega_1 \times \Omega_2 \times \ldots \times \Omega_T$. We use the common notation $\Omega_{[1,t]}$ to denote the set of all sequences of events that can occur from stages 1 to t. The dependence of gross margin, fixed cost, and resource usage on random variables is then denoted by referring to these quantities as $GM_{pt}(\omega_{[1,t]})$, $FC_{pt}(\omega_{[1,t]})$, and $W_{rpt}(\omega_{[1,t]})$, where $\omega_{[1,t]} \in \Omega_{[1,t]}$. We discuss in Section 5.3 the exact manner in which the quantities are derived, but for purpose of model discussion it suffice to know that for each project and resource we have all the necessary information to assess project's value and determine resource requirements under different sequences of realizations of random events.

4.3 Decision Variables

In stochastic programming, decisions are made in stages, and in-between stages information about the state of the business becomes available. In general the decision process has the form:

$$1^{st} \text{ stage decisions } \rightarrowtail \text{Information } \omega_2 \rightarrowtail 2^{nd} \text{ stage decisions } \rightarrowtail \ \ldots$$

$$\ldots \rightarrowtail \text{Information } \omega_T \rightarrowtail T^{th} \text{ stage decisions}$$

To ease the exposition, we will assume that there is one stage in the stochastic programming model for each period in the decision making process. In the actual decision support system developed for the company, several periods might be aggregated into one stage. As described in Section 4.2, the data defining the decision model are functions of random variables. The decision variables in the model are also functions of the random events that might occur. There are two classes of decisions that the company must make: strategic and operational.

The strategic decisions are modeled with the variables $x_{pt}(\omega_{[1,t]})$, indicating whether or not project p is to be included in the portfolio in time period t. An important characteristic of the model is that $x_{pt}(\cdot)$ is solely a function of the random variables that occur from periods 1 to t. The decision of whether or include p in the portfolio at t is *non-anticipative* of random events that occur after period t. At this point, we model the nonanticipativity implicitly by simply stating $x_{pt}(\cdot)$ is a function only of $\omega_{[1,t]}$. Algorithmic mechanisms for enforcing nonanticipativity will be described in Section 4.5. Associated with each project $p \in P$ is a *begin time* $b_p \in \{1, 2, \ldots, T\}$, which indicates the beginning time of project's life cycle). If a project is not selected at the beginning of its life cycle, then it cannot be selected later. However, if a project is selected at the beginning of its life cycle, it can be killed later.

Completion of a project requires resources to be allocated to the project. The main resources necessary in the model for the company are human resources. As such, there is a cost associated with increasing the level of resources (hiring) and with decreasing the level of resources (firing) from period to period. There are two types of human resources that affect project completion: design team members and administrative staff. The initial available level of resource r is defined by I_r. There are three types of operational decisions, each dependent on the random events that occur during the PPM process:

$y_{rt}(\omega_{[1,t]})$: Total available level of resource r at the end of stage t, under scenario $\omega_{[1,t]}$. It is the sum of previous stage resource level and current period adjustments (increase or decrease in the resource level). The periodic cost of keeping one unit of resource r is given by c_r. At the end of first stage $y_{rt}(\omega_{[1,t]})$ becomes y_{r1} because of the perfect information at the first stage, and y_{r1} is the results of first stage adjustments over I_r.

$h_{rt}(\omega_{[1,t]})$: Amount of increase in the resource level r during stage t, under scenario $\omega_{[1,t]}$. There is a cost associated with the unit increase of resource level r described by η_{rt}. Level of increase in the first stage is given by h_{r1}.

$f_{rt}(\omega_{[1,t]})$: Amount of decrease in the resource level r during stage t, under scenario $\omega_{[1,t]}$. There is a cost associated with the unit decrease of resource level r described by ζ_{rt}. Level of decrease in the first stage is given by f_{r1}.

4.4 Project Risk Measure

The model also includes a mechanism for controlling the company's risk in choosing collections of projects to include in the portfolio. The mechanism is based on calculating "risk coefficients" for each project in each stage. The risk coefficients are based on mean-Gini considerations, whose exact calculation is detailed in Section 5.3.1. For purposes of the model discussion, it suffices to know that for each project and time period a coefficient λ_{pt} is computed that quantifies the risk of project p in time period t. The total risk of the portfolio is then taken to be a linear combination of the individual risk coefficients chosen to be included in the portfolio. It should be noted that this is not the *true* mean-Gini measure of risk of the chosen portfolio. In fact, the mean-Gini risk is a *nonlinear* function of the chosen projects, and modeling this consideration exactly led to an intractable model. Justification of the linear combination of risk coefficients can be taken from the fact that the company was happy to have an adjustable measure to control the risk they took on, and the solutions obtained by the model fit the company's notion of risk perfectly well.

4.5 Model Objective and Constraints

Given the definition of problem data and decision variables, we can write an optimization model for the project portfolio problem as follows. Our objective is to maximize the total expected operating income:

$$\max \mathbb{E}_\omega \left[\sum_{t \in \mathcal{T}} \sum_{p \in P} (GM_{pt}(\omega_{[1,t]}) - FC_{pt}(\omega_{[1,t]}))x_{pt}(\omega_{[1,t]}) - \sum_{r \in R} (c_r y_{rt}(\omega_{[1,t]}) + \eta_{rt} h_{rt}(\omega_{[1,t]}) + \zeta_{rt} f_{rt}(\omega_{[1,t]})) \right] \tag{1}$$

The first summation of Equation 1 represents the total net profit associated with doing projects, and the second summation accounts for the operating expenses of adjusting the resource levels.

Most constraints of the model contain random variables, and we enforce these constraints in a probabilistic sense by saying that the constraints will hold with probability one, or *almost surely*, represented by the notation *a.s.* in the constraints of the model.

The required level of each type of resource should be enough to continue selected projects at all stages and can be described by the following constraint:

$$\sum_{p \in P} W_{rpt}(\omega_{[1,t]}) x_{pt}(\omega_{[1,t]}) \leq y_{rt}(\omega_{[1,t]}) \qquad \forall r \in R \ t \in \mathcal{T}, \text{a.s..} \tag{2}$$

The headcount level for a type of resource at the end of a specific stage is the sum of headcount level in the previous stage and number of people hired during the current stage minus the number of people fired at the current stage. For the first stage, headcount level is adjusted over the initial number of headcount level, I_r.

$$y_{r1} = I_r + h_{r1} - f_{r1} \qquad \forall r \in R \tag{3}$$

$$y_{rt}(\omega_{[1,t]}) = y_{r,t-1}(\omega_{[1,t]}) + h_{rt}(\omega_{[1,t]}) - f_{rt}(\omega_{[1,t]}) \qquad \forall r \in R,\ t \in \mathcal{T} \setminus \{\infty\}, \text{a.s.} \tag{4}$$

For each project $p \in P$, the requirements imposed by its prerequisite and mutually exclusive project sets must be imposed: When a project p is selected at stage t, all the projects in Q_p should be also in the portfolio at stage t. This is described by equation 5 below. When a project p is in the portfolio, none of the products in set E_p should be in portfolio as described by Equation 6 below:

$$
\begin{aligned}
x_{pt}(\omega_{[1,t]}) &\leq x_{qt}(\omega_{[1,t]}) \qquad \forall p \in P,\ q \in Q_p,\ t \in \mathcal{T} \text{ a.s.} \qquad (5)\\
x_{pt}(\omega_{[1,t]}) + x_{lt}(\omega_{[1,t]}) &\leq 1 \qquad \forall p \in P,\ l \in E_p,\ t \in \mathcal{T} \text{ a.s.} \qquad (6)
\end{aligned}
$$

We can only select a project p to perform when the stage of the decision coincides with the beginning period of that project's life cycle b_p. When a product is not selected at the beginning of its life cycle then it cannot be selected later. We ensure these restrictions by the following constraints:

$$
\begin{aligned}
x_{pt}(\omega_{[1,t]}) &= 0 \qquad \forall p \in P \forall t \in \{1, \ldots b_p - 1\},\ \text{a.s.} \qquad (7)\\
x_{p,t+1}(\omega_[1, t+1]) &\leq x_{pt}(\omega_{[1,t]}) \qquad \forall p \in P, \forall t \in \{1, 2, .., T-1\} \qquad (8)
\end{aligned}
$$

As explained in Section 5.3.1, there is a risk coefficient $\lambda\left[GM_{pt}(\omega_{[1,t]}), FC_{pt}(\omega_{[1,t]})\right]$ for each product p at each stage $t, (t \in \mathcal{T} \setminus \{\infty\})$. The company is interested in limiting their total risk to a level K in the following fashion:

$$
\sum_{t \in \mathcal{T} \setminus \{\infty\}} \sum_{p \in P} \lambda\left[GM_{pt}(\omega_{[1,t]}), FC_{pt}(\omega_{[1,t]})\right] x_{pt}(\omega_{[1,t]}) \leq K \qquad (9)
$$

Again, an important consideration in the model is that the strategic and operational decisions made at a stage t are *independent* of the random events $\omega_{[t+1,T]}$. These nonanticipativity constraints are given implicitly here by defining the $x_{pt}(\cdot)$ as functions of the proper arguments $(\omega_{[1,t]})$. Algorithmic methods for enforcing this nonanticipativity will be discussed in the next section.

