A Modern Simulation Approach for Pharmaceutical Portfolio Management

Kirk Solo and Mark Paich
SimNexus LLC

Keywords: portfolio, real options, optimization, hybrid simulation, pharmaceutical

Abstract

By creating an integrated simulation environment that models the underlying structure of a pharmaceutical enterprise portfolio it becomes possible to identify the optimal longitudinal allocation of finite resources across the constellation of available investment opportunities. The implementation of a hybrid approach that integrates multiple modeling techniques and analytic disciplines allows for a comprehensive environment that captures the underlying dynamics that drive observed market behavior. The implementation of an object oriented model structure constrains the model's complexity by supporting dynamic re-use of both structure and logic. By wrapping such a simulation with advanced optimization approaches it becomes possible to evaluate the pharmaceutical enterprise in a holistic fashion that avoids local optima that can be detrimental to the enterprise. The end result is an analytic approach that not only identifies the optimal enterprise investment portfolio, but provides detail around the structure and dynamics that create the observed behavior.

A DYNAMIC MODELING APPROACH TO PHARMACEUTICAL ENTERPRISE PORTFOLIO EVALUATION

A fundamental challenge in managing a pharmaceutical or biotechnology enterprise is identifying the optimal longitudinal allocation of finite resources across the infinite constellation of available investment opportunities. Recent advances in management theory, software development, and computer hardware have opened the door to a new class of longitudinal simulation that supports an enterprise model capable of identifying optimal portfolio management strategies.

Because the dynamics being modeled exhibit behaviors that call for different simulation approaches, this paper proposes the utilization of a hybrid simulation environment that can simultaneously and dynamically incorporate agent based, discrete event, and continuous equation approaches. The environment that we have used for this simulation is called AnyLogic™, and it also provides integrated Monte Carlo analysis and advanced optimization algorithms.

The Advantages of Dynamic Modeling for Enterprise Portfolio Evaluation

The simulation approach outlined in this paper presents many advantages in comparison with traditional enterprise portfolio evaluation approaches. Many efforts into product portfolio analysis have been deeply rooted in financial theory. For those familiar with traditional options and portfolio analysis certain questions should immediately emerge when confronted with the idea of applying these approaches to a pharmaceutical enterprise portfolio structure. In equity markets there is a vast history of data that can be leveraged to determine expected levels of return, risk, and interdependence. This data plus the terms of the option structure such as stock price, option price, strike price, and time to expiration are used to determine the components of the optimal portfolio. How then can we apply these principles in a context where we have no such data?

The answer to this lies in the application of dynamic simulation which models the business processes that underlie the constellation of potential investments that will make up the enterprise portfolio. Dynamic simulations can be used to create various scenarios that take the place of historical data in the options and portfolio models. Real Options models can be used to capture the New Product Development (NPD) structure. Market Dynamics models can be implemented to capture new product adoption, product sales, and any interdependencies between compounds and/or indications. This ability to capture interdependencies provides a measure of covariance that is intrinsically used by the simulation to identify optimal portfolios. Probability trees or other similar algorithms mirror real processes and can provide a measure of risk that is also included in the overall portfolio evaluation. Process delays can be integrated to capture the time elements needed to adequately assess the value of an option. Market simulations can then be used to asses additional measures of risk, interdependence, and most importantly return.

The utilization of a dynamic simulation results in measures of covariance, variance, and expected return that evolve from the actual structure of the business environment. This provides the information necessary to evaluate and identify optimal project portfolio structures, and also helps to identify the root drivers of success that are not readily apparent in traditional portfolio analysis. This actually provides potential for altering the risk structure of the underlying investment opportunities, and improving the potential risk/reward balance of the entire portfolio.

Defining the Challenge

In essence this approach is based on the compilation of several management science disciplines that have been distilled into software objects that can be parameterized using real-world data. These object-oriented structures can then interact in a virtual environment that is capable of supporting Monte Carlo analysis and advanced forms of optimization under uncertainty.

In order to provide a clear understanding of how the environment works a clear definition of the business objectives must first be established. For the purposes of this paper, the constellation of potential investment opportunities will be limited to product investments and will include products in the NPD pipeline, products already available in the market, and products that are available either through mergers and acquisitions or in-licensing.

