An Aggregate Capacity Model of the Pharmaceutical Research and Development Process

Paul R. Bunch
Eli Lilly and Company
Lilly Corporate Center
Indianapolis, IN 46285 - USA
Phone: (317) 277-4828
Fax: (317) 277-9271 (FAX)
bunch@lilly.com

Denny Park
Andersen Consulting
Andersen Consulting Northbrook Technology Park
3773 Willow Road
Northbrook, IL 60062
Phone: (847) 714-3344
Fax: (847) 326-3344
denny.j.park@ac.com

Abstract

The task of quickly bringing novel pharmaceutical products to market at reasonable cost is daunting. While both pharmaceutical sales and R&D expenses are on the rise, the latter is outpacing the former. In order to increase top line revenues, pharmaceutical firms have been employing a range of strategies to supplement the vitality of their R&D portfolios. While these strategies have the potential to increase top line revenues, they typically increase the demand for development resources that are already in short supply. We posit that the ability to accurately estimate and intelligently allocate scarce development resources is critical to controlling the cost of bringing new pharmaceuticals to market. In this paper, we consider development and implementation of steady-state and dynamic models to estimate resource requirements for the R&D process. Simulation results for a sample R&D pipeline will be considered.

1 Introduction

Bringing new pharmaceutical compounds to market is a challenging process involving many complex decisions. Issues that make the process challenging include long timelines from discovery to product launch, high attrition of molecules throughout the development process, and high development costs. For example, the time required to discover and develop a compound can exceed 10 years while the associated costs can exceed $500 million per compound (Pisano(1996); Drews(1998)). Fewer than 25% of the molecules that begin clinical testing actually make it to the marketplace. In addition to discovering and developing compounds internally, most pharmaceutical companies have sought to diversify their development portfolios by relying on external partners to co-develop and market their products. Furthermore, most pharmaceutical companies are attempting to market products in multiple dosage forms and for multiple disease states. While these diversification strategies create a richer set of opportunities to pursue, they also create demand for scarce R&D resources required to bring

products to market. Clearly those firms that are able to use their R&D resources most productively will be rewarded in the marketplace.

Wheelright and Clark (1992) propose a framework for product development that contains as a key construct the *aggregate project plan (APP)*. A primary purpose of the APP is to establish the types and mix of projects that should comprise a firm's development portfolio over a time horizon. The APP also allows one to balance the demand for critical development resources that are required to execute the plan with available capacity. Wheelright and Clark consider only specific development projects in their definition of the APP. We relax this definition to include not only specific projects currently in our consideration set, but also projects that might enter our consideration set as a result ongoing research and development or partnerships, for example. The key concept is that the resource planning process requires that one have an estimate of the type and mix of projects that will be pursued over some time horizon.

The approach suggested by Wheelright and Clark for determining the resource requirements corresponding to a set of proposed projects is *project centric*. That is, at any instant in time, one can determine the resource requirements for the set simply by summing the requirements across all projects in the set. This approach simply requires that one know the resources required for the remaining work for each project in the set.

We propose a *task centric* approach to determining resource requirements for a set of projects under consideration. That is, we consider the R&D process to consist of a set of tasks, each of which consumes resources and whose function is either to filter diverse opportunities or eliminate uncertainty in some feature of the product such as safety or efficacy in the case of pharmaceuticals. This view of the R&D process is depicted in Figure 1.

Error! Not a valid link.

Figure 1: A Staged R&D Process

Note that each stage in this abstraction of the R&D process could consist of any number of more detailed stages. In simple terms, as projects proceed from idea to commercialization, they are either terminated due to some undesirable feature or they progress to the next stage of development. We assume then that all projects are either launched as commercial products or are terminated. Determining the resource requirements for projects in this process at any instant in time simply consists of summing the requirements across all stages of R&D. This approach requires that one knows the number of projects in each stage over time and the resources required to execute a single project in each stage of the development process.

In this paper, we present both steady state and dynamic capacity models for multiple key resources required for efficient completion of a set of R&D projects. The steady state model is used to determine nominal resource requirements to achieve a steady output of products to the marketplace from R&D. The steady state model assumes that the projects are completely balanced along the R&D chain so that a steady output can be achieved. While this is clearly an idealization, it is useful to help determine what fraction(multiple) of resources currently available are necessary to achieve a given output from R&D. The dynamic model is used to determine the time variant resource requirements as individual projects progress through the development process. The dynamic model accounts for the distribution of projects along the R&D chain in an initial project set. In general, this distribution will not be balanced so that a steady output can be achieved from R&D. Additionally, the dynamic model accounts for the addition and deletion of projects from the set. Projects can be eliminated from the set due to undesirable product qualities such as adverse safety or efficacy results. Projects can be added to the set from internal discovery efforts, in-licensing efforts, or initiation of a line-extension project.