To make the optimization model (1—9) tractable, there must be some reasonable approximation of the uncertainty Ω. We chose a sample-path, or sample-average approach, in which set Ω is replaced by a randomly sampled approximation consisting of a finite number of *scenarios*. This approximation, or scenario generation, forms the first phase of our three phase DSS that we describe in the next section. The sample average approach has been used as of late on a variety of practical planning problems, including: [14] [19]. Further, recent theoretical and empirical evidence suggests that an accurate answer to the true problem can be obtained by approximating the uncertainty set with a surprisingly few number of scenarios. ([10], [17]) The use of scenarios also enables us to enforce the nonantcipativity of decisions by only creating decision variables that can depend on the appropriate scenarios and previous decisions.

5 A Three-Phase Decision Support System

Rarely is a real-life problem situation simply "solved" by applying an optimization model. Rather, the road from real-life problem to working solution is an iterative process. The model we propose in Section 4 is no exception; it is part of a larger decision support system (DSS) now in place at the company. Figure 1 depicts the various components and stages of the DSS, as the user is guided through the project portfolio planning process. Note that in order to instantiate an instance, user input is required both *before* and *after* data is drawn from the company's database.

5.1 Problem Specification

Initially, decision makers specify the planning horizon (in fiscal quarters), the set of projects to be analyzed (P), the customer set (C), and the set of target markets (M) for these candidate projects. One individual project may be split over multiple customers and market segments. The decision maker defines the exposure of project p to its markets m as a fraction α_{pm} and the exposure of each project to its customers c as a fraction γ_{pc}.

Once the initial instance specification data is obtained from the users, the DSS queries the company database to retrieve the resources (R) consumed by the projects in set P, the periodic unit cost of each resource (c_r), the set of technologies used in the product development phase (E), and the forecast gross margins (GM_{pt}), fixed costs (FC_{pt}), and resource usage (W_{rpt}) for the specified projects, resources and periods. Further, all important time-line information about the project life cycle, such as design-win date, are obtained.

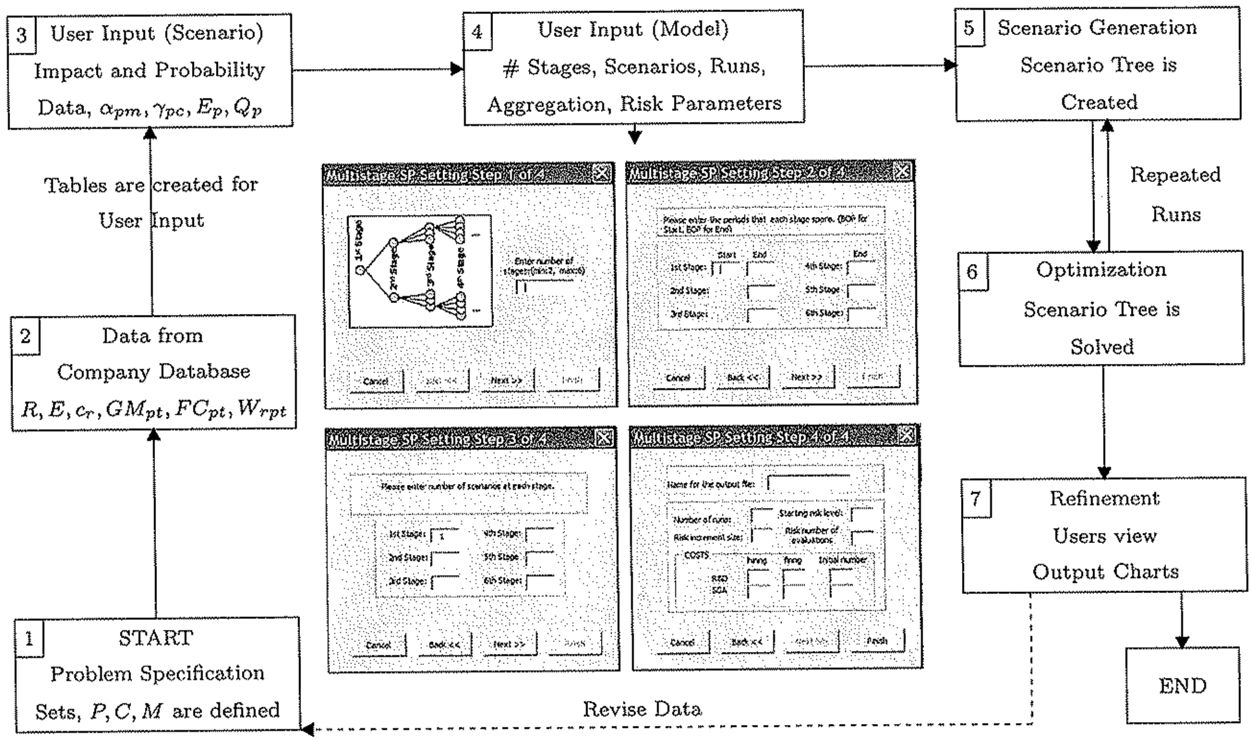

Figure 1: Event Flow Diagram of the DSS

This information is used to help describe a scenario, an instance of all possible sequences of random events relating to projects, resources, and technologies from the beginning until the end of the planning horizon. The set of all possible scenarios is denoted as $\Omega_{[1,T]}$. From analysis of the project life cycle and feedback from project managers, we concluded that seven *Risk Groups* were sufficient to accurately describe $\Omega_{[1,T]}$. The risk groups relate to

1. Product Performance,
2. Resource Performance,
3. Technology Status,
4. Intellectual Property (IP) Development Status,
5. Design Win Status,
6. Market Performance, and

7. Customer Performance.

The risk groups and possible outcomes of events in each risk group are the same regardless of the instance being considered. What varies from instance to instance is the *impact* associated with each outcome. Each outcome is associated with an impact parameter (ρ) that will be used to adjust forecast data. The outcome of events in one risk group may be correlated through time or correlated with outcomes of events from different risk group. These event relationships, the definitions of the risk groups, and the impact factors for each group will be described in the next section.

5.2 User Input

A key characteristic of the DSS in place at the company is that the project managers and high-level executives are responsible for quantifying the impact of various random events that may occur during the planning horizon. This domain expertise is spread among various individuals in the company. For example, an executive in charge of market segment m is best-positioned to make the assessment that if market m performs "well", then the company should expect to see 40% higher return on projects associated with that market. Similarly, project managers who work closely with specific large clients are able to assess that if that client c does "poorly", then the impact on return for projects associated with client c will be -20%. Further, the project manager may be able to assess that the likelihood of client c doing poorly is roughly 5%.

While the impact coefficients and associated probabilities certainly may seem arbitrary, it is important to note that they are *used* in a systematic manner and obtained from people within the company who are best-poised to provide such information. The DSS is a leap forward for the decision makers in the company who routinely did "what-if" analysis of various scenarios, but lacked the machinery necessary to properly capture the correlations between various events and to act intelligently in the face of uncertainty. The new DSS overcomes both of these obstacles: a simulation is used to ensure that all events are correlated through time, and an optimization model is used to simultaneously consider many different uncertain scenarios to choose the best action.

Typically, to create an instance, strategy meetings among executives and individuals with domain knowledge are held, and impact factors and probabilities are assigned for each event in each

Table 1: Risk Groups and Possible Outcomes

RISK GROUPS	OUTCOMES				IMPACT
Product Performance	Superior Performance	Expected Nominal	Poor Performance	Product Failure	$\rho_{pt}^{PP_{gm}}, \rho_{pt}^{PP_{fc}}$
Resource Performance	Over Performance	Expected Nominal	Under Performance	Loss of Key Resources	ρ_{rt}^{R}
Technology Status	Development Goes Well	Nominal	Behind Schedule	Failure	ρ_{et}^{E}
IP Development Status	Successfully completed		Could not be completed		Late Time to Market or continue
Design Win Status	Got the Design Win		Could not get the Design Win		Zero GM_{pt}'s or continue
Market Performance	Market Expands	Expected Nominal	Market Contracts	Market Collapses	ρ_{mt}^{M}
Customer Performance	Superior Performance	Expected Nominal	Poor Performance		ρ_{ct}^{C}

risk group. Another unexpected and pleasant benefit from the DSS in the company is that they have found that it has focused the discussion among executives and project managers of the impact of various events for various long-term plans. The risk groups and their probable outcomes are summarized in Table 1.

