THE IMPACT OF PROJECT INTRODUCTION HEURISTICS ON R&D PERFORMANCE

Autoria: Paulo Soares Figueiredo, Fernando Augusto Almeida Neves

Abstract

We use a dynamic model to explore managerial behavior in the pharmaceutical industry, with a focus on the impact of project introduction heuristics on product development performance. In particular, we focus on the impact on performance, of three types of heuristics: 1)Gradual increase or decrease, 2)Random-normal choice, and 3)Target-based search. We find that a gradual decrease of project introduction rates results in convergence, but the size of the adjustments and the delays in the pipeline can limit the precision. A random choice is detrimental to performance even when the average value is the optimal. A target-based search results in oscillation.

Keywords: Managerial behavior; project portfolio management; product pipeline management; heuristics; behavioral operations management.

1. Introduction/ Hypotheses

Previous studies in the innovation and product pipeline management (PPM) literature have examined factors influencing various dimensions of R&D performance, such as quality, cost, lead time and value created. However, still little is known about how managerial decisions affect performance in a dynamic setting and across the New Product Development (NPD) Pipeline. Assuming a fixed total R&D budget, the PPM problem is twofold: first, decide which projects to start, and then, decide which projects to continue and which to terminate at various stages of development, deciding how much to invest on each project at each phase and how to allocate people across the stages of the process. In making these decisions, managers face a set of tradeoffs between risks, returns, and time horizons for payoffs (Gino and Pisano, 2006). In theory, such tradeoffs are optimization problems that can be tackled with a technique such as dynamic programming. In reality, the complexity, ambiguity, and uncertainty of most companies' R&D portfolios make this an essentially impossible optimization problem to solve (Lockett and Gear,1973), *at least in closed form*. A few studies focused on behavior and heuristics on the scheduling of projects at a specific stage, such as Varma et al. (2007), Yan and Wu (2001), Kavadias and Loch (2002, 2004) and Gino and Pisano (2005). But these studies did not focus on project introduction policies across the product development pipeline, at the portfolio level.

Most decision-theoretic models proposed in the literature, however, are themselves highly complex and, as a result, they have not become a tool that is commonly used in management practice (Loch and Kavadias, 2002). Given the complexity of the problem of both portfolio selection and management, and individuals' bounded rationality (Simon, 1956), it is not surprising that companies utilize heuristics for managing their R&D portfolios rather than trying to optimize (Gino and Pisano, 2006). The idea that heuristics rather than optimization drive R&D decision-making was introduced by March and Simon (1958). While behavioral studies on PPM have become well accepted, research on the impact of specific heuristics on R&D performance is limited (Gino and Pisano, 2006).

Organizations often commit to more product development projects than they can handle. Over-commitment of development resources (i.e. when too many projects are introduced into the pipeline) is a common phenomenon as evidenced by case studies. Evidence suggests that many organizations have far more product development projects in progress than their capacity (Gino and Pisano, 2006). For instance, Wheelwright and Clark (1992) mentioned organizations tend to pursue a larger number of projects than they have the resources to fund and suggested that companies often operate their development organizations at 200-300% capacity utilization. Ash (2009) finds that loading a resource pool to 300% or 400% of capacity while allowing preemption may be good for engineering talent utilization rate; however, it is very bad for hitting individual project due dates. The latter author, however, did not focus on the relationship between capacity utilization and the quality of development activities.

For most firms that operate with high capacity utilization rates, the simplest form of heuristics would be to gradually decrease the project introduction rates (here also called "starts") from the high levels to lower levels, aiming to balance the pipeline and increase value creation.

Yu et al. (2010) developed a simple "static" model of the product development pipeline that establishes upper limits for the capacity to develop and launch new product families. This ''right number'' of projects may work as an alarm for firms that are trying to develop and launch too many product families. Figueiredo and Loiola (2012) and Figueiredo and Joglekar (2007) reached similar conclusions with a dynamic model that established a concave relation between number of projects started and total value created in the pipeline. Based on another dynamic model, Repenning et al. (2001) showed how a surge of resource demand can cause havoc in the NPD process in the phenomenon known as firefighting.

The traditional approach to the problem of over commitment is to develop better models for project management and more sophisticated in-process management tools (such as real time scheduling) and to undertake more planning activities. Gino and Pisano (2006) suggest that these models will be more useful if they rest on cognitively and behaviorally compatible assumptions, i.e. incorporating into the models elements that will reduce common cognitive biases which people incur in their decisions.

