R&D Portfolio Strategy and Performance: A Behavioral Model

Francesca Gino
Morgan Hall T86,Harvard Business School, Boston, MA 02163

Gary Pisano
Morgan Hall T89,Harvard Business School, Boston, MA 02163

Draft: 10/02/2004

Abstract

This paper explores the underlying causes of volatility in R&D performance over time at the firm level. R&D performance volatility has not been deeply examined in the innovation literature despite the fact that it plays a critical role in industries such as pharmaceuticals or the movie industry, where firms often undergo "hot" and "cold" streaks in R&D output. In this paper, we use a simulation model to explore such phenomenon, building on insights from behavioral theories of the firm: *we argue that the swings in performance, while rooted in uncertainty, are exacerbated by the behavioral influences in how decision makers deal with risk and uncertainty in R&D.* In particular, we propose and formally argue that a firm can affect both R&D output and R&D volatility by modifying its behavior toward risk or by changing the policy it uses to manage its projects portfolio. We simulate a multi-stage R&D process characterized by uncertainty in lead times and projects success. We explore the effect of scale, resource allocation strategies, risk preferences and behavior towards available information on R&D output and volatility. Our results highlight an important tradeoff firms face when managing their R&D portfolio, i.e., the tradeoff between output and volatility. We find that given a fixed budget, priority rules that assure higher R&D outputs come with a "burden": higher volatility. Analogously, we show that the same tradeoff between output and volatility exists when the changing policy regards firms' risk preferences instead of prioritization rules.

I. Introduction

R&D performance plays a critical role in the financial performance of firms across a wide range of industries. The past 20 years have produced a plethora of studies exploring the underlying organizational variables influencing such dimensions of R&D performance as lead times (time to market), productivity, and commercial success of new products (Clark and Fujimoto, 1991; Pisano, 1996; Iansiti, 1997; Fleming, 2001; Ulrich and Eppinger, 2003; Henderson and Cockburn, 1996; Thomke, 2003). Such studies have yielded many important insights of use to practitioners about appropriate strategies and structures for managing individual R&D projects and for managing portfolios of projects. The past several years have also witnessed the development of more sophisticated analytical tools, such as real option valuation, to aid investment decisions in highly risky R&D projects (e.g. Dixit and Pindyck, 1994).

Surprisingly, far less attention has been paid to one of the most vexing issues facing practitioners of R&D in many contexts: *the volatility of R&D output over time* at the firm level.[1] Volatility refers to apparent "streakiness" of R&D performance over time or the tendency for good performance to be followed by poor performance, and vice versa.[2]

[1] For an exception, see Fleming (2001), who focuses on the variance of inventive search. Basically, Fleming argues that one must consider the variance of inventive search when thinking about breakthrough. Inventions that combine new technologies and cross technological boundaries are worse on average, but their variance is higher, and this increases the possibility of a breakthrough.

[2] We should note that volatility is different from the issue of once successful firms failing in the wake of major shifts in technology or markets (Tushman and Anderson, 1986; Henderson and Clark, 1990; Christensen, 1997). Such studies focus on the failures of once successful firms to sustain their advantage in the marketplace. Generally, such failures can be explained by some of organizational inertia in the wake of major shifts in either technology or markets, and the focus is clearly on unilateral changes in performance (from success to failure). Volatility runs in both directions, i.e., from success to failure and back again. Later in this paper, we present examples coming from the movie industry and we show data from the pharmaceutical context to better illustrate the phenomenon.

There are two explanations typically provided for volatile R&D performance. The first is uncertainty in the R&D process itself. Given the technical or commercial uncertainty surrounding most types of R&D projects, it should not be surprising that R&D output is uneven over time. One might expect short-term fluctuations in performance, in much the same way a star baseball player can hit well one day and poorly the next. Short-term fluctuations are random noise. However, uncertainty by itself is not enough to explain volatility. With a large enough portfolio of R&D projects, a firm could presumably "smooth" its R&D output over time. While the prospects of any given project might remain uncertain, we should presumably see performance at the portfolio level even out. Small period-to-period variance can be explained by the random component of R&D, but not large variances that persist over multiple periods.