- Product Performance: Events related to the quality of a product. For each product $p \in P$ and for each stage $t \in \mathcal{T}$, the product performance may take one of the four possible outcomes in Table 1. The impact of the product performance on the fixed cost FC_{pt} is denoted by $\rho_{pt}^{PP_{fc}}$ and its impact on gross margin GM_{pt} is denoted by $\rho_{pt}^{PP_{gm}}$. The outcomes of product performance can affect the likelihood of the "Design Win" event if the project is in the development phase.

- Resource Performance: Events related to the performance of R&D and administrative staff. There is one outcome for each resource $r \in R$ and for each stage $t \in \mathcal{T}$. The outcome has an impact ρ_{rt}^{R} on the forecast resource usage rate W_{rpt}. If resources are lost in a stage, then projects using that resource cannot be completed unless new resources are obtained.

- Technology Status: Events related to the timing of the technology and project schedule. Each project is built using one type of technology, and there is an outcome for each technology $e \in E$

and for each stage $t \in \mathcal{T}$. The technology impact ρ^E_{et} is used update the probabilities of IP development status events.

- IP Development Status: A binary and one time event indicating successful IP development for project p. Failure of the IP development delays the time to market of the project and hence all other projects for which p is a prerequisite. The delay of the market release time is handled by shifting the gross margins GM_{pt} further out in the planning horizon. The number of stages to shift the release and any increase in development cost FC_{pt} is specified by the users. The probability of completing the IP development at time t, π_{IPt}, depends on the outcome of the previous period's technology status.

- Design Win Status: A binary and one time event indicating if the customers for the project are satisfied with the prototype design. If not, the project and all its dependent projects are killed. Otherwise, the product is released to the market and production phase starts. The probability of a design win at time t, π_{DWt} depends on the previous period's product performance outcome.

- Market Performance: Events related to market condition. There is one outcome for each market $m \in M$ for each stage $t \in \mathcal{T}$. A product might serve more than one market. In this case, the percentage exposure of product p to market m, α_{pm}, is used to obtain the overall effect of the market related outcomes. The outcomes have an impact ρ^M_{mt} on the gross margins GM_{pt} and are independent of outcomes of all other events.

- Customer Performance: Events related to customer performance. There is one of three potential outcomes for each customer $c \in C$ for each stage $t \in \mathcal{T}$. A product might serve more than one customer. In this case, the percentage exposure of product p to customer c, γ_{pc}, is used to obtain the overall effect of the customer related outcomes on instance parameters. The customer performance outcomes have impact ρ^C_{ct} on the product gross margins GM_{pt} and are independent of the outcomes of all other events.

After the probability and impact data is collected, the product-market (α_{pm}) and product-customer (γ_{pc}) exposures are provided by the decision makers. Last, project dependencies are specified in terms of prerequisite set Q_p and mutually exclusive set E_p for each project p. These sets are taken from the database, but at this point the user may over-ride the information from the database.

Data collection for large-scale problems is always a burden. However, the user-friendly interface of the DSS greatly eases this burden. The DSS is connected to the company database, so the forecast data for the project, technology and resource sets over the planning horizon are obtained instantly. The DSS generates tables to gather probability and impact data, and when the data is collected, error-checking routines ensure that all the data are present and fall within reasonable nominal ranges.

Even though the space of outcomes has been discretized, so that $\Omega_{[1,T]}$ is a finite set, the cardinality of this set is too large for us to consider *all* possible combinations of outcomes. To create a tractable instance of our math programming model, we sample from the set $\Omega_{[1,T]}$. To create an instance of the mathematical model that can be solved, the user needs to specify the number of stages of the stochastic programming model, the number of random event realizations at each stage, and information on how the fiscal quarters (periods) are spanned by stages. DSS gathers model data through a four-step wizard. (See Figure 1 Step 4). In the first step, users specify the number of stages (T) of SP model. (Maximum 6 stage model is allowed). In the next step, users determine how periods are spanned by stages, by specifying the beginning and end period for each stage. Then system aggregates the periodic data into stage data, by summing GM_{pt}, FC_{pt}, and W_{rpt} over the periods in that stage. In the third step, users specify number of random event realization at each stage $M_t(t \in \mathcal{T})$, for a $\prod_{t\in\mathcal{T}} M_t$ scenario model. At the last step, users specify number of runs (number of scenario trees to be solved) and risk parameters. Each scenario tree is solved for different risk levels (see Eq: 9). The starting, ending and the increments of risk level is specified by the users. When above information is acquired, the DSS checks, through an error check routine, whether all the data required for scenario generation is complete and logical. If the information is error free, then DSS goes to Scenario Generation Phase, otherwise users need to fix the error.

5.3 Scenario Generation

Conditional sampling is used to create a manageable sample of the discretized scenario space $\Omega_{[1,T]}$. The size of the scenario tree depends on the number of stages (T) and the number of random event realizations at each stage $t(M_t)$. The conditional sampling procedure works by first selecting a random sample of size M_2 for the second stage. The probability of each realization in the second stage is $1/M_2$. The second stage realizations are given by

$$\zeta_2^i = \left(GM_{p2}(\omega_2^i), FC_{p2}(\omega_2^i), W_{rp2}(\omega_2^i) \right), \ i = 1, ..., M_2.$$

Formulae for dependence of gross margin, fixed cost, and resource consumption $(GM_{pt}(\cdot), FC_{pt}(\cdot)$, and $W_{rpt}(\cdot)$ on the outcome of the random event ω^i are given subsequently.

Next, for every $i \in \{1, ..., M_2\}$ a random sample of size M_3 is generated. Thus, there are M_2M_3 realizations of events in the third stage, each with probability $1/(M_2M_3)$. The third stage realizations given the history of events up to this point $\omega_{[1,2]}$ are given by

$$\zeta_3^{ij} = \left(GM_{p3}(\omega_3^j|\omega_{[1,2]}^i), FC_{p3}(\omega_3^j|\omega_{[1,2]}^i), W_{p3}(\omega_3^j|\omega_{[1,2]}^i) \right), \ i = 1, ..., M_2, \ j = 1, ..., M_3.$$

The procedure is called conditional sampling due to the dependence of ζ_3^{ij} on ζ_2^i.

The procedure continues in this fashion until the T^{th} stage realizations given a history of events $\omega_{[1,T-1]}$ are generated. At the T^{th} stage, we have $M_2M_3 \cdots M_T$ realizations of events, each with equal probability $1/(M_2M_3 \cdots M_T)$. Each path from root node to leaf node in this tree is a *scenario*.

Conditional sampling is required to capture the correlations between random events. The events may be correlated through time and within risk groups. The relation between the outcomes is shown in Figure 2. For example, the figure depicts that the outcome of the product performance event in period t will affect the outcome of the design win status event in period $t+1$. The figure also shows the dependence of the gross margins, fixed costs, and resource consumption on the outcome of random events. For example, the outcome of the customer performance event in period $t+1$ will affect the gross margins in that stage. The exact functional dependence is given in Equation (10).

$$GM_{pt}(\omega_t^i|\omega_{[1,t-1]}) = \rho_{pt}^{PP_{gm}} GM_{pt} \left[\sum_{m \in M} \sum_{c \in C} (\alpha_{pm}\rho_{mt}^M)(\rho_{DW}^C \gamma_{pc} \rho_{ct}^C) \right] \quad \forall p \in P, \forall t \in \{2, 3, \ldots, T\}, \quad \forall i \in \{1, 2, ..., M_t\} \tag{10}$$