In this paper, we use a dynamic simulation model to explore such phenomena in the pharmaceutical industry, with a specific focus on the impact of project introduction heuristics on NPD performance. The use of heuristics is a way of *searching* for the optimal policy and/or of making necessary adjustments whenever there are changes in the shape and performance of the pipeline. In particular, we focus on two types of heuristics: 1) Gradual decrease or increase, and 2) Random Normal choice. The impact of these heuristics on performance, and their efficiency in optimizing the PPM process are discussed. We argue that the over commitment of resources distances managers from the optimal decisions. But *under commitment* of resources is also detrimental to performance. A search for the optimal number of projects to be introduced into the pipeline, based on a gradual increase (or decrease) of project introduction rates, results in convergence, but the size of the periodic adjustment and the delays in the pipeline can limit the precision of the process. A small periodic adjustment is more precise; however it slows down the search process. A random normal choice is detrimental to performance when the average value is optimal. A target-based adjustment has poor performance due to the long delays in the pipeline, which can cause oscillation in the chain. Such effect was detected in other chains, such as in supply chains (Sterman, 1995; Goodwin and Franklin, 1994).

2. The model

The basic structure and logic of the model are simple; every month, a certain number of projects is started and are introduced into the pipeline. These projects are developed and screened in sequence, before being released into the marketplace. The NPVs of the population of projects are tracked, enabling managers to decide which fraction of projects will be terminated and how much value will be lost due to termination. Value creation happens while projects are developed at each stage, and this value creation depends on how intensively the teams are working. The *average* NPV of the population of projects is also increased by the process of screening, since only those projects with higher value will be approved to the next stage. Besides deciding on which projects will be terminated (i.e., defining a screening threshold or minimum allowable NPV for a project to be approved), managers also decide on three variables: capacity adjustment, resource allocation across stages, and average complexity of the projects. Each of these variables affects capacity utilization (how intensively the teams are working) and therefore the value creation. For the purposes of this study, we assume that managers adapt the work intensity of the development teams as needed, and work faster/slower depending on the number of projects at hand. This adjustment of work intensity has an impact on the value creation rates at each stage (see figure 2.3).

The *resource allocation* bias reflects managers' tendencies to allocate more people to work on the initial, mid or final stages of the pipeline. Managers also have a bias towards *allocation of complexity*, i.e. they can have a tendency to increase/decrease the complexity of the projects in any stage of the process. The complexity of projects can be measured in many ways depending on which kind of product is being developed (lines of code for a software, number of parts for a car, etc.) but in this study the average size of the projects (in terms of investments) is adopted as a measure and proxy for complexity, meaning that a more complex project would require more design and developing activities to be performed. Such a proxy was also adopted by Yu, Figueiredo and Nascimento (2010).

The performance variables in the model are total value created (NPV) at the end of the pipeline, value creation rates at each stage and respective flows of projects. The adoption of NPV as the only performance criteria for project screening is a necessary simplification; in most companies, however, more than one factor is used to enable the decision to terminate a project, and different factors may be used depending on the stage of development of the project. For example, a pharmaceutical company might be more concerned with the safety of a substance at the early stages and with manufacturability at later stages.

The PPM problem is structured as a dynamic process in the shape of a chain, therefore it is reasonable to assume that accumulation and/or starvation might happen in such a chain. Depending on the decisions made by managers, projects may accumulate in early stages, or the later stages may starve in case too many projects are terminated in early stages. The dynamic aspect of the pipeline adds complexity to the problem, and to the optimization effort. The model structure is comprised of three processes: capacity management, value creation and screening in any stage of the pipeline. See figure 2.2 the stock and flow structure of a typical stage. These processes are described in detail in Figueiredo and Loiola (2012) and Figueiredo and Joglekar (2007).