The second usual explanation for changes in R&D performance over time is change in firm capability or management. The popular business press is full of stories of how specific managers have turned around ailing firms (and vice versa). One wonders, however, whether there is an attribution problem at work in these accounts. While there is an overwhelming amount of evidence that "management matters" in R&D performance (see the studies above), one must be careful to attribute the turnaround of an individual firm to a specific manager. There may be many instances where the "heroic" manager is simply the one lucky enough to be in charge when the firm's fortune improved and the "goat" is the one unlucky enough to be in charge when the firm begins in regression to the mean. In addition, the capability argument, while interesting, does not fully explain why firms swing back and forth from "hot" to "cold". Theories of firm capability (e.g. Nelson and Winter, 1982; Teece, Pisano, and Shuen, 1997) would suggest that firm

capabilities evolve slowly and are not subject to radical short-term swings. It certainly may well be the case that a change in management or some other change in organizational capabilities can improve (or destroy) R&D performance, but it is unlikely that such an impact would happen quickly. In addition, such a change in a positive (or negative) direction can not explain why the firm's R&D performance might decline a relatively short-time later. Presumably, the impact of such capability changes should be relatively enduring, unless a major shift in technology or markets made those capabilities obsolete.

In this paper we explore a third alternative based on behavioral theories of the firm: *we argue that the swings in performance, while rooted in uncertainty, are exacerbated by the behavioral influences in how decision makers deal with risk and uncertainty in R&D.* That is, we assume that there is a certain level of "native" uncertainty in R&D processes that generate some degree of volatility. However, the effect of such uncertainty is magnified by managerial influences. In particular, we propose and formally argue that a firm *can* affect both R&D output and R&D volatility by modifying its behavior toward risk or by changing the policy it uses to manage its R&D portfolio. Thus, the R&D performance differences we observe in industries such as motion pictures or pharmaceuticals do not emerge (only) by chance.

We develop a simulation model of multi-stage R&D process where firms can make decisions about investments in projects at each stage subject to total budget constraints. There is uncertainty in how long each stage takes, as well as in the prospects for "success" at each stage. The objective of the firm is to maximize its R&D output over time. With this objective "in mind", at each point in time firms have to make decisions

about whether to invest additional resources to move a project forward and how much to invest on projects that are at different stages of development. The main policies firms can use in making such decisions regard (a) rules to prioritize projects, i.e., rules to determine which projects get funded first; and (b) rules firms use to deal with risk and uncertainty, e.g., uncertainty around the nature of a project (i.e., whether it will be a success or a failure). We focus on these two types of policies in our simulations.

The exploration of the impact of several rules on R&D performance is aimed at a deep understanding of *what causes volatility*. The present work addresses this issue. In particular, by varying key parameters, we are able to interrogate the model to shed light on three questions as: (i) what is the effect of total portfolio size (scale) on R&D performance (output and volatility)? (ii) what impact do different resource allocation heuristics have? iii) what impact do different preferences toward risk and behaviors toward available information have?

Our findings show that firms' policies and behavior *do* affect R&D volatility. Specifically, our results suggest an important tradeoff firms face when managing their R&D portfolio, i.e., the tradeoff between output and volatility. We find that, given a fixed budget, priority rules that assure higher outputs come with a "burden": higher volatility. Analogously, we show that the same tradeoff exists when the changing policy regards firms' risk preferences instead of prioritization rules. We suggest that these two types of behavior we model can be seen as two sides of the same coin. Indeed, they are both mechanisms for influencing the "balance" of a firm's pipeline (i.e., how the projects are distributed among the phases of the development process), but in different ways.

While resource allocation affects the capacity at each stage of the R&D process, risk preferences influence firms' behavior in screening projects from phase to phase.

II. R&D Performance Volatility: When Success Breeds Failure and Failure Breeds Success

R&D performance has been widely studied in previous research. The bulk of this literature has focused on various dimensions of lead time performance, productivity, and overall "quality" of the output of R&D processes at specific points in time (Clark and Fujimoto, 1991; Pisano, 1996; Iansiti, 1997; Fleming, 2001; Ulrich and Eppinger, 2003; Henderson and Cockburn, 1986; Thomke, 2003). Far less well researched are inter-temporal variations in R&D performance. There is a literature on declines in performance over time. Research by Tushman and Anderson (1986), Henderson and Clark (1990), and Christensen (1997), for example, has explored the underlying reasons why successful firms fail in the wake of major changes in their technological or market environments. This work has highlighted the critical role played by organizational capabilities and resource allocation processes in constraining companies' attempts to adapt.

The current paper explores a different phenomenon: *R&D performance volatility*. We are interested not only in why performance may decline, but also why it might improve. Our motivation for this research is based on empirical evidence that in some contexts performance volatility appears to be the norm. That is, individual firms appear to experience "hot" and "cold" streaks in their R&D performance. High performance is followed by poor performance and poor performance is followed by high performance.