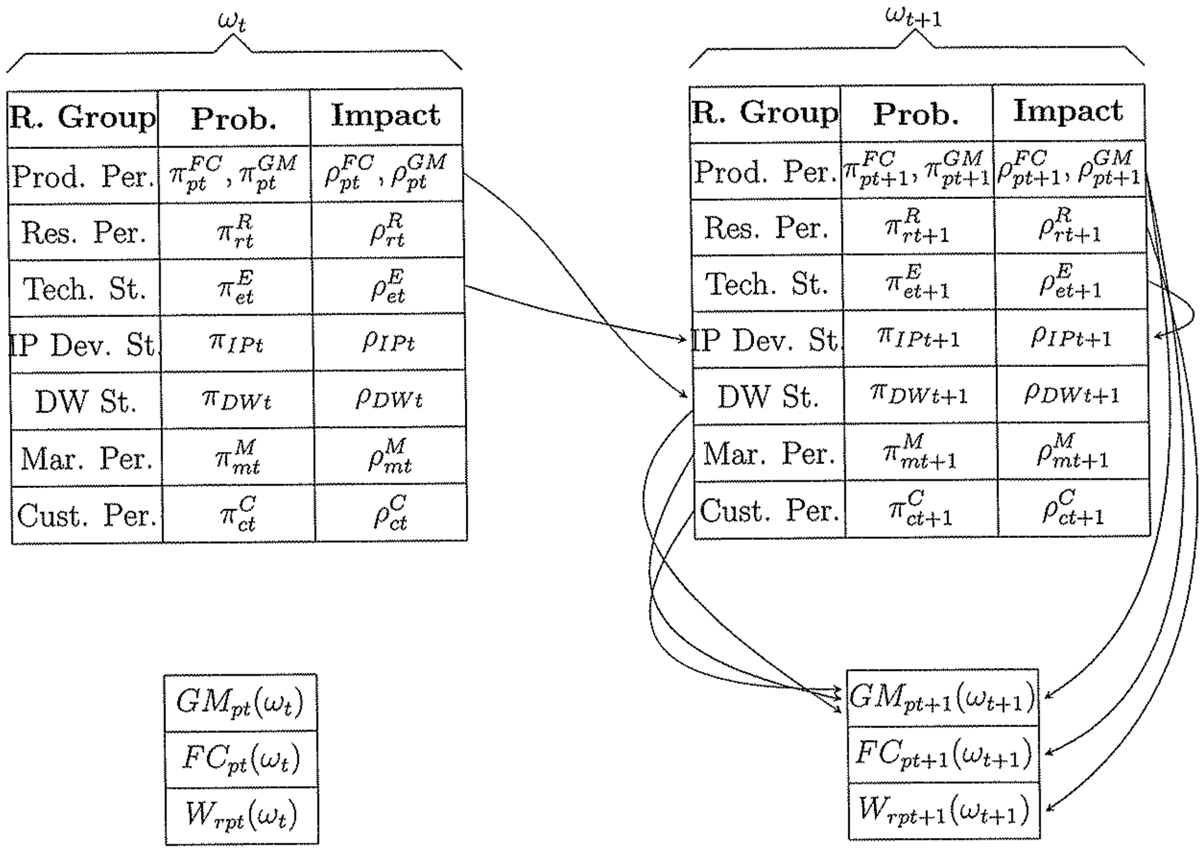

Figure 2: Outcome Relations

The impact factors for product performance ρ_{pt}^{PP}, market performance ρ_{mt}^{M}, and customer performance ρ_{ct}^{C} all depend on the random events. If the project has more than one customer or market, the effective performance is calculated using the market and customer exposures of the product. Similar functional relationships for the impact of the random outcome on impact factors, and hence on the scenario-dependent values for product fixed cost and resource usage are also used, and given in Equations (11) and (12).

$$FC_{pt}(\omega_t^i|\omega_{[1,t-1]}) = \rho_{pt}^{PP_{fc}} FC_{pt} \qquad \forall p \in P,\ \forall t \in \{2,3,\dots,T\},\ \forall i \in \{1,2,...,M_t\}, \tag{11}$$

$$W_{rpt}(\omega_t^i|\omega_{[1,t-1]}) = \rho_{rt}^{R} W_{rpt} \qquad \forall p \in P,\ \forall r \in R, \forall t \in \{2,3,\dots T\},\ \forall i \in \{1,2,...,M_t\}. \tag{12}$$

The outcome of scenario generation is a scenario tree (Figure 3). Each node of the tree carries information regarding to the system state at that stage. We will denote the set of nodes that belong to stage t as S_t. From this point forward, we will use node notation to describe gross margins, fixed

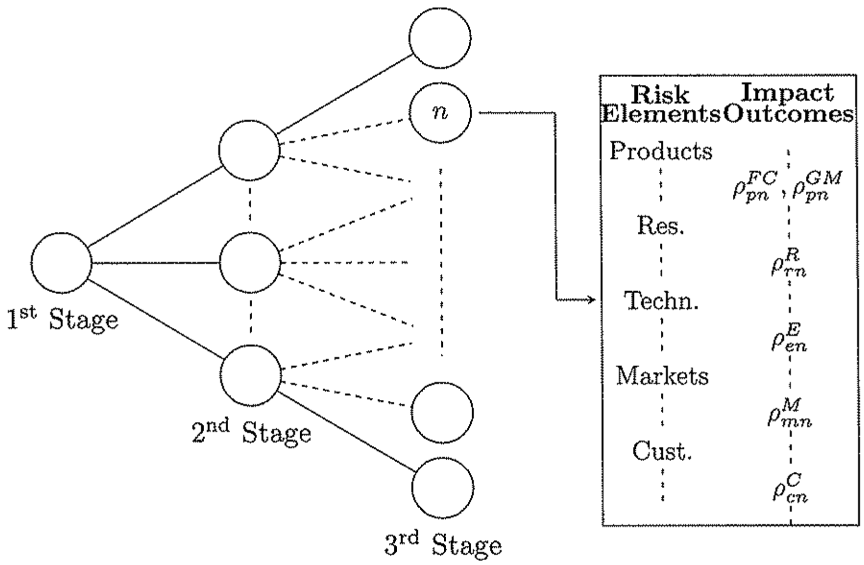

Figure 3: Scenario Tree

costs and required headcount levels. That is at stage t, GM_{pn}, FC_{pn} and W_{rpn} will be the updated gross margins, fixed costs and required headcount levels at node n ($\forall\ n \in S_t$).

5.3.1 Risk Measure

Once the scenario tree is created, Gini coefficients λ_{pt} are calculated for each product p ($p \in P$) and stage t ($t \in \mathcal{T}$). Similar to variance, the Gini coefficient is a measure of dispersion and is defined by twice the covariance of a random variable and its cumulative distribution. [16] use Gini coefficients to decide on stock portfolios, and [13] using Gini coefficients in the context of deciding on R&D projects. However, both [16] and [13] use the Gini coefficients in a different manner than proposed here. Previous work has based portfolio decisions on marginal conditional stochastic dominance—for a given portfolio of stocks/projects, the decision to include a new stock/product opportunity is made by comparing the mean-Gini adjusted returns of the stock/project pairs. In our case, we will use weighted average Gini measure to limit the product risks to some user-defined level (via Equation 9).

To compute the Gini coefficients λ_{pt}, first the return of each available product for each realization of the uncertainty in that stage (θ_{pn}) is a calculated

$$\theta_{pn} \stackrel{\text{def}}{=} GM_{pn} - FC_{pn},\ \forall p \in P_t\ \forall n \in S_t,$$

and the total return of all products for that realization (θ_n) is computed

$$\theta_n \stackrel{\text{def}}{=} \sum_{p \in P_t} \theta_{pn}\ \forall p \in P_t \forall n \in S_t,$$

where $P_t \subset P$ is the set of all products that can be selected in stage p (i.e. $b_p \leq t$). Next, the values θ_{pn} and θ_n are assembled into collections of length $|S_t|$ and each sorted in ascended order into vectors ϕ_{pt} and ϕ_t. The Gini coefficient is then computed as

$$\lambda_{pt} = 2\,\mathrm{Cov}(\phi_{pt}, F_t(\phi_t)) \quad \forall t \in \mathcal{T} \setminus \{\infty\},\ \forall_{\surd} \in \mathcal{P}_{\sqcup},$$

where $F_t(\phi_t)$ is the cumulative probability of the total return ϕ_t. Since our returns our generated via a sampling procedure, and the probability of each realization is equal, $F_t(\phi_t)$ is found by dividing the rank of each element in ϕ_t by the cardinality of set S_t. For further explanation of the meaning and calculation of the Gini coefficients, see [20].

5.4 Optimization

Once the sampling approximation of the random process has resulted in the scenario tree, we can create a tractable approximate optimization model. The solution of the model produces a "best" portfolio and required headcount levels for each resource. In the model the decision variables fall into two categories:

- Strategic decisions (x_{pn}): Binary decision variables indicating whether or not product p is included in the portfolio at node n of the scenario tree. These are the decisions associated with $x_{pt}(\omega_{[1,t]})$ in the true model presented in Section 4.

- Resource Level Decisions (y_{rn}, h_{rn}, f_{rn}): The number of resources r required at node n in the scenario tree, the increase in resources r at node n of the scenario tree, and the decrease in resources r at node n of the scenario tree, respectively. These variables correspond to the decisions $y_{rt}(\omega_{[1,t]})$, $h_{rt}(\omega_{[1,t]})$, $f_{rt}(\omega_{[1,t]})$ in the true model.