Figure 2.1: A Multi Stage Product Development Pipeline

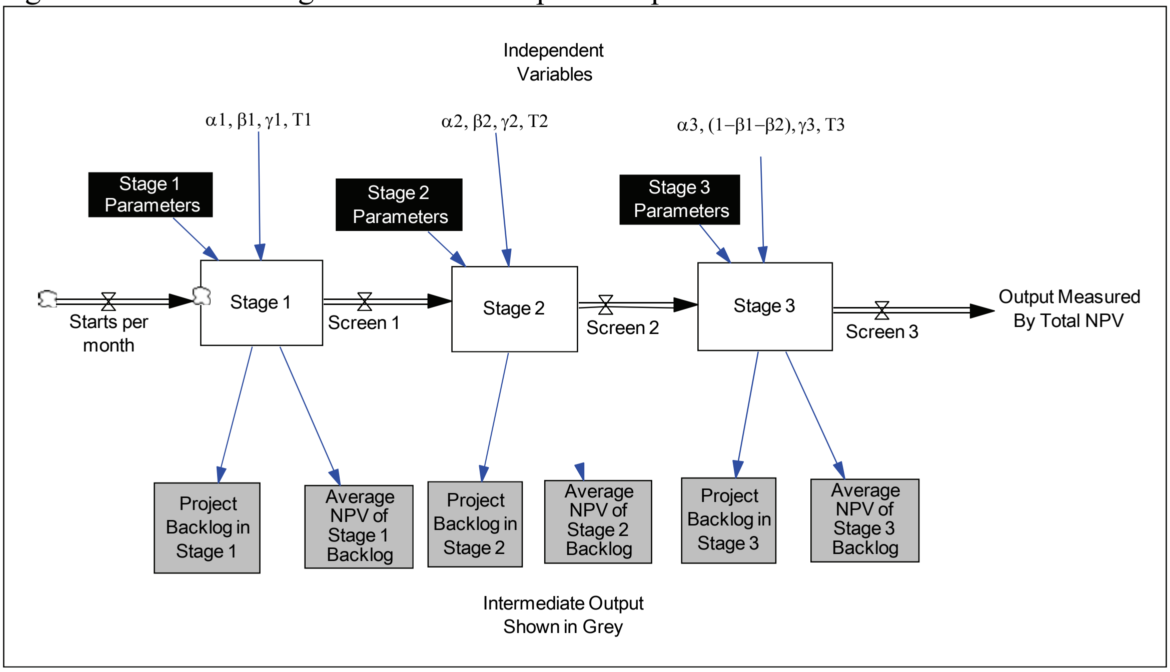

Figure 2.2: Stock and Flow Structure of a Typical Stage

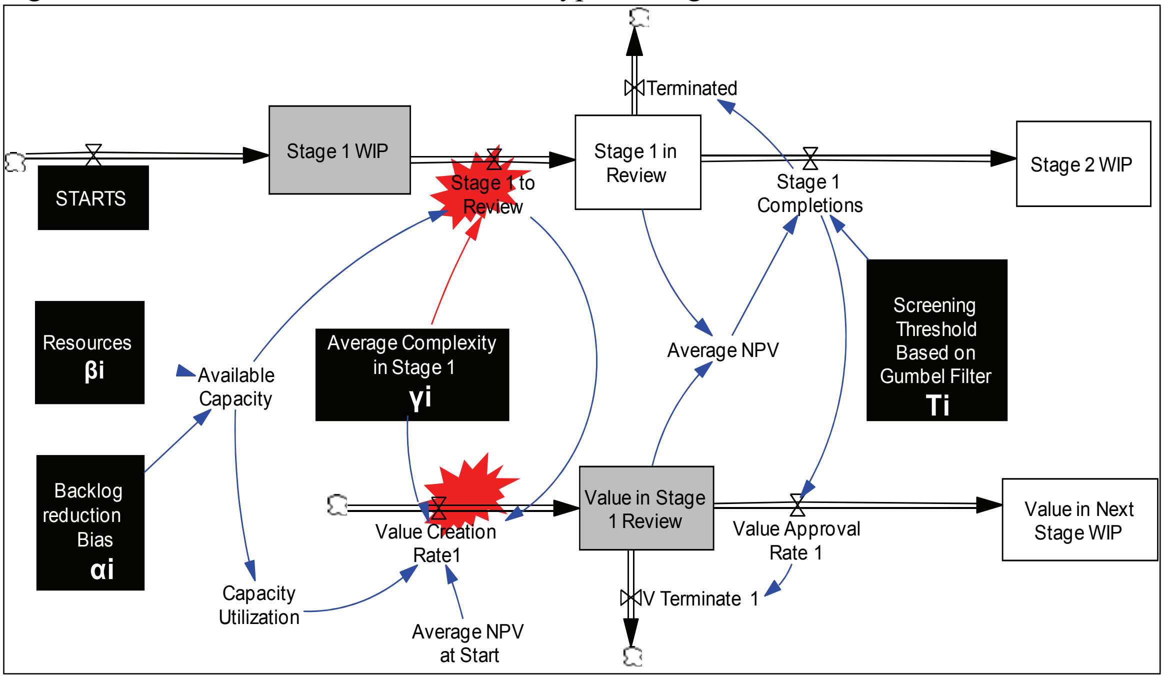

2.1 Model Operationalization

This section presents the three basic processes in the PPM model. A more detailed description can be found in Figueiredo and Loiola (2012) and Figueiredo and Joglekar (2007).