The simulation must ultimately answer the question: "What longitudinal resource allocation across the available product investment opportunities yields the best results for the enterprise?"

To answer this question in a holistic fashion we must model each of these components in a manner that provides simulated behavior that is consistent with observed behavior in the market. Once this is achieved we can then implement decision points and resource constraints that will form the basis for our evaluative process.

Once the simulation environment is accurate and active, we can perform Monte Carlo simulations on specific strategies which can in turn be wrapped in an advanced optimization algorithm to ultimately identify optimal enterprise longitudinal investment strategies.

Simulation Approach

To reasonably simulate such an expansive problem set it is necessary to find a way to simplify the underlying structure. To this end we have utilized an object-oriented modeling approach that allows similar entities to be based on a standard structure. Rather than having to re-create the structure for each and every entity, it becomes possible to dynamically generate objects that can be parameterized based on their unique attributes.

This paper details a simulation approach that incorporates a wide range of management science disciplines that have been integrated into an object-oriented simulation environment that also supports advanced forms of optimization. By evaluating the enterprise in a holistic fashion it becomes possible to identify those strategies that will provide the greatest level of performance for the enterprise rather than identifying local optima that apply to only one piece of the puzzle.

DEFINING THE BUILDING BLOCKS

In order to accurately model the enterprise portfolio it is necessary to distill business processes and activities into a set of objects that when incorporated into a simulation provides a fair representation of reality. To this end we must start by defining the structure and parameters that will define each of the individual components that will be included in the simulation.

Two generalized product simulation structures will be used to define our model. The first structure will capture the behavior of compounds that are or will be in the NPD pipeline structure. This structure can capture licensing dynamics as well as M&A dynamics as they relate to compounds in or entering the pipeline. The second structure will capture the market dynamics of compounds that are or will be available in the market. Again, licensing and M&A dynamics can be captured in the same structure simply by adding the new compounds to the simulation, and incorporating the financial dynamics of the deal.

Both of these structures will capture the longitudinal consumption of resources as well as the longitudinal generation of revenues.

NPD Pipeline Structure

Considerable academic effort has been put into providing an analytic framework for pharmaceutical and biotechnology NPD pipelines (Loch and Kavadias 2002). Recent evaluations of NPD pipelines have focused on an adaptation of a binomial options approach that incorporates a probabilistic option tree (Ding and Eliashberg 2002). This structure is used to derive algebraic equations that are combined with probabilistic success factors for each NPD phase. This is coupled with an evaluation of expected operating profits from successful compounds to yield an estimation of the optimal number of initial compounds starts for an individual indication.

Probability Tree

While a formulaic approach is simple to implement, it is more flexible and of more use in a simulation environment when it is represented as a probability tree. This structure will be used as the founding of our NPD pipeline analysis.

For our purposes compounds in the NPD pipeline will potentially flow through five distinct phases. The first phase will be defined as the Discovery phase. This will be followed by three clinical evaluation phases called Phase I, Phase II, and Phase III. The probability tree will conclude with the product Launch.

Compounds will flow through the five stages based initially on a probabilistic structure that is calibrated to historical data gathered from similar compounds and/or approaches. The compound will be modeled as an object, and will have unique parameters that will be contained in a common structure.

There will be two potential probabilistic flows out of the Discovery phase that will correspond to success or failure. The sum of the two probabilities will equal 1. In the event of success the compound will proceed to Phase I, and in the event of failure it will be discarded.

The three subsequent trial phases will have an additional possibility that will incorporate the probability of human error relating to dosing or other application factors and will capture titration dynamics by passing the compound back through the current phase. The success and failure transitions will behave the same as they do in the Discovery phase. These three probabilities will sum to 1.

The following graphic represents a probability tree that can be parameterized for an individual compound.