For any node n of the scenario tree, π_n is the "path probability", or probability that the sequence of events leading to node n will occur. Due to our sampling methodology, if $n \in S_t$, then $\pi_n = 1/(M_1 \cdots M_t)$. With these definitions, we can model the product portfolio management problem as follows:

$$\text{Maximize: } \sum_{t \in T} \sum_{n \in S_t} \pi_n \left[\sum_{p \in P} (GM_{pn} - FC_{pn}) x_{pn} - \sum_{r \in R} (\eta_r h_{rn} + \zeta_r f_{rn} + c_r y_{rn}) \right] \tag{13}$$

subject to

$$\sum_{p \in P} W_{rpn} x_{pn} \leq y_{rn} \qquad \forall r \in R,\ \forall t \in cT,\ \forall n \in S_t \tag{14}$$

$$y_{r1} = I_r + h_{r1} - f_{r1} \qquad \forall r \in R \tag{15}$$

$$y_{rn} - y_{r\rho(n)} = h_{rn} - f_{rn} \qquad \forall r \in R,\ \forall t \in \mathcal{T} \setminus \{1\}\ \forall n \in S_t \tag{16}$$

$$x_{pn} \leq x_{qn} \qquad \forall p \in P,\ \forall q \in Q_p,\ \forall t \in \mathcal{T},\ \forall \backslash \in \mathcal{S}_{\sqcup}, \tag{17}$$

$$x_{pn} + x_{un} \leq 1 \qquad \forall p \in P,\ \forall u \in E_p,\ \forall t \in \mathcal{T},\ \forall n \in S_t, \tag{18}$$

$$x_{pn} = 0 \qquad \forall p \in P,\ \forall t \in \{1, 2, ..., b_p - 1\},\ \forall n \in S_t, \tag{19}$$

$$x_{p\rho(n)} \geq x_{pn} \qquad \forall p \in P,\ \forall t \in \mathcal{T} \setminus \{\infty\},\ \forall \backslash \in \mathcal{S}_{\sqcup}, \tag{20}$$

$$\sum_{t \in \mathcal{T} \setminus \{\infty\}} \sum_{p \in P} \lambda_{pt} \sum_{n \in S_t} \pi_n x_{pn} \leq K \tag{21}$$

$$x_{pn} = 1 \qquad \forall p \in F_1,\ \forall t \in \mathcal{T},\ \forall \backslash \in \mathcal{S}_{\sqcup}, \tag{22}$$

$$x_{pn} = 0 \qquad \forall p \in F_0,\ \forall t \in \mathcal{T},\ \forall \backslash \in \mathcal{S}_{\sqcup}. \tag{23}$$

Equation (13) is the equivalent to Equation (1) in section 4.5 with the scenario tree and sampling, and maximizes the expected portfolio profit. Equation (14) (counterpart of Equation (2)) is the resource level requirement for the portfolio at all stages. Equation (15) (equivalent of Equation (3)) is the initial adjustment of the resource levels, and Equation (16) (equivalent of Equation (4)) is the adjustment of resources for the subsequent stages, where $\rho(n)$ represents the parent node of node n in the scenario tree. Equation (17) (Counterpart to Equation (5)) ensures that when a product p is selected for the portfolio, all the prerequisite products of it are also selected. Equation (18) (equivalent of Equation (6)) ensures that when p is selected for the portfolio, each of its mutually exclusive products, E_p, are not selected. Equation (19) (the analogue of Equation (7)) and Equation

(20) (the analogue of Equation (8)) ensure that if a product is not selected at the beginning of its life cycle it is not selected later. Equation (21) (the equivalent of Equation (9)) keeps the total average Gini measure of the selected projects below a specified risk level, K. Additionally, we include two constraints (Equations (22) and (23)) to include or exclude a specific project or set of projects for the whole planning horizon. All the products in set F_1 should be in the portfolio, and all the products in set F_0 should be omit from portfolio $\forall t \in \mathcal{T}$.

All of the steps in Figure 1 up to the solution of above model are handled via Excel VBA. However, we solve each SP model using AMPL modeling language with CPLEX vx.x solver. When a scenario tree is created, DSS also creates the required .mod, .dat and .run files for AMPL. Then, DSS activates AMPL and solves models for a specified number of runs. (See Figure 1 Steps 5 and 6). When all the runs are complete, DSS proceeds with the Refinement Phase.

5.5 Refinement

A key component of the DSS is the feedback given to decision makers about solutions obtained from the optimization model (equations (13)-(23)). Once an instance is created and solved, the sampled outcomes of uncertainty and the decision made for each outcome of the uncertainty are written to a database. Next, charts and other visual aides are constructed from this database to help the decision maker see the cause and effect relationships of model inputs. From our experience, the management found four types of charts useful in their decision making process.

- Efficient Frontier: The problem (13)-(23) is solved for increasing risk levels (K), and the optimal expected portfolio profit is plotted against K. Also, using the results from multiple runs depending on the users' choice various levels of Confidence Intervals (CI) on the expected profit are constructed.

- Initial Portfolio: Since the multistage stochastic program is often solved at different risk levels (K), there may be many different suggestions as to the optimal initial portfolio. We graphically depict the optimal portfolio with respect to the risk level.

- Products in Time: Similar to the initial portfolio, we graphically depict the portfolio at each stage of the planning horizon.

- Robustness of Products: This chart summarizes the how often project included in a portfolio were killed at later stages of development due to unfavorable scenarios.

The three-phase DSS has been tested and validated with real case studies. In the next section, we describe one of these case studies, in which Semiconductor Manufacturer's management is on the verge of constructing project portfolio for their one of major business units, in which the decisions span a four year planning horizon. Each of the different charts providing solution feedback will be demonstrated via this case study example.

6 Case Study

The company was faced with a strategic portfolio decision charting the course of one of their major divisions over a four year time period. The strategic decision involved creating a portfolio of projects to undertake from a candidate set of size 21.

For this particular business decision, inferring reasonable expected project returns from forecast values was turning out to be very difficult for management. The various "information points", such as IP development completion and the design win date, during a project's life cycle had a significant impact on the project return for this instance. Further, there was a significant amount of interdependence (pre-requisite and exclusivity) information for the projects that could be undertaken. For these reasons, this particular portfolio decision seemed like an ideal case study for the DSS in place at the company.

In the model for this portfolio selection, to more accurately represent some of the business decisions that could feasibly be taken in subsequent years, the management decided to put two more projects as "business exit" projects indicating total or partial business exits.

The timeline, cash flows (GM_{pt}, FC_{pt}), and total resource cost of the various candidate projects is depicted in Table2. The numbers in parenthesis are negative values. A few points in the figure are noteworthy. First, the business exit projects are represented by the first two lines. P1 indicates total exit from the business. If P1 is selected, then the projects in the fifth column cannot be selected. The project P2 is for a partial business exit. If P2 is selected, the subsequent projects (in the mutually exclusive project column for P2) cannot be considered for the portfolio.

Table 2: Data from Company Database

					Forecast Fixed Costs and Gross Margins									
					Year 1		Year 2		Year 3		Year 4			
No.	Proj.	IP Depend.	Preq. Proj.	M. Ex. Proj.	FC	GM	FC	GM	FC	GM	FC	GM	$\sum_t \sum_r W_{rpt}$	Tot. Value
1	P1	-	-	2,4,7,8 9,10,11	0.00	0.00	0.00	15.97	0.00	0.00	0.00	0.00	(35.30)	(19.33)
2	P2	-	-	1,4,7 8,9,10	0.00	0.00	0.00	2.45	0.00	0.00	0.00	0.00	(14.70)	(12.25)
3	P3	-	-	-	(0.30)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	(14.70)	(15.00)
4	P4	-	3	1,2,5,11	(5.00)	3.50	(3.30)	4.80	0.00	0.00	0.00	0.00	(1.30)	(1.30)
5	P5	-	2,3	4	(5.00)	3.50	(3.30)	4.80	0.00	0.00	0.00	0.00	(1.30)	(1.30)
6	P6	-	3	-	(0.50)	0.00	(0.30)	7.00	0.00	0.00	0.00	0.00	(1.10)	(5.10)
7	P7	P6 Pass	3	1,2,11	(2.80)	(0.10)	(6.30)	14.50	(6.50)	23.00	(0.90)	0.90	(17.90)	3.90
		P6 Fail			(2.80)	(0.10)	(8.80)	12.00	(6.50)	23.00	(0.90)	0.90		(1.10)
8	P8	-	-	1,2,11	(1.00)	0.00	(6.80)	0.60	(11.30)	46.60	(9.30)	44.90	(32.50)	31.20
9	P9	-	8	1,2,11	0.00	0.00	0.00	8.00	0.00	0.00	0.00	0.00	0.00	8.00
10	P10	P8 Pass	-	1,2,11	0.00	0.00	(0.60)	0.00	(5.70)	5.20	(13.10)	60.40	(23.30)	22.90
		P8 Fail			0.00	0.00	(0.60)	0.00	(12.10)	5.20	(24.60)	60.40		5.00
11	P11	-	2	1,4,7 8,9,10	0.00	0.00	(1.10)	0.00	(1.00)	0.00	(1.90)	0.00	(22.40)	(26.40)
12	P12	-	-	-	(1.10)	0.90	(0.20)	0.50	0.00	0.00	0.00	0.00	(1.90)	(1.80)
13	P13	-	12	-	(1.20)	1.00	(0.80)	1.40	(0.40)	0.20	0.00	0.00	(1.50)	(1.30)
14	P14	-	13	-	0.00	0.00	(1.20)	0.30	(0.90)	2.30	(1.10)	1.90	(1.40)	(0.10)
15	P15	-	3,12	-	(0.60)	0.00	(0.19)	3.50	(2.30)	6.30	(0.80)	2.00	(3.20)	4.71
16	P16	-	15	-	(0.30)	0.00	(1.60)	5.90	(1.80)	5.80	(0.60)	0.90	(0.30)	8.00
17	P17	-	11,15	-	0.00	0.00	(2.80)	(0.10)	(4.10)	13.60	(4.80)	24.00	(5.90)	19.90
18	P18	-	16	-	0.00	0.00	(0.10)	0.00	(1.40)	5.80	(3.80)	24.50	(3.00)	22.00)
19	P19	-	3	-	(8.00)	(0.50)	(15.40)	7.30	(18.10)	17.00	(21.60)	46.00	0.00	6.70
20	P20	-	8,19	-	0.00	0.00	0.00	0.00	0.00	0.00	(13.00)	25.00	0.00	12.00
21	P21	-	11	-	0.00	0.00	0.00	0.00	(27.90)	48.90	(31.00)	60.30	0.00	50.30
22	P22	-	11	-	0.00	0.00	0.00	0.00	0.00	0.00	(3.30)	15.20	0.00	11.90
23	P23	-	11	-	0.00	0.00	0.00	0.00	0.00	0.00	(3.62)	18.40	0.00	14.78