2.1.1 Capacity Management Process

A central construct of the model is the utilization of capacity. Research shows that employee productivity (percent of time spent on *value-adding* tasks) initially increases and then decreases as the number of development activities assigned concurrently to each engineer increases (Wheelwright and Clark 1992, pg. 91). This effect is captured in a function that links utilization and value created.

Managers have at each stage a fixed amount of resources (employees). Assuming that the allocation of resources is fixed, an increase in capacity, measured by man hours per month, is only possible by using the existing resources more intensively, thereby increasing their utilization. In case of overcapacity, the utilization equals to the demanded capacity based on the backlog.

Capacity is adjusted continuously, depending on the value of the target capacity and on the time to adjust capacity. Target capacity is defined as the demanded rate of development in each gate based on the backlog. If the backlog is filled with projects, the target capacity will be higher, resulting in more work intensity or capacity utilization by the teams.

2.1.2 Value Creation Process

The available capacity is used within each stage as shown in figure 3.2, during the process of value creation. A certain number of projects enter stage 1 backlog. The value of the projects is tracked by the model, along with their number. The NPV value of the projects is multiplied by a factor, depending on the capacity utilization, as the projects that were in the backlog are developed and go to the next phase to be reviewed. The rate "stage to review" is equal to the available capacity, unless there is overcapacity. The projects then reach gate 1, or "stage 1 in review". In this phase projects are reviewed, and depending on the average NPV some fraction will be terminated and the rest will "follow the flow" to the next stage, the backlog of stage 2.

Projects that are approved in the third phase are launched to the market. The values of total NPV created, number of projects and average NPV of finished projects are tracked and used as performance measures. These calculations have been simplified by assuming that the time discounting effect is built into the "Average NPV at Start" parameter.

2.1.3 Project Screening Process

The average NPV of the projects feeds into the screening process: the decision to proceed or terminate a fraction of projects is made depending on the average NPV and a predetermined threshold. The population of NPVs of projects after a review is assumed to follow a Gumbel distribution, because project screening is a search process that selects NPV extreme values (Gumbel 1958, Galambos 1978, Dahan and Mendelson 2001). The Gumbel distribution is the probability distribution for the maximum of multiple draws from exponential-tailed distributions. It applies to NPD problems especially well when there are no specific limits on the potential NPV of a project (Dahan and Mendelson 2001).

3. Project Introduction Heuristics

We focus on three types of heuristics for project introduction; these heuristics are simplified ways of trying to balance the pipeline, searching for the "right" level of project introductions. First, we examine how gradual, monotonic adjustments of the project introduction rates affect performance, and if such adjustments are effective in reaching the optimal level of "starts". Since most development teams operate at a high level of work intensity or capacity utilization, it is interesting to check if a gradual decrease from a high level of starts can be effective in the search for the right policy. The impact of a gradual increase, from an initially low level of starts, is also studied.

Second, we investigate whether a random choice for project introduction rates can be effective, once the *average* value of the random choices is close to the optimal choice. It is assumed that managers have a benchmark for the project introduction rate, but there is uncertainty in the process, so that the number of starts varies around the benchmark. A normally distributed choice for starts is used.

Third, we discuss the impact of target based adjustments on performance. Such heuristic takes into account present and previous performances, changing the direction of the adjustments whenever performance decreases. We describe how the PPM heuristics were modeled next.

3.1 Gradual Monotonic Decrease

This simple heuristics adjusts the number of starts gradually, reducing their number by a constant value ("delta adjustment") every period. The number of starts has a high initial value of 60 projects per year, and is gradually reduced. The performance of the previous period, in terms of value creation rate at the end of the pipeline, is compared to the present performance, and the adjustment stops whenever the value creation rate at the end of the pipeline ceases to increase.

3.2 Gradual Monotonic Increase

This heuristics adjusts the number of starts by increasing their number by a constant value ("delta adjustment") every period. The number of starts has a low initial value (20 projects per year), and is gradually increased. Once again, the performance of the previous period, in terms of value creation rate at the end of the pipeline, is compared to the present performance, and the adjustment stops whenever the value creation rate at the end of the pipeline ceases to increase.