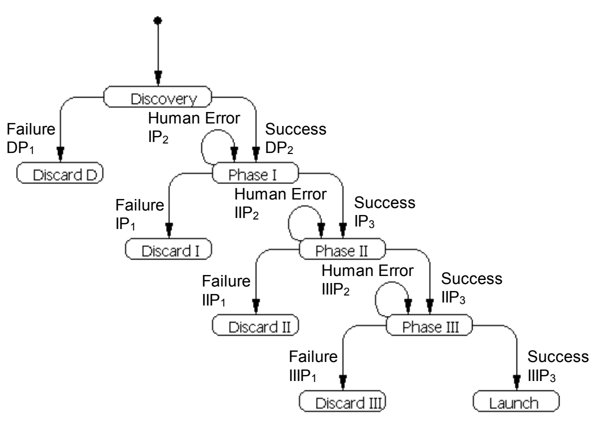

Figure 1. NPD Pipeline Probability Tree

Delay Structure

The above structure can also incorporate delay elements within each of the phases. For this implementation these delays are parameterized by an expected delay value coupled with a probabilistic distribution structure. This will provide support for a variety of analytic techniques that will be detailed later in the document. It also ensures that we will capture the temporal dynamics in as accurate of a fashion as possible.

Decision Tree

In the form above the structure can be used to create a simulation that can determine the optimal number of compounds that should be started through the NPD pipeline for a given indication. While this is beneficial, it falls far short of a comprehensive portfolio model. In order to meet our objectives we must add another layer of functionality.

We are going to add to the structure a series of decision opportunities relating to intentional discarding of compounds that have successfully passed a phase, as well as the ability to buy or sell rights to compounds that are in one of the NPD pipeline phases. Also implicit in this structure is the ability to dynamically add new compounds to the simulation. This is important when analyzing M&A efforts as they relate to the enterprise portfolio.

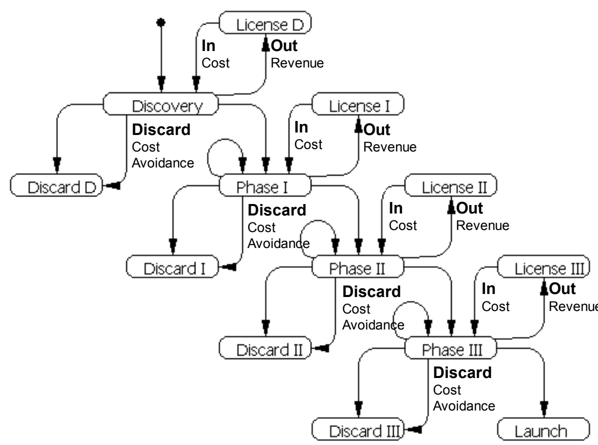

Figure 2. NPD Pipeline Decision Tree

Object Oriented Structure

The hybrid probability-decision tree above is a visual representation of an individual flow. These diagrams can be parameterized to define the delays and if-then-else algorithms that will guide the system's behavior. This logic can be embedded in the objects that represent NPD pipeline compounds. These objects can be dynamically created and destroyed. It is also possible for each individual object to incorporate different data parameters as well as different decision algorithms that will ultimately determine their behavior in the simulation.

Resource Consumption

While a compound is in one of the NPD pipeline stages, it consumes resources at a rate consistent with the parameterization that has been applied for that specific compound. This allows the simulation to capture the longitudinal resource consumption of each entity that is active in the simulation in a manner that is consistent with real-world behavior. By establishing the quantity of available resources it becomes possible to use resource consumption as a driver of systemic constraints.

Active Compounds

The simulation approach for modeling compounds available in the market comes from the world of dynamic simulation (Senge) (Warren) and will include concepts and structures significantly different from the NPD pipeline evaluation structure. However, the approach still captures key temporal dynamics as they relate to longitudinal resource allocation.

The approach that will be adopted for compounds that are active in the market or just being introduced to the market comes from the book Dynamic Pharmaceutical Marketing Models: Principles and Applications, by Mark Paich, Corey Peck, and Jason Valant which will be available through CRC Press in 2004. This approach is characterized by continuous equation analysis of patient flow dynamics that is coupled with physician adoption models. This approach allows for a time series simulation that captures the diffusion rate of compounds for a specific indication.

Patient Flow Models

The first component of a Dynamic Pharmaceutical Marketing Model is the Patient Flow model. This continuous equation simulation captures the movement of patients through a series of states.

Critical aspects of the model include the concepts of incidence rate, diagnosis rate, compliance, persistence, and treatment re-initiation.