Second, to depict the dependency a project's cash flows on another project's IP development status, an alternative cash flow line is given under the original one. For example, the projects P7's cash flow depends on P6's IP development status. If P6 cannot complete IP development on time, then second line of P7 gives the future cash flows for the project. A similar relationship holds for projects P8 and P10. For this business decision, the IP development status does *not* affect the time to market, because IP development can be completed for P6 and P8 at higher costs. However, if this were not the case, the gross margin streams would be shifted to the right at the second line of cash flow data.

Third, the project's market release time b_t is the period when nonzero gross margins are observed and it is the same time period when the design win event should be considered for the project in the simulation described in Section 5.3. If a project passes the design win for all of its customers then the forecast gross margins GM_{pt} would be used as a base to calculate the scenario dependent GM_{pn}. However, if the product fails from all of its customers then in all other scenarios GM_{pn} would be 0. In the medium case, the GM_{pn} would be calculated using percentage product-customer relations γ_{pc} as in Equation (10).

The total headcount cost for each project is calculated from $\sum_{(t \in \mathcal{T}, \nabla \in \mathcal{R})} W_{rpt} c_r$ based on the

Table 3: Resource Related Figures (in $ millions)

	Overhead cost (c_r)	Hiring cost (η_r)	Firing cost (ζ_r)	Initial level (I_r)
R&D	$0.256	$0.200	$0.300	80
SGA	$0.150	$0.100	$0.200	20

overhead cost, (c_r), of R&D and SGA (administrative) employees. Table 3, provides cost and initial count data related to both resources.

Management considers three projects vital for the company strategy. Independent of the portfolio P3, P12 and P15 are "must do" projects. These projects are mainly technology development projects, and have very little or zero gross margins expected throughout their life cycle. However, there are multiple high return products depending on completion of these projects. (i.e. P16, P17, P18, P20). Thus, F_1 in Equation (22) is the set of these three projects.

Before undertaking a careful, quantitative analysis using the DSS in place at the company, the managers had thought that there were likely two different "business lines" that were attractive options for this portfolio decision, but that capital and manpower resource limitations would limit the ability to undertake them both. The general consensus was that the first business line (BL1) was a more conservative option, that would yield reasonable returns at low risk levels, and that the second business line (BL2) was a riskier, but potentially more profitable, undertaking. Management was particularly interested in gaining insight from the DSS in support of one of these business lines, or in learning if there were portfolio options from the set of candidate projects with more desirable risk/return characteristics that either BL1 or BL2.

For the analysis, in addition to the project specific information from the company database (Table 2), management defined the key markets (M), key customers (C), and probabilities and impact data for each risk group outcomes defined in Table 1. For this instance, there are 4 markets, 5 customers, 3 technologies, 2 type of resources and 23 projects (including the 2 "business exit" projects). With these data, the product-market (α_{pm}) and product-customer exposures (γ_{pc}) were formed to indicate the percentage relationship, between products, markets and customers. Then 'must do' projects are indicated to form $F1$ set (Equation (22)). The data collection and scenario generation process are handled with Excel's VBA module. For this instance, it is quite reasonable to consider a four-stage stochastic program, (every stage being an end of each year), as the company can decide to start or end projects roughly at the beginning of its Fiscal Year.

Even with the relatively coarse discretization of the uncertainty space described in Table 1, there are still on the order of a quadrillion (10^{15}) outcomes for each stage. From these outcomes, the optimization model was built to contain 10,000 scenarios: (with $M_2 = 50, M_3 = 20, M_4 = 10$ and $N_2 = 50, N_3 = 1000, N_4 = 10,000$). Each of these 10,000 scenario sampled instances with $Z = 25$ replications. And for each sampled instance, a risk frontier was created by solving with 12 different risk levels (starting from 10 with increments of 5).

Each of the optimization instances was solved by CPLEX (v9.0) directly through calls from the AMPL modeling language. Solving these instances took roughly 6 1/2 hours, and there was not significant variation in solution time from instance to instance. The value $N = 10000$ scenarios was chosen for a sample size in part due to the fact that this was the largest instance that solved in a reasonable amount of CPU time. While 10,000 scenarios is a very, very small fraction of the total number of scenarios, we have seen even for this small fraction, both the optimal solution and its value obtained from different sampled instances tends to be very "stable", indicating that it is likely to be the true optimal solution. Recent theoretical evidence showing that there is a good probability of obtaining the true optimal solution from a small sample size is given by [17] and [9].

Figure 4 and Figure 5 shows the optimization phase results and graphs described in Refinement Phase. In the large majority of the 25 samples instances, at lower risk levels (10-25) and high risk levels (50-65), the optimal initial portfolio is shown on the initial portfolio chart of Figure 4. The optimal initial portfolio projects for these risk levels are very similar to projects in BL1. At risk levels (30-45), the optimal initial portfolio is shown on the upper right chart of Figure 5, and the optimal initial portfolio at these risk levels are very similar to BL2. The solution of the optimization instances did not indicate that there were significantly better portfolios to consider than those already under consideration at the company (BL1 and BL2). As seen from the efficient frontier charts of Figure 4 and Figure 5, as the risk level increases more projects are added to the portfolio. The whole portfolio for BL1 and BL2 can be seen from the project robustness charts for each risk level. The lower left charts provide the summary statistics for both BL1 and BL2 portfolios, and show which projects are continuing for each planning year.