The remainder of this paper is organized as follows. Section 2 provides background and concepts, including an abstraction of the R&D project flow and a method for estimating resource requirements for executing R&D projects. Steady state and dynamic resource estimation models are introduced in Section 3. Model structure is provided along with a description of the function of various model substructures. Implementation of the models will also be described briefly. In Section 4, simulation results are considered for several sample R&D pipelines.

2 Background and Concepts

The abstraction of the R&D process shown in Figure 1 serves as a good starting point, but it does not go far enough. Specifically, in addition to working on projects that will produce new products generated from internal discovery efforts, most pharmaceutical R&D programs undertake efforts that extend the use of existing products. The extension could include efforts to pursue using a single product (molecule) to treat multiple disease states or to administer the product in multiple dosage or delivery forms (capsules vs. intravenous). Furthermore, R&D programs may add projects to their set by licensing molecules or technologies from partner companies. Thus an R&D program could have in its set of projects several types of projects originating from several sources. This is depicted in Figure 2. As before, each stage in this diagram could consist of any number of more detailed stages. Note that the characteristics of projects will most likely depend upon their source. That is, the demand for resources by line extensions may differ from the demand for resources from new products developed internally.

Error! Not a valid link.

Figure 2: A Staged R&D Process with Line Extensions and In-Licensed Projects

Critical functional resources are required to complete work at each stage of the development cycle. In the pharmaceutical R&D process, these functional resources include clinical pharmacologists, toxicologists, chemists, statisticians, and clinical research physicians, among others. In most instances, these resources are drawn from functional areas in which areas of technical expertise are concentrated. Completion of work in a given stage requires a pool of resources, which will generally be comprised of a variety of functional resources. Any given function then can be called upon to support multiple projects over a range of development stages in the R&D process. This situation is depicted in Figure 3, in which solid lines represent the flow of projects and dotted lines indicate the input of resources to complete a project at a given step.

Given the representation shown in Figure 3, one can estimate the resource requirements for executing a set of projects at any instant in time by knowing the number of projects by type in each stage of development and the resources required per project by type and stage of development. Specifically, resource requirements can be determined as

$$r_{i,Tot}(t) = \sum_{S} \sum_{T} WIP_{S,T}(t) \times r_{i,S,T} \qquad (1)$$

where

$r_{i,Tot}(t)$ = total number of resource i required at time t

S,T = indices referring to development stages and project types, respectively

$WIP_{S,T}(t)$ = work in process (number of projects) of type *T* in development stage *S* at time *t*

$r_{i,S,T}(t)$ = number of resources of type *i* required to complete a single project of type *T* in development stage *S*

Note that equation (1) can be used to determine the resource requirements at any instant in time *t*. The essence of estimating the resources required at time *t* is the determination of $WIP_{S,T}(t)$. Furthermore, the resource requirements can be compared to the available resources and planned additions to quantify utilization levels over time. Significant differences in resource utilization from target levels indicate that interventions might be necessary. These could include the addition or reduction of headcount, re-engineering key business processes, modifying the project portfolio by the level of in-licensing, terminating projects, or outsourcing selected activities of key functional resources.

3 Capacity Models

3.1 Steady State Capacity Model

The aim of the steady state capacity model is to determine the resources required to support a constant desired output from R&D. That is, the steady state model could help one determine the number of each functional resource needed to support two product launches per year, for example. The data required to drive the steady state model consist of the task specific data – cycle time, success rate, and resource requirements for each functional resource, and global process data – number of product launches required over a time horizon and project mix (line extensions vs. new products).

Calculation of the quantity $WIP_{S,T}(t)$ for the steady state model is straightforward. Hopp and Spearman(1996) present Little's Law, which relates the throughput *(TH)* of a process to its

Error! Not a valid link.