3.3 Random Normal choices

Even though managers may make an intuitive judgment and educated guess while deciding on a project introduction policy, uncertainty plays an important role in PPM. Once a benchmark is chosen, it is expected that the uncertainty inherent to the innovation process will affect their decisions. Even though Novartis introduces 40 new projects every year *on average*, this number is not kept always constant (Reyck et al. 2004). It is therefore interesting to determine how different levels of uncertainty in the rate of "starts" affect performance. For such a purpose, a random normal distribution was applied to the decision on project introduction, and different levels of variation were added to the optimal value.

3.4 Target-based Heuristics

The most obvious choice for a decision rule that searches for the optimal choice of starts is one that takes into account the effects of previous choices, changing the direction of the adjustment (increase/decrease) if the previous adjustment resulted in worse performance. Such an optimization effort would be based on a target for performance (value creation), and the search would stop once the performance is near the target level. This heuristics is based on a) a rule to determine the *direction* of the next adjustment, taking into account the effect of a previous decision on present performance; and b) The magnitude of the periodic adjustments ("delta adjustment"). The direction of adjustment is calculated by comparing the present value creation rate with the previous one. If the difference is positive, then the direction of the adjustment should be kept the same. Otherwise, the direction is changed.
‘

Starts: IF THEN ELSE(Value Approval Rate>= target VALUE Outflow BASE CASE, previousstarts , previousstarts + (Direction of Adjustment * Delta Starts))

Direction of Adjustment: IF THEN ELSE(Value Approval Rate >Previous Performance , 1*Previous direction, -1*Previous direction)

4. Analysis of Results

The model used on this study was calibrated to the Novartis innovation pipeline (Reyck et al. 2004). The Novartis pipeline has four stages, but the first stage (basic research) was excluded and only three stages were considered. For the purposes of the study, all variables except the project introduction rate (Starts) were kept constant at "base case" values, taken directly from the data. Base case values are optimal decisions, meaning that value creation at the end of the pipeline is maximum and teams operate at nominal capacity (100%), the point where value added to tasks is maximum.

Initial values for the stocks were obtained by running simulations, so that the steady-state (equilibrium) values were determined. For example, when studying the "random normal" heuristics, one of the initial values for starts is 40. A simulation was run, using a constant value of starts at 40. The model reaches equilibrium after a certain period, and these values are used as the initial values for the stocks in the study.
The calibrated model has an optimal level of starts set at 40 projects per year, which is the "base case" value for the variable. If more of fewer projects are initiated, total value created decreases.

4.1 Gradual Monotonic Decrease

While gradually reducing the number of starts, managers have to decide on the size of periodic adjustment; if the size is too large, the number of starts may "jump" to a value beyond the optimal level of the concave curve, since the adjustment is large and stops whenever the value creation rate ceases to improve. The delays in the pipeline also contribute to distance the final choice of starts from the optimal value. This happens because it takes some time between making a change in starts and noticing the results of such a change. But there is a trade-off in this decision process; even though a small periodic adjustment is more precise, it also *takes longer* to find the "right" value for the decision variable. This indicates that a "mixed" policy could result in a better performance; managers could initiate with larger adjustments and reduce them as performance goes near a benchmark.

Figures 4.1 and 4.2 below show the adjustment process and the impact of the heuristic on performance, for different values of "delta adjustment".