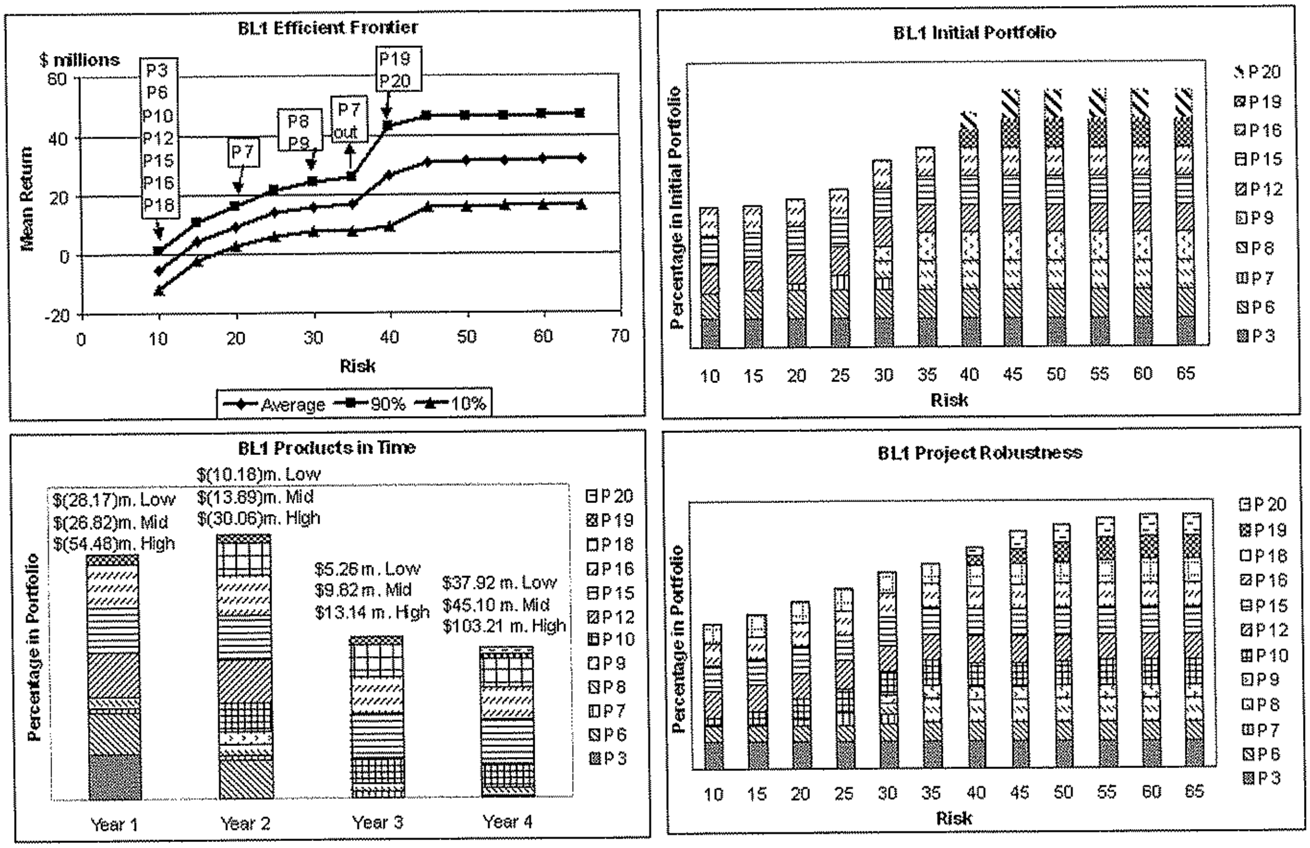

Figure 4: Business Line 1

Figure 6 displays the average efficient frontier line for BL1 and Bl2 together. The figure is the average of 25 runs and statistics of this figure is given in Table 4. BL2 portfolio is not feasible for the risk levels (10-20), then on the average for the risk levels 30 and 35, it becomes optimal. For all the other risk levels BL1 portfolio is optimal.

Table 5-7 display the expected yearly required headcount levels from each type of resources when BL1 or BL2 portfolio selected, for low medium and high risk levels, respectively. One important

Table 4: Comparison in Risk

	Low Risk (10-25)		Medium Risk (30-45)		High Risk (50-65)	
	BL1	BL2	BL1	BL2	BL1	BL2
Return	5.38	2.63	22.43	24.72	31.32	29.45

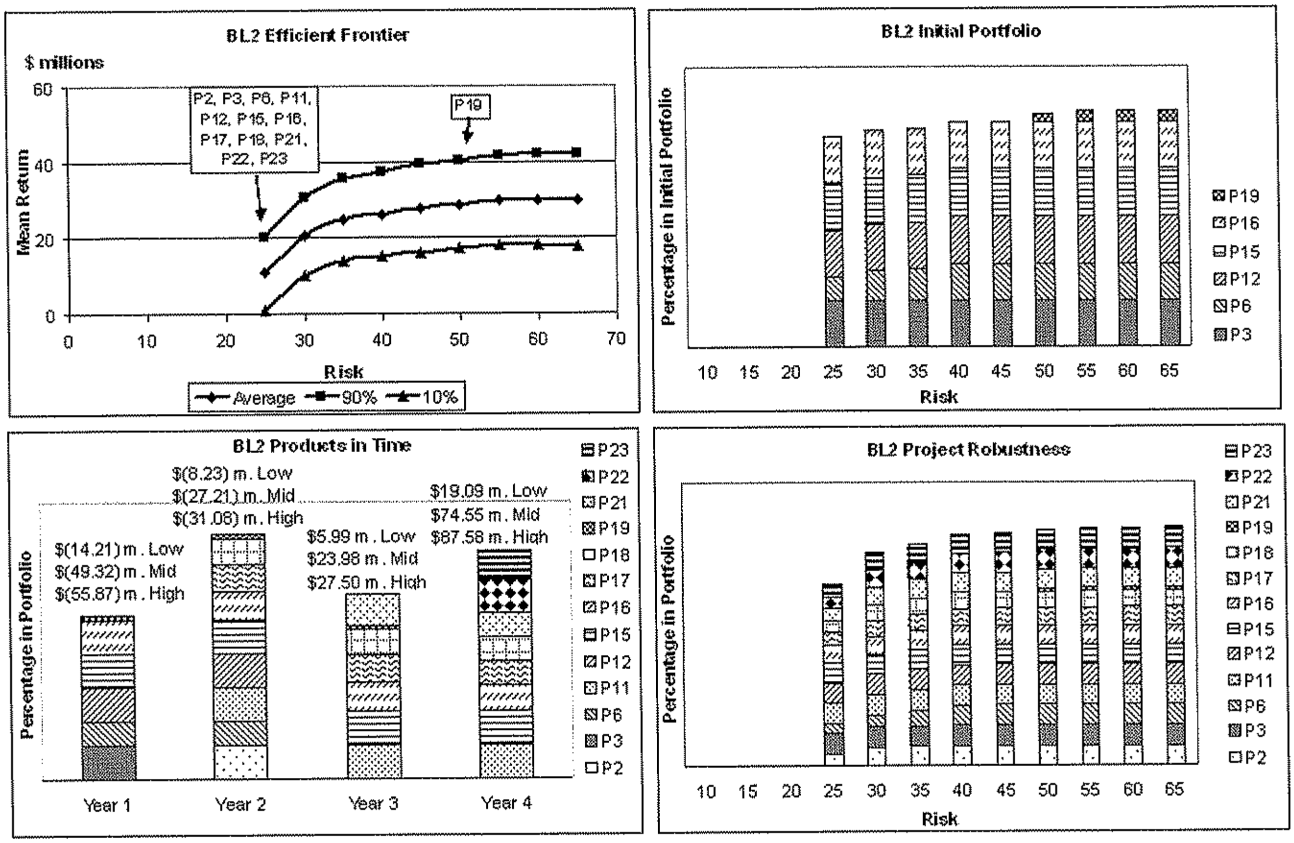

Figure 5: Business Line 2

observation is that the required level of resources for BL2 dramatically reduces after the first year, yielding a high turnover rate. Relatively required level of resources for the BL1 is more stable. The last line depicts the yearly profit of the portfolios, which are shown in the lower left charts of Figure 4 and Figure 5. According to this both BL1 and BL2 do not earn profits for the first two years of the planning horizon.

One of the key results of the solution is that even at the highest risk levels both BL1 and BL2 reveal lower expected returns than the forecasted figures in Table 2. This is because forecasted figures do not adequately take into account the possibility of zero gross margins when the project fails. Another driver of the difference in forecasted figures and the expected optimal figures is the probability and impact data specified by the management. It seems that for this particular instance, the management consensus was to be slightly "pessimistic" about a few key impact events.

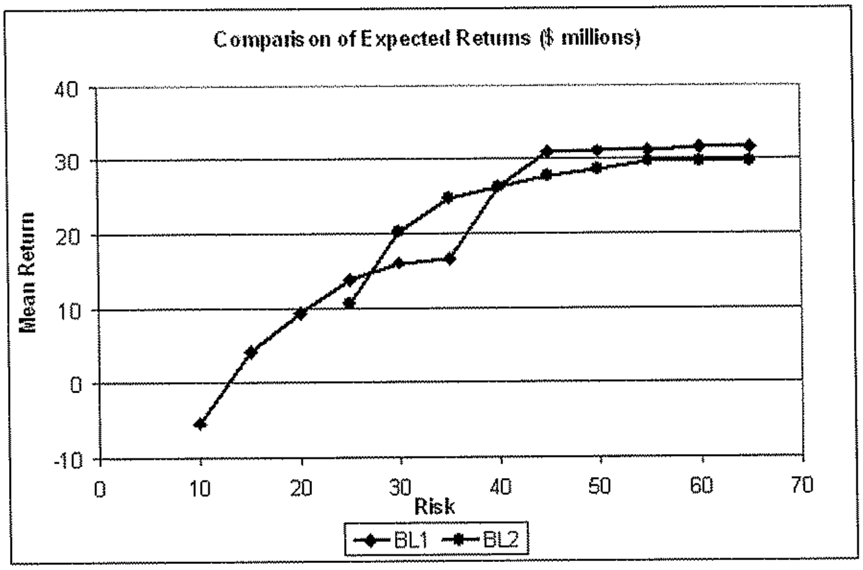

Figure 6: Optimization Phase Results

Table 5: Headcount and Profit Comparison in Time: Low Risk levels (10-25)

	Year 1		Year 2		Year 3		Year 4		Total	
	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2
R&D	80	142	35	54	27	36	27	36	Mean:**42**	Mean:**67**
SGA	18	31	9	14	7	9	7	9	Mean:**10**	Mean:**16**
Return	(28.17)	(14.21)	(10.18)	(8.23)	5.26	5.99	37.92	19.09	**4.83**	**2.63**

The choice of the portfolio depends on the degree of risk averseness. The optimization problem suggests BL1 at all the risk levels except 30 and 35, and BL2 for the risk levels 30 and 35. With the aid of efficient frontier chart, one can conclude project specific risk information from the risk levels at which each product first appears in the portfolio. For example P6, P10, P16 and P18 are low risk and P7, P8,P9 are moderate risk and P19, P20 are high risk projects for BL1, and almost every BL2 projects are moderate risk projects and P19 is the high risk project. After viewing the results, management did not need to make further analysis with input change and decided to invest in BL1 projects. Portfolio BL1 was more attractive from the following points: 1) On the average it yields higher profit, 3) Portfolio and expected profit is more responsive to the degree of risk averseness, which, by the management, is considered more realistic compared to real life situations 3) Required level of resources for the portfolio is less fluctuating,