Perhaps nowhere is R&D volatility more pronounced than in the pharmaceutical industry. Figure 1 below shows the average yearly launches of new chemical entities (NCEs) over discrete 5-year intervals for a sample of major pharmaceutical companies over a 20-year period. The variability of R&D performance is quite clear from this graph. For instance, in the 5-year period 1979-1983, Wyeth had the highest yearly average R&D output of the sample. In the next period, its rank dropped to second from the bottom. However, during the latest period (99-03), its ranking jumps back to the top. All of the firms experience significant swings in their "short-term" (5-year interval performance) despite the fact their R&D spending during this period was increasing monotonically. A comprehensive analysis of pharmaceutical company R&D performance by Grabowski and Vernon (1994) reveals a similar pattern: the rank ordering of company innovation performance varies significantly across time.

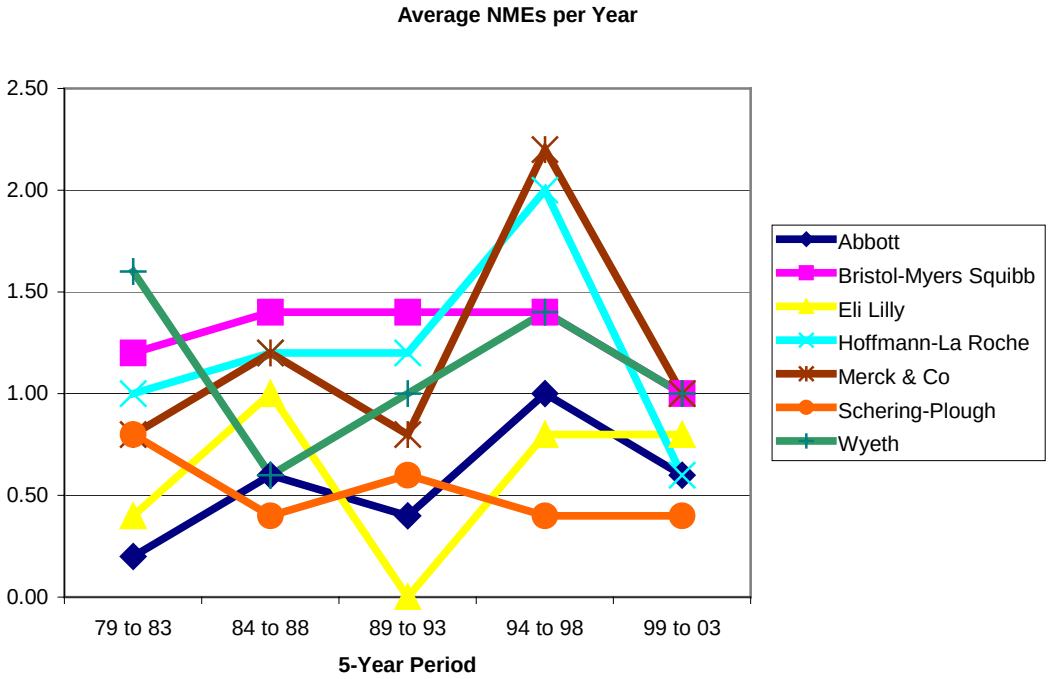

Figure 1. Average yearly launches of new chemical entities.

It is also hard to imagine that such volatility is the result of changes in internal organizational capabilities. While such changes could explain uni-directional improvements (or declines) in performance, it seems highly unlike the internal capabilities were cycling rapidly. This would certain run counter to what is generally believed about the relatively slow rate at which internal organizational capabilities evolve (e.g. Nelson and Winter, 1982).

One might argue that because pharmaceuticals is highly uncertain (e.g. fewer than 1 in 5 candidate drug compounds that are initially tested in people ever reach the market), such volatility should be expected. However, all of the firms in the above example are quite substantial and have a relatively large number of development programs running in parallel. For instance, in 2000, Merck publicly reported 51 internal R&D programs[3] in various stages of development; Eli Lilly had 46; Bristol-Myers had 27 (Paraxel, 2000: 34). Given the proprietary nature of drug company R&D portfolios, it is also quite likely that these publicly reported figures actually *understate* the total. In essence, these firms appear to be doing exactly what they should be doing: they are pursuing highly diversified R&D project portfolios. Yet, despite this level of diversification, we still see a relatively high degree of volatility in output.