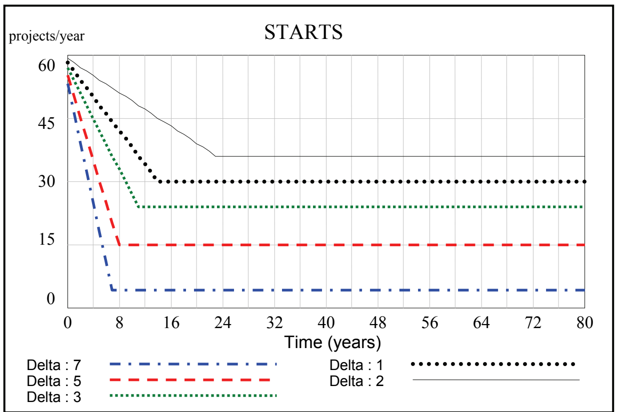

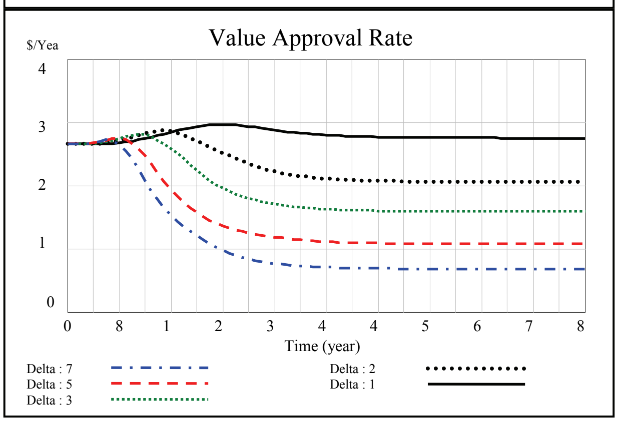

4.2 Gradual Monotonic Increase

A gradual monotonic increase of starts exhibits similar behavior to a gradual decrease, as shown in figures 4.3 and 4.4. Smaller adjustments are more precise in reaching a value near the optimal, of 40 projects per year, but they slow down the search for the right policy.

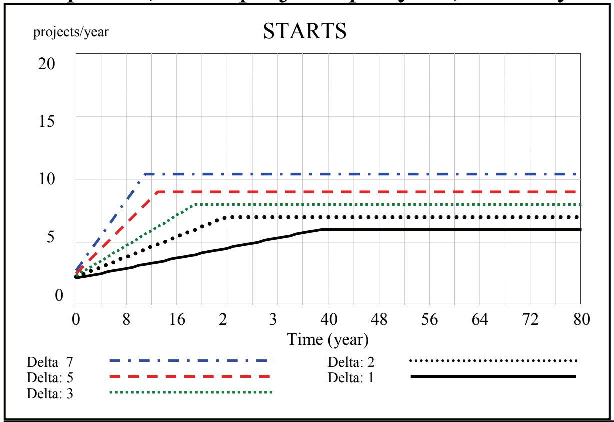

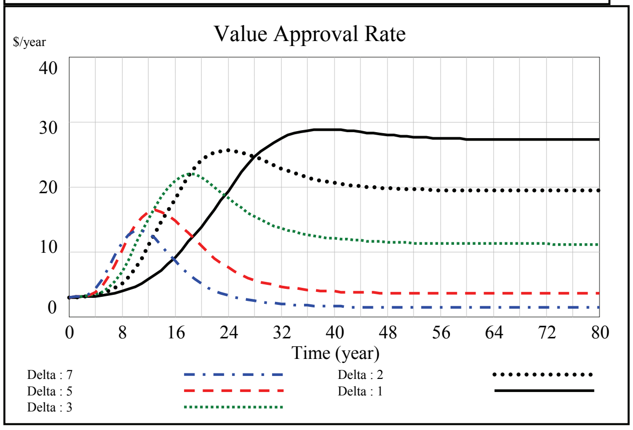

4.3 Random Normal choices

In order to determine the effect of uncertainty in project introduction policies, we compare a fixed, stable introduction of projects with a normal random one. In both conditions, the average project introduction rate is set at the optimal level. The graphs below show how performance is affected by different values of the standard deviation (σ) of the normal distribution. As the standard deviation of the random normal distribution increases, performance decreases. This phenomenon is shown in the graphs below.