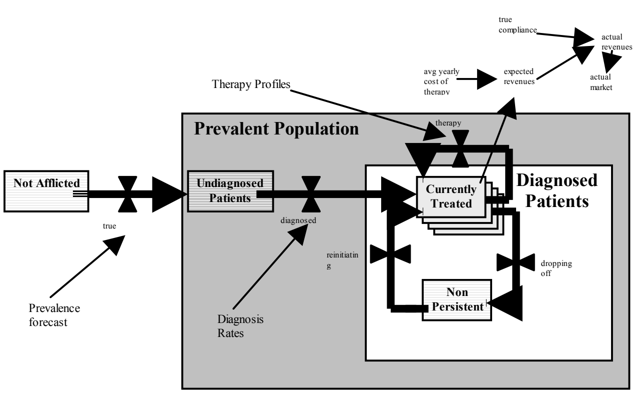

Figure 3. Patient Flow Model

By capturing the flow of patients, it becomes possible to establish the volume of potential market opportunities over time. However, this is not adequate for an accurate representation of market share because patients do not determine what compound they use for a given indication. This decision is made by a physician, so for this reason a physician adoption model is also needed.

Physician Adoption Models

It is important to understand the process by which physicians become aware of a compound, and the process by which they come to prescribe a compound, and the dynamics that effect how often they prescribe the compound for a particular indication.

To capture the dynamic of compound awareness another stock/flow structure is used to model the awareness diffusion.

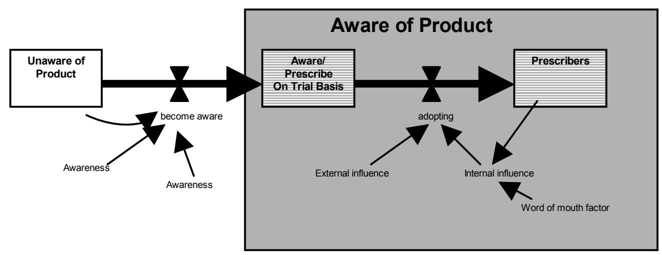

Figure 4. Physician Awareness Model

This structure provides an excellent framework for evaluating the impact of awareness marketing on the speed and magnitude of awareness diffusion.

While this provides a story of how many physicians are prescribing a compound, it falls short of establishing what kind of market share is being achieved. To this end we must combine the physician awareness model with the patient flow model, and incorporate a third factor which is treatment attractiveness.

Treatment Attractiveness

By establishing the relative longitudinal utility of the available compounds for a specific indication and combining this information with the relative awareness of available compounds, it becomes possible to establish the market share that a compound will capture through new prescription decisions. These utilities are defined in a longitudinal fashion to capture marketing impact dynamics as well as the impact of competing compounds newly introduced to the market. For patients that are switching or re-initiating treatment we can also capture historical information around switching behavior to provide a comprehensive estimation of total treatment attractiveness.

Evaluative Metrics

Patient flow, physician adoption, and treatment attractiveness can be combined into an integrated simulation for a specific indication that will track the sales of compounds over time. The resulting longitudinal revenue structure can prove critical when establishing the overall objective function.

By also tracking investments in marketing that will impact how compounds perform in the market, it is possible to create another strategic lever than can be evaluated through the simulation process.

Parameterized Bass Diffusion

It is probably obvious at this point that creating market models for each indication addressed by the enterprise would be a very daunting task. Fortunately we can incorporate a shortcut that involves using a parameterized Bass Diffusion equation that can be based on previously evaluated indications that share similarities with the indication being considered.

The nature of the indication as well as the marketing approach taken will have direct impact on the shape and magnitude of the diffusion curve. By parameterizing the Bass Diffusion curve we can provide a reasonable estimate of behavior without the full analysis.

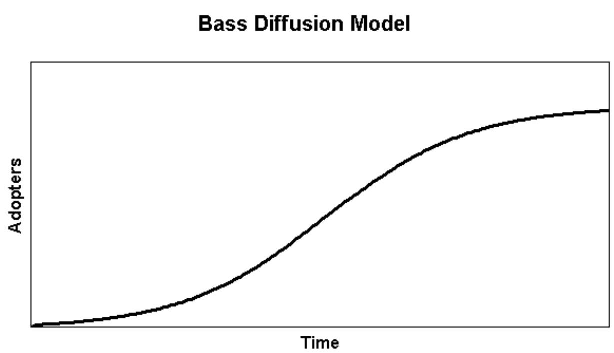

Figure 5. Bass Diffusion Curve

Modeling Enterprise Resources

The structure of the simulation allows for the consumption and creation of enterprise resources to be tracked over time. These resources can include cash, personnel, or Property, Plant, and Equipment.