Figure 3: R&D Process Flow with Supporting Functional Resources

work in process (*WIP*) and cycle time (*CT*). We use a generalization of this relationship to determine WIP at each stage of the R&D process. Specifically, the relationships used to calculate WIP at stage *j* is:

$$WIP_j = \frac{TH_j}{SR_j \times CT_j} \tag{2}$$

$$TH_j = \frac{TH_{j+1}}{SR_{j+1}} \tag{3}$$

In this context, we define the throughput of stage *j*, TH_j, to be the number of projects that move from stage *j* to stage *j+1* for further development. Thus for the N-staged process shown in Figure 4, one can calculate *WIP* for all stages by setting the desired output rate from R&D, TH_N, and recursively applying equations 1-3.

Error! Not a valid link.

Figure 4: An N-Staged Sequential R&D Process

The steady state capacity model is a useful aggregate tool to assess the resource requirements to under a range of conditions to achieve a given output from R&D. Furthermore, it can be easily implemented in a spreadsheet and its concepts are easily communicated. Its shortcomings lie in its simplifying assumptions. In particular, the steady state model assumes at all times that the R&D process is balanced with respect to WIP and as a result does predict fluctuations in requirements over time. R&D project portfolios are rarely perfectly balanced, however, since there is uncertainty whether a given project will move to the next stage or be terminated. At times more projects than expected will progress, at other times fewer than expected will progress. Furthermore, the steady state model treats the flow of projects through the process as continuous when the reality is that partial projects do not move from one stage to another – a project either moves to the proceeding stage or it is terminated. Additionally, many factors affecting resource requirements change over time and are difficult to effectively track. These factors include the inclusion of specific projects in the project set, headcount on-board, and the number and mix of projects entering the R&D process.

3.2 Dynamic Model

A system dynamics capacity model was developed to address the shortcomings of the steady state model. An additional goal of developing a system dynamics model was to constrain the movement of projects through the process based on resource availability so that a true assessment could be made of the throughput capability of R&D. Furthermore, this would allow us to de-bottleneck the R&D process. Our strategy was to develop two complementary models that would represent the same process structure, but that would treat predict the performance of different sets of projects. In particular, we track the performance of specific projects in our current project portfolio in one model and the performance of projects not yet in our portfolio in another model. These sub-models are then combined to determine the aggregate performance of the R&D portfolio. Additionally, a resources sub-model was incorporated so that the time dependent resource requirements and availability could be tracked and used to control the flow of projects through the system.

Conceptual stock-flow diagrams representing the basic structures of the current/future projects sub-models and resources sub-model are shown in Figure 5 and Figure 6, respectively. The stock-flow diagrams use the convention that shaded variables originate in other models and thus capture interactions between models. Variables that are underlined and italicized are exogenous data. As an example consider the resources sub-model that is shown in Figure 6. The stock named “Stage” is shaded and represents the aggregate of the stocks from the current projects and future projects sub-models that track the number of projects in each stage of development. Furthermore, the variable “Resources Reqd Per Project/Stage” is input data read from an exogenous data file.

3.2.1 Master Model

The function of the Master Model is to interact with external components such as databases, spreadsheets, and results files, control the execution of the sub-models, send/retrieve data to/from the sub-models, and calculate summary metrics. The detailed model calculations are performed

in the sub-models, but the Master Model allows the sub-models to interact in a meaningful way with each other and with external components.

3.2.2 Current Projects and Future Projects Sub-Models

A conceptual stock-flow diagram that represents the Current and Future Projects Sub-Models is shown in Figure 5. The purpose of the Current Projects Sub-Model is to simulate the movement of projects currently in our R&D pipeline to their ultimate destination. The Future Projects Sub-Model serves an identical purpose for projects that are likely to enter our R&D pipeline. All projects currently in development or that enter over time will either exit as a launched product or be terminated due to some unfavorable finding. At each simulation time step, the Current and Future Projects Sub-Models inform the Master Model of the location of each project in the original and potential future projects list, respectively. This information is needed to determine the resources required by projects in each list.

The Current Projects Sub-Model is initialized with information sent by the Master Model. This information includes a list of the current projects in our portfolio, the current location of each project(stage of development), the estimated time each project will spend in each subsequent stage, and the likelihood that each project will progress through each of the remaining stages of development from its current stage.

Although we are not aware of the identity of projects in the future projects set, we make assumptions about the number of these projects that might enter our pipeline from various sources. Furthermore, we assume a range of values for the cycle times, success rates, and resource requirements for these projects. These estimated or assumed values are based on either our corporate performance objectives or reasonable assessments of what is achievable.