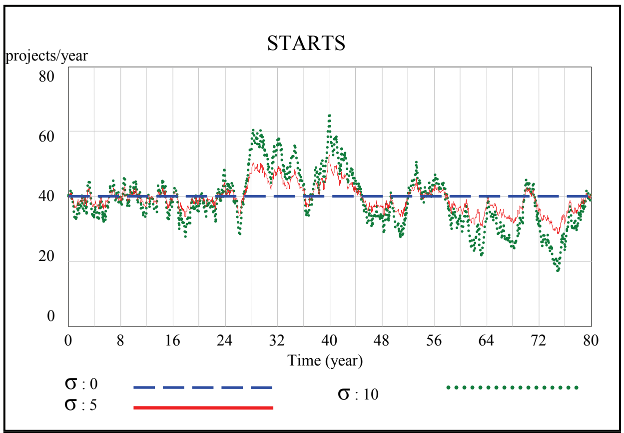

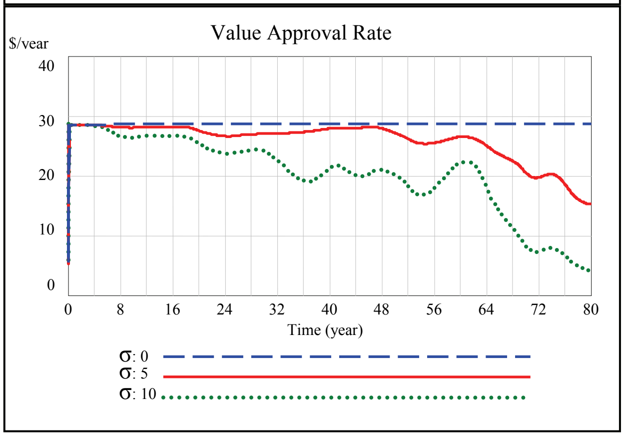

4.4 Target-based Heuristics

The long delays in the NPD process make an ideal target based search quite complex; there isn't a simple rule of thumb to tackle the problem efficiently. The average lead time of a pharmaceutical project in the Novartis innovation chain is approximately 11 years. Because of such large delay, it becomes extremely difficult to determine the right direction of adjustments. A present decision will have an effect on final performance after a long period, and the combined effect of policies from different periods is unknown. A heuristics that takes into account the change in performance and the direction of the adjustment made in the previous period will have very limited efficiency and will create oscillations in the choice of stars. But this phenomenon is not unexpected. The presence of oscillation caused by delays has been well documented in other processes modeled as ageing chains, such as the Beer Game (Sterman, 1995, Goodwin and Franklin, 1994, Croson and Donohue, 2005, Steckel, Gupta and Banerji, 2004).

The results of such heuristic are shown below. The target value creation rate is set at the maximum possible value ($28.0662 per month), and the search process is set to stop whenever value creation is between the ± 10% interval relative to the target. However, the search process does not converge and the decision on starts exhibits oscillation and amplification of variance.

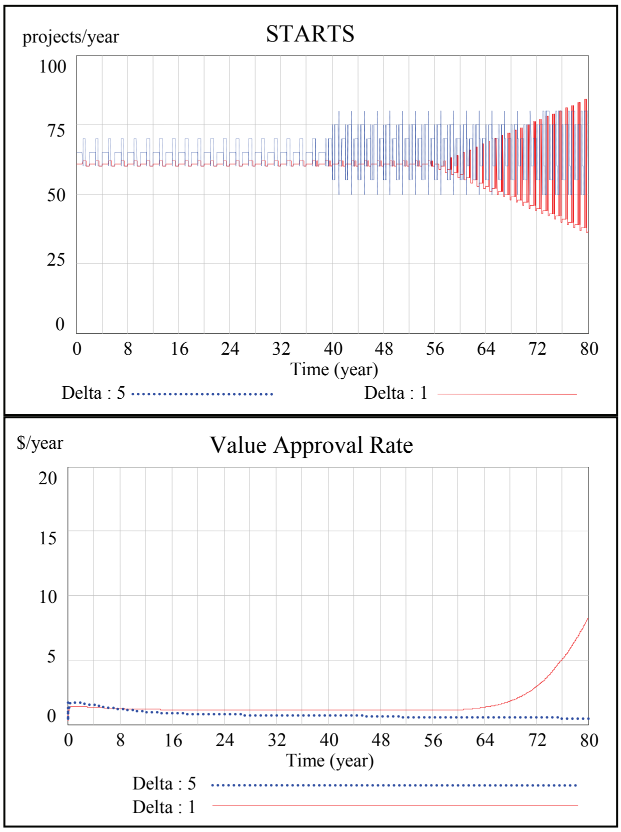

5. Conclusions and Implications

In the pharmaceutical industry, value is being destroyed through longer product development times. It usually takes (on average) 12-13 years to bring a new product to market as opposed to around 8 years a decade or so ago (Paich et al. 2004, Cook 2006). Given that patent lives are (normally) fixed at 20 years, the double hit of increasing "time to market" is evident – higher R&D costs and less time at market before generic competitors are able to enter the market. Pharmaceutical companies have been dealing with this problem by shifting away from internal R&D towards partnerships, licensing deals and acquisitions of more innovative biotechnology companies. Major deals are being made between "old" pharmaceutical companies that have empty R&D pipelines but possess the infrastructure to market new drugs, and "new biotech" companies having technology but no infrastructure (Wilson, 2010). A key problem with such policies is to determine the right amount of projects to acquire or create internally. A too small number of projects may compromise the future sales revenue of the company, since there won't be enough products in the company's portfolio. An excessively high number of projects can reduce efficiency and create bottlenecks in the process, delaying the release of new products and reducing their net patent life. This condition would also compromise future revenues.