For the purposes of this simulation enterprise resources are going to serve as a constraint on the system. Depending on the type of resource they will either be consumed by the underlying simulation activity or temporarily engaged in an activity. An example of a consumed resource would be cash. Resources such as human capital and PP&E will be temporarily engaged by an activity, and then made available again once the activity is completed.

The enterprise resource structure can be arbitrarily complex, but should be as simple as possible while still capturing the behavior of the system.

MODEL INTEGRATION AND PARAMETERIZATION

Up to this point several disparate approaches have been defined that can be simulated in their own right, and would prove very useful. However, an enterprise portfolio model hasn't yet been integrated.

The next step in the simulation creation process is to integrate and parameterize the components to form a baseline model that can be simulated. This process will involve the instantiation and parameterization of simulation objects and may include a longitudinal plan for M&A, licensing, or new product efforts that will result in the dynamic instantiation of model objects.

To simplify this process, it is often best to first replicate the current state of the enterprise being evaluated.

Instantiation and Parameterization

Each object that has been created can be though of as a shell that is waiting to be populated with data and decision rules. In this step objects must be created for every compound currently in the NPD pipeline or available in the market, complete with their specific parameters as well as the baseline decision rules that will determine the initial strategy. Additionally, products added to the pipeline or market through discovery, licensing or M&A must be "scheduled" for dynamic creation at the appropriate point in time.

On the surface this seems to be a daunting task, but fortunately this can be programmatically driven and only requires that critical values for each compound be entered in a spreadsheet.

Baseline Run

Once the model has been set up it can be simulated for the baseline run. This can involve a single simulation to establish some descriptive understanding of how the structure of the model impacts its behavior.

Critical pieces of information such as aggregate longitudinal resource consumption and revenue creation can be captured through this process.

LONGITUDINAL DYNAMICS

There are a variety of longitudinal dynamics that will be observed in a baseline run. One of the most important pieces of information is an understanding of how each individual object in the model interacts with other objects, and what the resulting aggregate values look like for variables being tracked.

There are three primary dynamics that can be observed in some combination over time. The first is a Smooth Dynamic that is characterized by a lack of peaks or troughs. The second is a Gap Dynamic that is characterized by a drop and subsequent recovery in level. The third is a Peak Dynamic that is characterized by an increase and subsequent decline. One or many of these may combine to form the overall longitudinal dynamics. The following figures show how data from multiple compounds can be aggregated to produce these dynamics.