Note that we represent the flow of projects through our process recursively. Consider the N-staged process shown in Figure 4, for example. When work is completed on a project at stage j, the project is either terminated (killed) or proceeds to stage $j+1$. If a project is successful after completing work in stage N, then it exits the process as a successfully launched product. Note that projects may enter the R&D process at any stage of development. In particular, line extensions of existing products and in-licensed projects only participate in a subset of the tasks in the entire R&D process flow. Furthermore, projects currently in the pipeline complete only the subset of steps remaining in the R&D process. Interactions in these models exist when the flow of material is constrained by resource availability since all stages can compete for the same functional resources.

3.2.3 Resources Sub-Model

The resources sub-model takes purpose in accounting for the resource requirements and availability at all times during the simulation. Resources can be added to the available pool when

Error! Not a valid link.

Figure 5: Stock-Flow Diagram for the Current/Future Projects Sub-Models

Error! Not a valid link.

Figure 6: Stock –Flow Diagram for the Resources Sub-Model

work in a given stage has been completed and when new hires are brought on-board. Note that we include only net hires – that is, we do not explicitly model resources that leave the pool due to non-project related reasons. We account for hiring freezes, maximum allowable hiring levels, and include explicit recruiting goals. Resources are depleted from the available pool when projects are allowed to pass from a queue to a development stage.

Note that we consider instances of the model in which resources either limit or do not limit the flow of projects through our development process. Although resources always limit the flow of projects, we wanted to determine the resource requirements in both the limiting and non-limiting cases. When run in the resource-limiting mode, the resource model serves as a feedback mechanism since common resources are required at various stages along the R&D process flow. It is this feedback that allows us to get a more realistic estimate of our throughput capabilities.

3.3 Implementation

Our implementation strategy was to make both the steady state and dynamic capacity models easily accessible to users with a wide range of skills. That is, the objective was to give a wide range of users access to the powerful modeling capabilities without burdening them with the arcane details. Accordingly, intuitive user interfaces were constructed in a spreadsheet. The interface consists of a flow diagram representing the R&D process. Parameters such as R&D performance data, functional resource data, and corporate objectives data can be entered directly from the interface or pulled in from enterprise databases. The interface also contains information on key output metrics such as number of projects in each stage and number of resources required for the projects in each stage of work. Data describing projects in the current project list were maintained in a separate spreadsheet. The dynamic model was run directly from the spreadsheet to shield users from the complicated stock flow diagrams of the Master Model and Sub-Models. Simulation results from multiple runs were collected, analyzed within the spreadsheet, and summarized in a separate output spreadsheet.

4 Results and Discussion

The essence of this work has been to develop a framework within which the resources required for R&D activities can be explicitly considered. The models presented in this paper can help estimate the resources implied by a plan or objective for R&D. Alternatively, the models can test the feasibility of R&D plans and objectives when resource constraints are considered. In particular, the models are used for both for long-term strategic planning and short-term business planning. The strategic planning process includes setting high-level corporate growth objectives and estimating the implied resource requirements. Both the steady state and dynamic models can be used for this purpose. The short-term business planning includes estimating resource requirements to efficiently complete currently active projects and projects that may be enter in the near future. The dynamic model is better suited to handle this problem since it can consider data for specific projects and can track their progress over time.

Simulation experiments were run under a variety of conditions to validate the dynamic model and to draw some general conclusions about the R&D process under consideration. For purposes of this discussion, we consider a specific instance of the process shown in Figure 4 in which N is equal to five (5). General processing characteristics, which will be used for the future projects list, are given below in Table 1. Note that the process described by the data in Table 1 is one that becomes increasingly complex as it progresses through development. This is evident by the increasing time and resources required for completion of work in the later stages. Data for projects in the current project list take the form shown in Table 2. For our present purposes, we conduct simulation studies by fixing the parameters in Table 1 and considering the impact of distributing projects in the current projects list. We argue that studying this system could suggest

what resource requirements might be along a transition path to implementing a long-term strategy. That is, we can imagine that at some future time, through process improvement initiatives, we are able to provide a more steady stream of projects into our pipeline. Consistent future process performance would help maintain a favorable balance of projects along the R&D process chain. Recall that a perfectly balanced pipeline is one that satisfies equations 2 and 3. Even if this future state of nirvana is reached, we still are faced with projects currently in our pipeline. Typically, the projects are not in perfect balance, and there is variability in their processing characteristics.