Although our analysis was focused on a narrow set of management heuristics, it does have some potentially interesting implications for Product Pipeline structure and project introduction strategies. It could help managers create strategies for project introduction through external acquisition and/or in-house development, and avoid the problem of coping with an excessively large number of projects. Specifically, it was found that a simpler heuristic of gradual decrease or increase of project introduction rates from a high/low work intensity condition can yield better results than a target-based heuristics in which the direction of the adjustment can be changed. Having a target for value creation adds considerable complexity to the problem, since the delays in the process make it impossible to decide on the right direction of the adjustments. This indicates that in the absence of a more sophisticated model to determine optimal policies, managers should adjust project introduction rates gradually.

The size of periodic adjustments is the key driver of precision in such a decision process; a smaller adjustment is more precise, but takes longer to converge. A larger adjustment converges faster, but with less precision. It was also found that consistent, stable project introduction policies yield better results than policies that vary over the optimal introduction rate.

The limits of the present paper suggest several lines of future research. First, in terms of future simulation work, it would helpful to explore a broader range of R&D contexts apart from pharmaceutical pipelines. Second, Even though simple rules can be effective in finding a near optimal policy for starts (under certain conditions), the PPM problem is a complex one. In a real company, updates would have to be made periodically to the "optimal" policies, because of changes in the pipeline configuration. It was assumed that all other decision variables and parameters in the model were kept constant; however, other parameters could change over time and add considerable complexity to the task, since there are many significant interactions between the key decision variables (Figueiredo and Loiola, 2012).

While simulation is useful for exploring specific effects in a controlled manner, some of the richness of an empirical setting was lost. For instance, it would be interesting to capture managerial behavior by means of an experimental, game-based study. We leave that for a follow-on study. We hope that this paper has highlighted some fruitful avenues for further empirical validation and exploration.

6. Bibliography

Ash, R. (2009): An Examination of Engineering Personnel Assignment Policies in the Multi-Project, Capacity Constrained Situation. *Engineering Management Journal,* Vol. 21, No. 4.

Croson, R., and Donohue, K. (2005). Behavioral Causes of The Bullwhip Effect and The Observed Value of Inventory Information. *Management Science*, 00(0), pp. 1–14

Figueiredo,P. Joglekar, N. (2007): Dynamics of Project Screening in a Product Development Pipeline. In: 25th International Conference of The System Dynamics Society, Boston. http://www.systemdynamics.org/conferences/2007/proceed/papers/JOGLE201.pdf

Figueiredo, P. and Loiola, E. (2012). Enhancing NPD Portfolio performance by Shaping the Development Funnel. Journal of Technology Management & Innovation, December. Vol. 7, S. 4, Pgs. 20-35. http://jotmi.org/index.php/GT/article/view/art280

Gino, F. and Pisano, G. (2006). Behavioral Operations. Working paper, Harvard Business School. http://hbswk.hbs.edu/item/5401.html

Gino, F., Pisano, G. (2005). Do Managers' Heuristics Affect R&D Performance Volatility? A Simulation Informed by the Pharmaceutical Industry. Harvard Business School Working Paper. http://www.hbs.edu/research/pdf/05-015.pdf.

Goodwin, J.S., Franklin, S.G. (1994): The Beer Distribution Game: Using Simulation to Teach Systems Thinking. *The Journal of Management development*, vol. 13, n.8

Loch C. H. and Kavadias S. (2002). "Dynamic Portfolio Selection of NPD Programs using Marginal Returns," *Management Science*, Vol. 48 (10), pp.1227-1241

Kavadias S. and Loch C. H. (2004). *Project Selection Under Uncertainty: Dynamically Allocating Resources to Maximize Value*, Hillier International Series in Operations Research and Management Science, Springer-Verlag, Boston, MA.

Steckel, J.H., Gupta, S., Banerji, A. (2004). Supply Chain Decision Making: Will Shorter Cycle Times and Shared Points-of-Sale Information necessarily help? *Management Science*, Vol.50, n.4

Sterman, J. D. (1995) The Beer Distribution Game. In Heineke, J. N. & Meile, L. C. (Eds.) *Games and Exercises for Operations Management.* Englewood Cliffs, NJ, Prentice Hall.

Wheelwright, S. and Clark K, (1992). *Revolutionizing Product Development: Quantum Leaps in Speed, Efficiency, and Quality*. New York, NY: The Free Press.