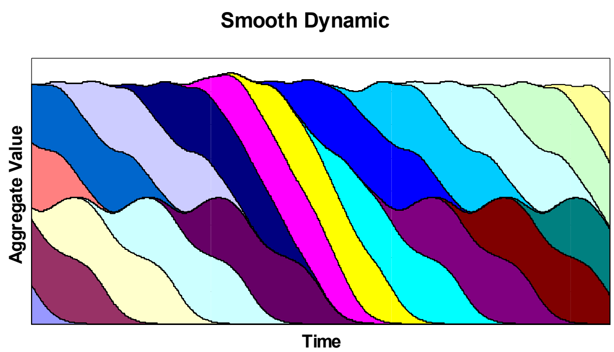

Figure 6. Smooth Aggregate Dynamic

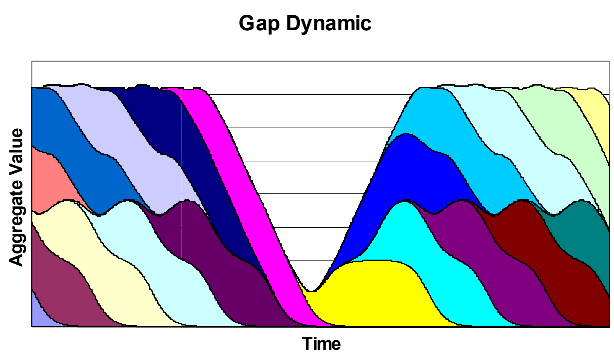

Figure 7. Gap Aggregate Dynamic

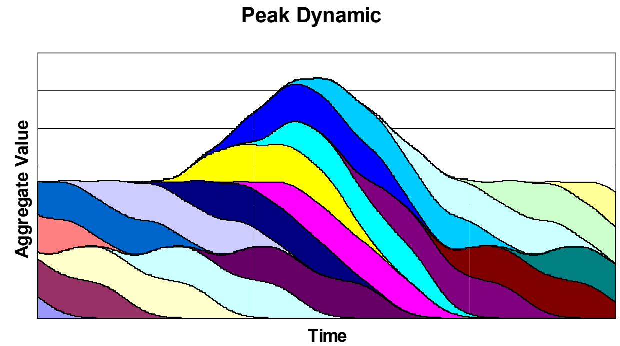

Figure 8. Peak Aggregate Dynamic

The complexity of these dynamics is evident by the number of objects that contribute to the aggregate value. It is critical that the longitudinal allocation of assets is smoothed in order to reduce volatility and achieve optimal enterprise behavior.

If it is assumed that investments will result in a delay followed by a corresponding and roughly commensurate revenue stream, then we can observe how each of these dynamics can affect aggregate earnings over time.

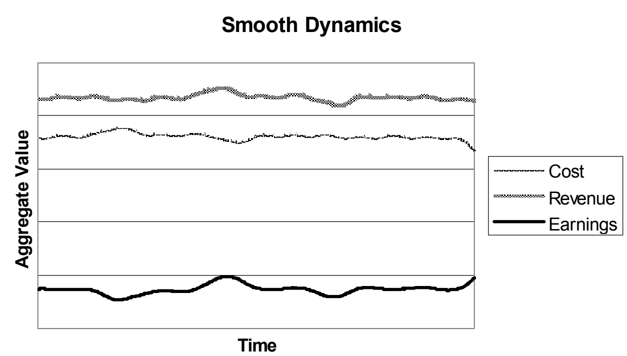

Figure 9. Smooth Longitudinal Dynamics

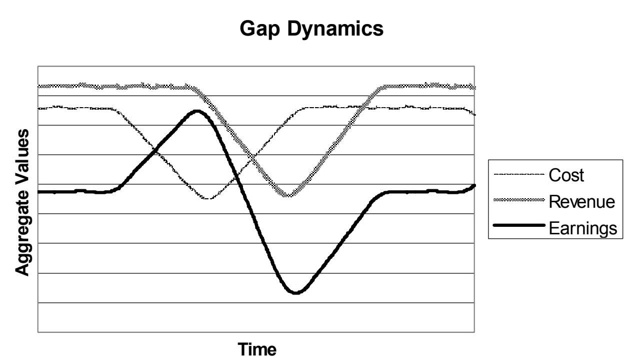

Figure 10. Gap Longitudinal Dynamics

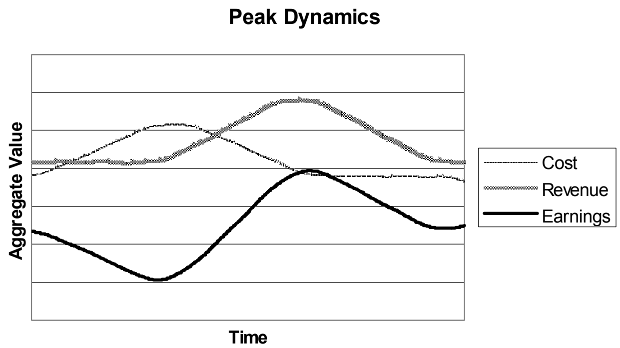

Figure 11. Peak Longitudinal Dynamics

The above graphs demonstrate the importance of smoothing investment and revenue performance to avoid major earnings swings that can be very damaging to an enterprise portfolio effort. This form of simulation allows the user to identify and correct Peaks and Gaps related to aggregate longitudinal dynamics.

ADVANCED ANALYTICS

Now that the structure of a baseline simulation is in place, it is possible to utilize some of the analytic approaches that are available. Stochastic Simulations, Real Options evaluations, Genetic Algorithm optimization, and Portfolio Theory can now be integrated with the enterprise portfolio simulation structure (Laguna) (Glover, Kelly, and Laguna) (Seget).

Stochastic Simulation

Because the simulation structure involves probabilistic factors, each time the simulation is run with a different random seed the result will be different. For this reason it is necessary to perform a stochastic analysis that captures the distribution of performance for a given strategy set.