Volatility is not specific to pharmaceuticals but it is an issue in other industries, such as medical devices, motion picture and, in general, the entertainment industry (Vogel, 2001). A good example occurred in the market for cardiovascular stents, where the leading companies (J&J, Boston Scientific, Guidant, and Metronic) have each dominated the market at different times and for different generations of stents. Or, in the

[3] These numbers only include "internal" R&D programs, and exclude drug candidates licensed in from external sources. If one includes externally licensed-in compounds, the portfolios are even larger.

movie business, this phenomenon is often in place and can be seen when a studio launches a major blockbuster, followed by a series of flops, and then launches another blockbuster.[4]

Common features in R&D across these contexts are high costs, high uncertainty (both technological and market uncertainty) and, most important, *low predictability of outcome*.[5] Because these characteristics of R&D in such contexts are all interconnected, a combination of the three ensures that even the largest companies must rely on the returns of a small number of very successful projects. For instance, in the video games industry in order of 1000 games per year released 90% of video games lose money and the top 20 games generate 50% of sales.[6]

Before exploring the underlying causes of such volatility, we should say a few words about why volatility matters. There are two reasons volatility matters for practitioners, and an additional reason it is important for research on innovation. On the practical side, high levels of R&D output volatility contribute to cash flow volatility, which can have deleterious consequences for a firm's ability to fund attractive investment opportunities (Froot, Scharfstein, and Stein, 1994). From an operational viewpoint, R&D output volatility can also cause inefficiencies given the fixed cost structures of the organizational infrastructure needed to manufacture and distribute new products. The

[4] The success of future films can not be predicted based on data about past movies "because the movie industry is very much a 'short-term contract' industry with high turnover where the combination of actors, directors, writers, and studios change all the time" (Ainslie, Dreze and Zufryden, 2002: 6). Even when a movie involves the same talented actors, falls in the same genre category, imitates or follows a previously released successful movie, there is still high uncertainty as to whether or not it will be a success (Craig, Greene and Douglas, 2003).

[5] In pharmaceuticals, there is an additional feature that might contribute to volatility, i.e., long time frames. Indeed, in pharmaceuticals R&D is lengthy (an average of 12 years elapse between the synthesis of new active substance and the placement of a medicinal product on the market), costly (895 million dollars is the estimate for new chemical or biological entity R&D cost in 2001), and of a high risk nature.

[6] MacCormack, A. and E. D'Angelo, "Activision: The 'Kelly Slater's Pro Surfer' Project", HBS case study, N9-605-020, 2004.

firm may find that it has significant excess manufacturing and marketing capacity during "cold" streaks where it has few new products to launch, but then lacks capacity in a "hot" year where it has multiple new product launches.

Volatility also has important implications for research on innovation. The general method for studying the factors contributing to R&D performance has been to find associations between various organizational and managerial practices (e.g. team structures, project management systems, problem-solving methodologies) and performance *at specific periods of time.* High performance volatility, however, should make researchers cautious about drawing conclusions based on snap-shots in time. Without an understanding of the underlying causes of volatility, it may be difficult to distinguish between "signal and noise" in R&D performance.

III. R&D Portfolio Management & Performance Volatility

This section develops the conceptual underpinnings of our argument. R&D is inherently uncertain. There are two practices that companies use to deal with this uncertainty. At the level of the individual project, it is extremely common to divide the process into discrete *stages* (for instance, research, concept development, detailed design, prototyping, manufacturing scale-up, etc.). The number of stages in an R&D process varies by context but the logic behind staging is essentially the same: staging provides the firms discrete points in the process to exercise an *option* to terminate further development. Thus, for instance, a firm might decide that after an early "concept development" phase further development is not warranted. R&D can be viewed as a process of acquiring and processing information about the technical feasibility and

commercial attractiveness of alternative designs or technologies.[7] Each stage represents, then, an investment in acquiring information that is then used to determine whether or not to proceed with the project. Real option analytics to uncertain investments (e.g. Dixit and Pindyck, 1994) provides a strong rationale for such a staged approach.

The second practice to deal with uncertainty is portfolio diversification. While individual projects may be highly uncertain, the firm should be able to reduce its overall exposure by investing in a portfolio of R&D projects. The size of a firm's R&D portfolio is generally dictated by the firm's overall size. Empirical studies suggest that total R&D budgets are typically set as a some fixed percentage of revenues (which may vary by firm and across industries). At any point in time, then, a firm can be viewed as holding a portfolio of R&D projects at various stages of development (e.g. some early, some closer to commercialization and market introduction). Given the large development costs for a single project in industries in which volatility is an issue, the strategy is to have a steady stream, or pipeline, of projects at various stages of development, so as to secure the companies' viability in a certain period of time.