It is worthwhile to consider how the R&D system might respond to various initial project portfolios. In particular, we note that if the dynamic model is initialized with a perfectly balance portfolio, then the time-averaged resource requirements should match the resource requirements implied by the steady state model. We assume that future projects exhibit average behavior. In particular, the operational parameters of future projects are identical to those of the steady state model shown in Table 1. Regardless of the distribution of projects in the pipeline initially, therefore, the long-term behavior of the dynamic model should be identical. More specifically, the infinite-limit time-averaged resource requirements of all dynamic models we will construct will match the resource requirements implied by the steady state model. The value of the dynamic model, then, is to estimate the resource requirements as the pipeline transitions from its current state to the steady state.

Stage	Processing Time (mos)	Success Rate (%)	# Resources Per Project	Steady State # Projects Per Unit Output	# Resources Per Unit Output
1	9	90	1	6.5	6.5
2	6	25	1	3.9	3.9
3	9	80	3	1.5	4.5
4	12	85	9	1.6	14.4
5	18	75	15	2.0	30.0

Table 1 : Staged R&D Process Data

Project	Current Stage	Time in Stage 1	…	Time in Stage 5	Success Rate in Stage 1	…	Success Rate in Stage 5	Resources For Stage 1 Work	…	Resources For Stage 5 Work
1	1	9	..	12	95	..	90	0.5	..	12
2	1	3	..	24	75	..	75	1	..	10
3	2	NA	..	18	NA	..	85	NA	..	15
4	4	NA	..	22	NA	..	80	NA	..	16
5	3	NA	..	20	NA	..	60	NA	..	18
.	..	NA	..	16	NA	..	50	NA	..	15
.	..	NA	..	18	NA	..	90	NA	..	12
.	..	12	..	12	80	..	75	1.5	..	10
.	..	NA	..	18	NA	..	70	NA	..	10
100	5	NA	..	20	NA	..	80	NA	..	15

Table 2 : Staged R&D Data for Current Projects

Figure 7 and Figure 8 show representative simulation data for the dynamic model. In particular, Figure 7 shows the initial distribution by stage of a fixed number of projects in the R&D portfolio. That is, this represents the current project list. Case 1 represents a portfolio that is initially perfectly balanced. Cases 2 and 3 represent pipelines in which all projects begin in the first and last stages, respectively. Case 4 represents a pipeline in which projects are distributed throughout, but which is not properly balanced.

The resource requirement profiles corresponding to these initial project distributions are shown in Figure 8. Note that since the current projects list for case 1 is perfectly balanced, it follows the steady state line. After a transition period, the resource profiles of the other cases follow the steady state requirements. These results match our intuition. Since projects in earlier stages require less resources, we expect the resource requirements for Case 2 to start low and gradually increase as projects progress to the later stages, where resource requirements per project are much higher. The resource requirements for Case 3 start high since all projects start at the last stage of development. As projects in the current projects list exit the process, the pipeline is depleted of high resource consuming projects. Therefore resource requirements drop off dramatically and gradually reach steady state as new projects make their way through the process. Case 4 is essentially a hybrid of the previous cases. Note resource requirements increase dramatically as the large number of projects initially in stage 3 progress to the latter stages of the pipeline and decrease as these projects exit the system.