The probabilistic distributions that result from the stochastic simulation of a single strategy set provide a window into the probable behavior of the systemic structure that is being evaluated.

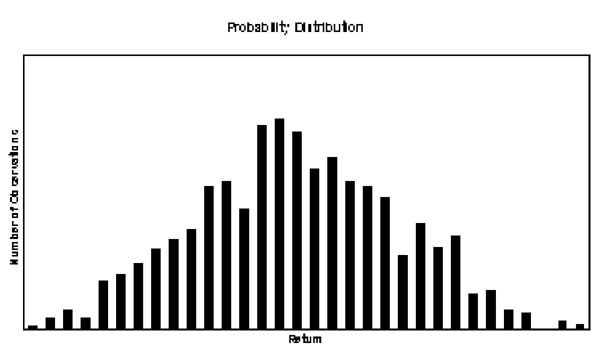

Figure 12. Stochastic Probability Distribution

Real Options Evaluation

NPD pipeline investments share many of the characteristics that define traditional financial securities. Because of these similarities it is possible to apply analytic methods that have deep founding in financial theory to the NPD process. Investment in a specific compound is similar to the investment in a stock in that it involves an investment coupled with an expected return that comes with an associated level of risk.

By breaking the investment in the individual compound out into a series of staged investments, it creates a structure that is very similar to a stock option. The initial investment will result in the opportunity to make an additional investment in the future that will result in an expected level of return associated with a certain level of risk.

Decision Algorithms and Hurdle Rates

Decision algorithms can be used to compare an expected NPV at a particular point in the NPD process with an arbitrary hurdle rate to determine whether investment in the compound should continue or be abandoned. Because there is uncertainty surrounding the likelihood of success for the compound, the expected NPV may differ from what will actually be observed. The simulation structure provides an environment where different hurdle rates can be tested under uncertainty, thus creating the ability to identify the optimal hurdle rate.

Abandonment Value

The underlying dynamics often result in a hurdle rate that is negative, meaning that compounds with a negative expected NPV that is greater than the hurdle rate will still be transitioned. This dynamic is the result of a staged investment structure where a current investment can advance the product until more information is known about the expected NPV. In the probabilistic circumstances where deviation from expected value is positive, the change in expected value can be capitalized on. Otherwise, the compound can be abandoned.

What is important here is that a full investment is not needed to get to the point of acquiring new information. Once the new information is available, the investment can be reevaluated. If the updated NPV has become more favorable, investment can continue. If the updated NPV is unfavorable, the investment can be abandoned without ever having to have made the full investment. The simple inclusion of the ability to abandon a compound at a certain stage increases the overall expected value of the NPD Pipeline.

Risk Management and Portfolio Optimization

Each of the preceding sections has built a framework that is necessary in order to support a more comprehensive evaluation of the entire constellation of investments in the enterprise portfolio. By taking the simulation structure that has already been created and combining it with stochastic simulation capabilities and an advanced non-linear optimization routine we are now able to evaluate strategy sets that can ultimately form the risk reward tradeoff that is defined by the Efficient Frontier.

Optimization Approach

At the core of our approach is a modified Genetic Algorithm that utilizes concepts gleaned from biological evolution to evaluate strategies in non-linear systems. While the system itself is non-linear, the highly efficient genetic algorithm can often yield near linear improvement from generation to generation as can be seen in the following graph.

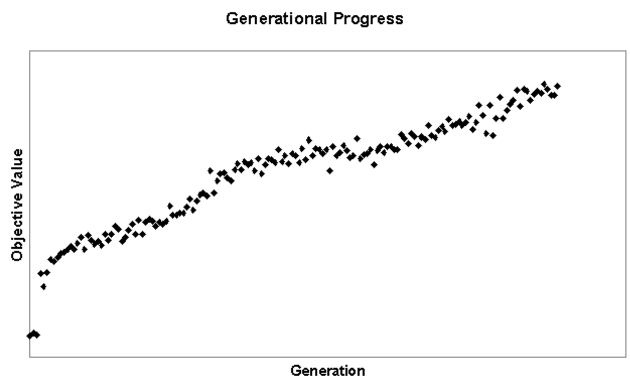

Figure 13. Genetic Algorithm Generational Progress

Efficient Frontier

By tracking the observed distributions that result from the stochastic simulations of each strategy set evaluated it is possible to derive the Efficient Frontier. This is simply an X-Y plot that graphs each strategy based on the mean value of its objective function and the standard deviation of its distribution.