Assuming a fixed total R&D budget, the management problem is two fold: first, decide which projects to start, and then, decide which projects to continue and which to terminate at various stages of development, and how much to invest at each phase. Up to this point, we have described the phases as part of the formal R&D process, but there is no reason why we could not extend the issue to consider additional "downstream" (post-commercialization) investments in such things as manufacturing capacity or advertising. Our model can be easily extended to consider such post-launch investment decisions and their ramifications for volatility. In addition, our model can also be extended to consider

[7] For a discussion of the information processing approach to R&D, see Clark and Fujimoto (1991).

decisions about investments in products at different stages of their life cycle (e.g., "next generation" technology (high uncertainty) versus incremental investments in improving the current technology). For example, a software company like Microsoft must decide how much incremental improvement it will make to its current operating system (version 1, version 2, etc.) and when to invest in a new platform. We might think of investments in next generation technologies as "early stage" and investments to incrementally improve current technology as "late stage". The decision to invest in incremental improvements of existing product lines leaves fewer resources available for development of new platforms.

In making these decisions, managers are making a set of tradeoffs between risks, returns, and time horizons for payoffs. Consider the example of the software company that has a successful program on the market. It can continue make incremental investments in improvements (e.g. adding features). However, there are likely to be diminishing returns to such improvements. At some point, the incremental benefit (in terms of additional sales) will not exceed the incremental cost of the improvement and the firm will need to introduce a "next generation" platform. If it waits too long to start work on the new platform, it will lose sales (as its product loses appeal in the market place). On the other hand, if it introduces a new platform too soon, it will cannibalize sales from the existing platform that might still have had many years of profitable life left in it.

In theory, such tradeoffs are optimization problems that can be tackled with a technique such as dynamic programming. In reality, the sheer complexity, ambiguity, and uncertainty of most companies' R&D portfolios make this an essentially impossible

optimization problem to solve. Most companies simply lack either the information or the capability needed to solve the problem optimally. Information about technical feasibility or market potential is often highly subjective or incomplete. In addition, the time and resources required to complete a given stage may also be far from certain. And then, we need to consider that for a large firm, they may need to manage all of these issues for a 50 or more different projects. The decision-theoretic models proposed in the literature 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, it is not surprising that companies utilize heuristics for managing their R&D portfolios rather than trying to optimize. The idea that heuristics rather than optimization drive R&D decision-making was first introduced by Nelson and Winter (1977).[8] While behavioral approaches to R&D have become well accepted, research on the impact of specific heuristics on R&D performance is surprisingly limited. Very little is known, for instance, about how specific resource allocation rules (policies) or heuristics regarding project termination/continuation may impact overall R&D performance.

In the remainder of this paper, we explore how two types of heuristics might shape R&D performance in general, and R&D output volatility in particular. One of these is the heuristic used for prioritizing R&D investments across stages of the development process. The second is the heuristic used for determining which projects move forward and which ones are terminated. This heuristic is a proxy for a firm's risk

[8] At the individual level, much work has been done in psychology. When optimal strategy cannot be computed a priori because of the problem complexity, individuals are likely to invoke simple heuristics and rules (Tversky and Kahneman, 1974).

preference and it deals also with the issue of how firms deal with information in R&D processes. How to prioritize R&D investments across stages of the development process and how to choose which projects progress forward and which ones do not are indeed the main policies firms use in managing their projects portfolio. Moreover, they give us the opportunity to study the effects of firms' behavior on performance.

In particular, our interest lies in exploring the impact of such heuristics on R&D performance in general, and on the volatility of R&D output in particular. To do this, we first constructed a model of the R&D process. We then used computer simulations to explore the impact of changes in different heuristics and firms' behavior on various dimensions of R&D performance. Simulations, indeed, allow estimation of otherwise intractable models by portraying the dynamic behavior of systems over time (Pritsker, 1986).

Before presenting the simulation methodology, we first describe our model of the R&D process. We have intended this model to be highly general and adaptable to a wide range of particular circumstances. As discussed later, in the simulations we parameterize our model using data from the pharmaceutical industry. However, the model is robust enough to be used to explore R&D performance in an extremely wide range of contexts.

IV. A model for R&D portfolio management

Companies face two relevant problems regarding their R&D processes. First, they need to select a portfolio of projects. Second, they need to manage it, i.e., decide which projects to continue investing in, which ones to terminate.