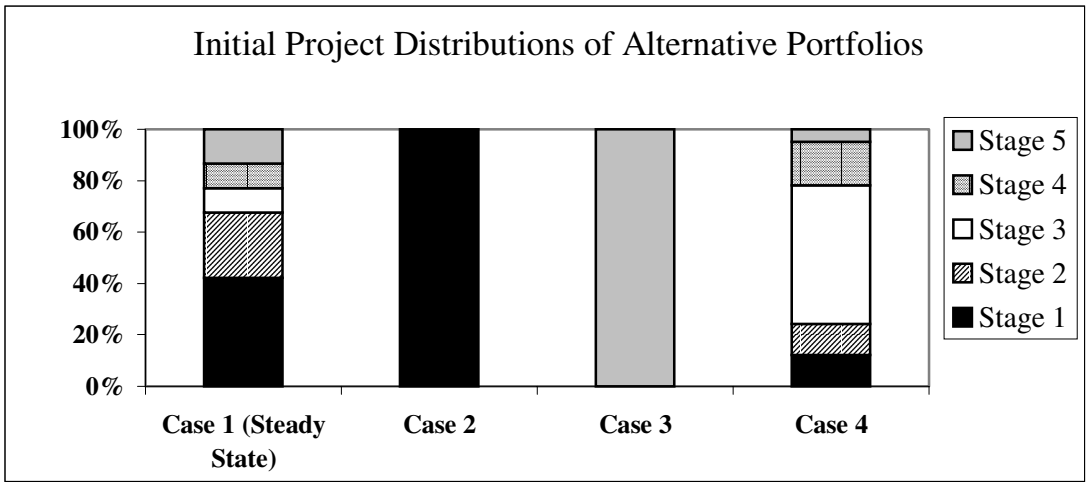

Figure 7: Initial Project Distributions for Alternative Cases

5 Summary and Conclusions

Aggregate steady state and dynamic models have been presented to estimate the time variant resource requirements to execute projects in a staged R&D process. These models can be used to predict resource requirements at an aggregate level for general stages R&D processes. The steady state model is attractive because of its simplicity and its ability to set target resource levels to achieve a given level of R&D output. The dynamic model is useful when incorporating both current projects and future projects in the consideration set. In all cases, the average resource requirements predicted by the dynamic model match those of the steady state model. When the current project set is not perfectly balanced, the dynamic model will reveal a resource requirement profile that is quite different from the steady state.

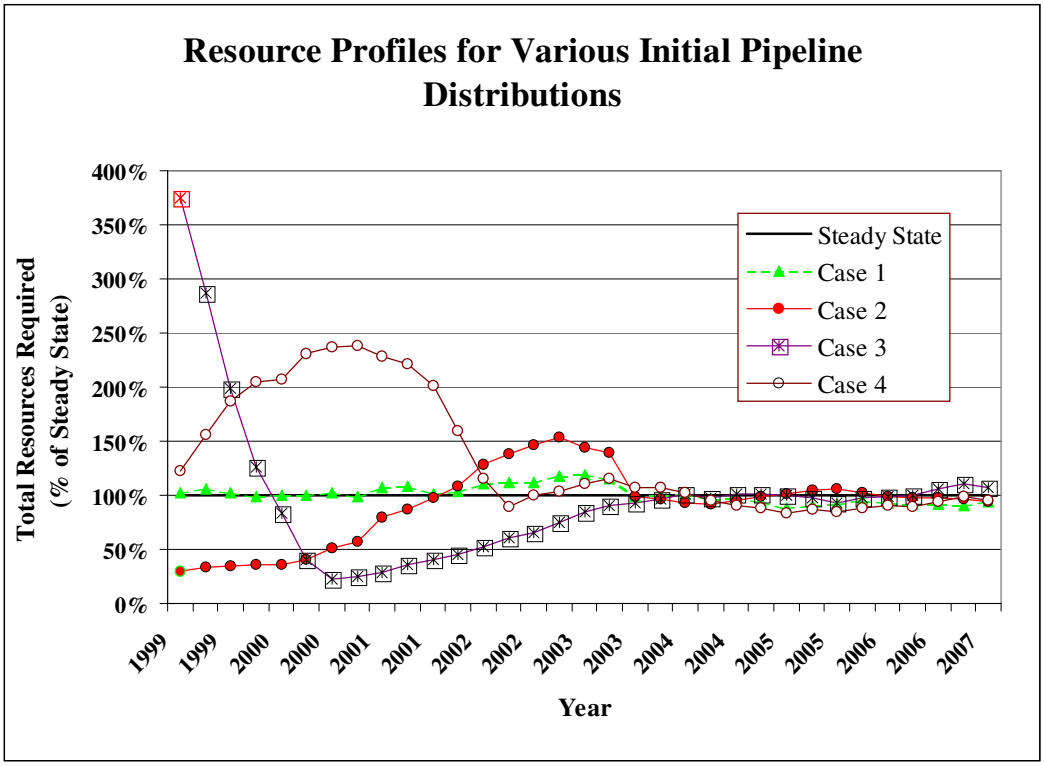

Figure 8: Time Variant Resource Requirements for Alternative Cases

References

Drews, J. (1998). Innovation deficit revisited: reflections on the productivity of pharmaceutical R&D. *Drug Discovery Today,* Vol. 3, No. 11

Hopp, W.J. and Spearman, M.L. (1996). *Factory Physics,* Irwin, Chicago

Pisano, G.P. (1996). *The Development Factory,* Harvard Business School Press, Boston.

Ford, D.N. and Sterman J.D. (1998). Dynamic modeling of product development processes. *System Dynamic Review.* Vol 14, No. 1

Wheelright, S.C. and Clark, K.B. (1992). *Revolutionizing Product Development,* The Free Press, New York