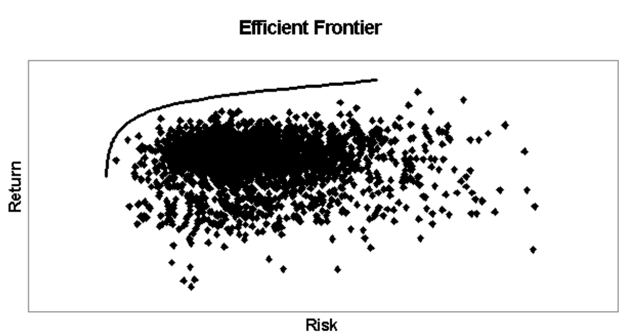

Figure 14. Efficient Frontier Plot

The simulation can keep track of the specific strategies that resulted in each individual plot on the Efficient Frontier. It is then possible to identify the optimal strategy set for a given level of risk. Furthermore, the shape of the Efficient Frontier can often lead to the identification of the "sweet spot" at which additional increases in risk result in minimal improvements to the mean of the objective function.

SIMULATION USING AnyLogic V™

Many of the approaches and techniques applied in this exercise have previously been exceedingly difficult to integrate due to the somewhat conflicting nature of the disciplines employed. By using functionality that is native to AnyLogic™, it is possible to seamlessly integrate all of these approaches into a comprehensive simulation.

Visual Object Oriented Development Environment

AnyLogic™ provides a visual object environment that facilitates rapid prototyping of simulations. Once an object is created it can be parameterized with minimal effort. For more sophisticated users it is possible to incorporate custom Java code or external Java libraries.

Multi-Approach Simulation Support

The simulation algorithms utilized by AnyLogic™ provide out of the box support for the incorporation of discrete event and continuous equation interactivity. The extension of this capability to include advanced optimization approaches further enhances the tool.

SUMMARY AND CONCLUSION

By combining multiple disciplines into a single simulation environment it becomes possible to create a comprehensive model that can evaluate an enterprise portfolio in detail previously not possible. This support for a holistic simulation eliminates the potential for locally optimal behavior that ultimately undermines the portfolio as a whole.

It should also be clearly noted that each aspect of the simulation can support varying levels of detail, allowing for the simulation to be customized to specific needs.

REFERENCES

Ding, Ming, and Eliashberg, Jehoshua. "Structuring the New Product Development Pipeline" Management Science 48, no. 3 (March 2002)

Glover, Fred, and Kelly, James P., and Laguna, Manuel. "The OptQuest Approach to Crystal Ball Simulation Optimization" University of Colorado, Graduate School of Business

Laguna, Manuel. "Optimizing Complex Systems with OptQuest" University of Colorado, Graduate School of Business

Loch, Christoph H., and Kavadias, Stylianos. "Dynamic Portfolio Selection of NPD Programs Using Marginal Returns" Management Science 48, no. 10 (October 2002)

Seget, Steven. "The Pharmaceutical Portfolio Management Outlook" Reuters Business Insight (2003)

Senge, Peter. The Fifth Discipline. Doubleday

Warren, Kim. Competitive Strategy Dynamics John Wiley & Sons

AUTHOR'S BIOGRAPHIES

Kirk Solo

Kirk is the founder and Managing Director of SimNexus LLC, a dynamic simulation consultancy. Kirk has over a decade of dynamic simulation of experience across several industry segments. SimNexus is a distributor of the AnyLogic^{TM} simulation package for business applications in North America.

Kirk holds an MBA from MIT's Sloan School of Management and a BA in Economics from The Colorado College.

Mark Paich

Mark has been a professor and practitioner of advanced forms of simulation for several decades. He holds a Ph.D. in System Dynamics from MIT, a Masters in Economics from The University of Colorado, and a BA in Economics from The Colorado College. Mark has published extensively in simulation journals, and has taught at the undergraduate and graduate levels.

Mark has implemented large simulations for several Fortune 100 companies and is widely respected as one of the leading simulation modelers in the world.