Table 6: Headcount and Profit Comparison in Time: Medium Risk levels (30-45)

	Year 1		Year 2		Year 3		Year 4		Total	
	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2
R&D	111	142	69	55	58	37	58	37	Mean:**74**	Mean:**67**
SGA	24	31	18	14	15	9	15	9	Mean:**18**	Mean:**16**
Return	(26.82)	(49.32)	(13.89)	(27.21)	9.82	23.98	45.10	74.55	**14.21**	**22.00**

Table 7: Headcount and Profit Comparison in Time: High Risk levels (50-65)

	Year 1		Year 2		Year 3		Year 4		Total	
	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2	BL1	BL2
R&D	114	137	72	52	61	36	61	36	Mean:**77**	Mean:**65**
SGA	25	30	19	13	16	9	16	9	Mean:**19**	Mean:**15**
Return	(54.48)	(55.87)	(30.06)	(31.08)	13.14	27.5	103.21	87.58	**31.81**	**28.12**

A final analysis was done to attempt to quantify the value of using a stochastic programming approach against a deterministic (or mean-value) model by computing the *Value of the Stochastic Solution* (VSS). Expected or mean value problem is simply solved by fixing all the random quantities to their mean values. Then the difference between the expected result of using stochastic solution and mean value solution is the VSS. [3] In our instance, we found a portfolio of products by fixing the impacts of outcomes to their expected values. Then, for each risk level, we computed the VSS by comparing the performance of optimal portfolio and mean value portfolio under a sample of scenario set. We measured VSS under 15 different scenario sets and took the average values, which are shown in Table 8. For nearly all levels of risk, the VSS was significantly greater than 0, indicating that by explicitly considering the cost of reacting to uncertainty in the model, better decisions are being made.

Table 8: Average Value of Stochastic Solution (VSS) under different risk levels

	Risk Level											
	10	15	20	25	30	35	40	45	50	55	60	65
VSS ($ millions)	10.11	4.50	0.49	1.85	7.41	12.37	12.81	12.87	11.46	10.29	9.55	9.35

7 Implementation Experience

The project portfolio selection DSS described here has been constructed, revised, and remodeled over a two-year period via continuous interaction with the firm's senior management. Through our interactions with the decision makers expected to use the tools, we observed that they do not feel comfortable tinkering with the DSS during the optimization phase, but it is intuitive for them to comprehend the system's recommendations through the tables and charts. Thus, we designed the DSS such that after the initial data input, the process flows automatically, and the requisite charts and the tables are displayed at the end. The management discusses the results using these charts and tables. If more "what if" analysis or in-depth statistics are needed, the DSS allow them to delve directly into the solution through a spreadsheet interface. The system also provide utilities for them to make additional analysis by changing part of the input and rerun the entire process.

During our implementation, the management decided to connect the DSS with the company database so that information that is already available can flow in directly. A main concern of the management is to make sure that the DSS generates the type of charts and tables that the whole decision team can easily interpret and understand. They were heavily involved in the design of the output format such that charts and tables generated were very similar, if not identical, to what they have already being using in the process. Overall, the senior management has been very satisfied with the way the DSS tool helped in their decision process. Since it is primarily data driven while taking into consideration some human judgements, the tool creates a level of formality and credibility to the process. Currently, the firm uses the DSS for the portfolio selection and management process, and they are planning to use the tool for higher-level strategic decisions such as the alignment of market potentials and business units within the company.

8 Conclusions

As is typical in the high-tech environments, the semiconductor industry faces a dynamic and volatile market. The development of intellectual property (IP) via R&D projects is among the most important decisions high-tech companies make, and the selection of a well balanced R&D portfolio require significant effort and data analysis. In this paper we introduce a three-phase decision support structure for the project portfolio selection process at a major U.S. semiconductor company. The key features of the decision support system we developed for the company are as follows:

- Flexible risk modeling via scenarios; this allows us to incorporate not only quantitative information in the company database, but qualitative information distributed among the decision makers: their assessment about different aspects of the business environment, likelihood of successful technology development, exposure to customer and market segment success.
- A multi-stage stochastic program that provides an effective way to synthesize large volume of information such as business constraints and project interdependence, while systematically evaluating the various sources of uncertainties present in the business environment.
- An effective interface with decision makers that provide access to all components of the DSS including detailed information gathering from company databases, and survey for key decision makers, wizard-like user interface to assist in model building, error checking routines that notify the users in case of missing or illogical data, interaction with sampling and optimization tools (for advanced users), and automatically generated charts, tables, and figures.
- Sensitivity analysis tools that allow decision makers to resolve particular instances with different parameters and evaluate the robustness of the outcome.

We demonstrated, through numerical examples, a real-world R&D project portfolio selection problem using our DSS. We show the typical output generated by the tools and different ways that the decision maker may use the tool to analyze the trade-off in different portfolio selection alternatives, balancing risk with expected return. The optimal portfolio constructed by the model has been implemented at a particular business unit at the company.

References

[1] P. Artzner, F. Delbaen, J. M. Eber, and D. Heath, *Coherent Measures of Risk*, Mathematical Finance (1999), no. 9, 203–228.

[2] R. Bekkers, G. Duysters, and B. Verspagen, *Intellectual Property Rights, Strategic Technology Agreements and Market Share Structure: The Case of GSM*, Reseach Policy **31** (2002), 1141–1161.

[3] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, Springer, New York, 1997.

[4] M. H. Boisot, *Knowledge Assets, Securing Competitive Advantage in the Information Economy*, Oxford University Press, Oxford, 1998.

[5] A. Chen, K. Subprasom, and J. Zhaowang, *Mean-Variance Model for the Build-Operate-Transfer Scheme Under Demand Uncertainty*, Transportation Research Record (2003), no. 1857, 93–101.

[6] R. G. Cooper, S. J. Edgett, and E. J. Kleinschmidt, *Portfolio Management for New Products*, Addison-Wesley, Reading, MA, 1998.

[7] M. W. Dickinson, A. C. Thornton, and S. Graves, *Technology Portfolio Management: Optimizing Interdependent Projects Over Multiple Time Periods*, IEEE Transactions on Engineering Management **48** (2001), 518–527.

[8] M. A. Greiner, J. W. Fowler, D. L. Shunk, W. M. Carlyle, and R. T. McNutt, *A Hybrid Approach Using the Analytic Hierarchy Process and Integer Programming to Screen Weapon Systems Projects*, IEEE Transactions on Engineering Management **50** (2003), 192–203.

[9] A. Kleywegt, A. Shapiro, and T. Homem-de-Mello, *The sample average approximation method for stochastic discrete optimization*, SIAM Journal on Optimization **2** (2001), 479–502.

[10] J. T. Linderoth, A. Shapiro, and S. J. Wright, *The empirical behavior of sampling methods for stochastic programming*, Annals of Operations Research **142** (2006), 219–245.

[11] H. Markowitz, *Portfolio Selection*, Journal of Finance **7** (1952), 77–91.

[12] M. M. Orman and T. E. Duggan, *Applying Modern Portfolio Theory to Upstream Investment Decision Making*, Journal of Petroleum Technology (1999), 50–53.

[13] J. L. Ringuest, S. B. Graves, and R. H. Case, *Conditional Stochastic Dominance in R&D Portfolio Selection*, IEEE Transactions on Engineering Management **47** (2000), no. 4, 478–484.

[14] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro, *A Stochastic Programming Approach for Supply Chain Network Design under Uncertainty*, European Journal of Operational Research **167** (2005), 96–115.

[15] H. Shalit and S. Yitzhaki, *Mean-gini, portfolio theory, and the pricing of risky assets*, Journal of Finance **39** (1984), no. 5, 1449–1468.

[16] ———, *Marginal Conditional Stochastic Dominance*, Management Science **40** (1994), no. 5, 670–684.

[17] A. Shapiro and T. Homem-de-Mello, *On the rate of convergence of optimal solutions of Monte Carlo approximations of stochastic programs*, SIAM Journal on Optimization **11** (2000), 70–86.

[18] C. Stummer and K. Heidenberger, *Interactive R&D Portfolio Analysis with Project Interdependencies and Time Profiles of Multiple Objectives*, IEEE Transactions on Engineering Management **50** (2003), 478–84.

[19] B. Verweij, S. Ahmed, A.J. Kleywegt, G. Nemhauser, and A. Shapiro, *The Sample Average Approximation Method Applied to Stochastic Routing Problems: A Computational Study*, Computational Optimization and Applications **24** (2003), 289–333.

[20] S. Yitzhaki, *More Than a Dozen Alternative Ways of Spelling Gini*, Reseach on Economic Inequality **8** (1998), 13–30.