The focus of our work is this second problem, namely, *portfolio management*. That is, even after a set of projects (a portfolio) is selected, the company must manage the allocation of resources to projects on an ongoing basis. Quite surprisingly, the problem of *managing* R&D portfolios has received much less attention than portfolio selection.[9] A common assumption in the studies addressing this issue is that a given set of projects is initially specified (see, for instance, the study by Hopp, 1987; Banerjee and Hopp, 2001; Subramanian et al., 2000). Our model looks at the problem of portfolio management but, different from previous studies, it allows the set of initiated projects to vary over time.

Before describing the formulation of our model in the next section, a few comments are in order. We started the paper by illustrating a phenomenon we have observed in different industries and which is not explored deeply in the literature: volatility in R&D performance. We have then asked the following question: what factors, if any, exacerbate or reduce such volatility? In asking such question, our interest lies in factors upon which firms have control (at least to a certain extent) and their interaction with exogenous characteristics of the R&D process. That is, given a certain level of "native" uncertainty of the R&D process, our attention falls on endogenous variables the firm can decide upon. We thus focus on firms' behavior. Indeed, a firm is asked to make choices during the development process. The main choices it can make to manage its portfolio regard how projects are advanced, that is which projects get funded

[9] A widespread literature exists on portfolio selection and new product development resource allocation. The problem has been studied through a qualitative approach (see, for instance, Liberatore and Titus, 1983; or Cooper et al, 1998) and analytical models (for a good overview see, for instance, Loch et al., 2001). What becomes clear after reviewing the different proposed models is that the problem of portfolio selection is both complex and multidimensional and, as a result, it cannot be reduced to a single model (Banerjee and Hopp, 2001).

first and what thresholds are set up to terminate projects at different stages of the R&D process.

The model incorporates two dimensions of uncertainty for projects. First, it includes "technical uncertainty", in the sense that some projects are discovered to be "technically" infeasible, and terminated. Technical uncertainty is captured in our model by attrition rates. Second, the model incorporates uncertainty with respect to development lead times.

Volatility (measured as variance of R&D output) results from the interaction of these two sources of uncertainty. Consider a simple example where the R&D process is composed of only one phase. If uncertainty is not in place, then projects progress through the phase based on their expected lead times and expected probability of success. In this case, longer or shorter lead times and higher or lower probabilities of success would affect only mean output over time and its "frequency"[10], but not volatility.[11] For instance, a higher mean R&D output would result from higher expected probabilities of success or from shorter lead times. Thus, what affects R&D output volatility is uncertainty around one of the two factors or around both of them. Volatility is an issue since it implies that output is *not predictable* based on previous performance.

A basic principle coming from the operations management literature might help us clarify what affects output and what instead affect volatility. In queuing theory, indeed, what is known as Little's Law states the following: the average output rate over some time interval is equal to the average number of projects (in progress) in the system

[10] The frequency of output can be measured as the ratio between the cumulative output in a certain period of time and the considered time range, i.e, mean time between launch.
[11] The same result would be in place in the case of a multi-stage process, with buffers between two consequent phases and having infinite capacity.

(i.e., WIP), multiplied by the inverse of their average time in the system (i.e., average lead time in the case of R&D processes we are considering). Thus, as this simple relationship shows, output rate (on average) increases with the average number of projects in progress and decreases with lead times. In our model, attrition rates of projects affect the (average) number of projects in progress. The lower the probability of success, the lower the average WIP, and thus the lower the expected output. As mentioned above, it is the uncertainty around attrition rates and lead times that causes volatility.

In the paper, we argue that volatility is exacerbated by firms' behavior. Thus, after reproducing in our model the uncertainty in the R&D process, using computer simulations we conduct controlled experiments: we change policies and study their impact with respect to a benchmark behavior (the one commonly observed in the industry). In particular, the focus in our simulations is on the comparison of different rules and decisions firms might use to manage R&D projects in the pipeline. Controlled experiments require only variables under study to vary (once at a time) and other elements that might confound the results to be constant. Thus, for instance, in our simulations the budget remains constant over time. In the dynamic of the model, projects compete for resources and the firm, at each point in time, has to decide which projects get funded first. The conflict among resources is better captured if we hold the budget constant over time. We expect firms' behavior to matter. That is, we propose that the firm can do something to affect both outcome and volatility by changing its decisions on how to manage the projects in its portfolio.

Are there different heuristics from the ones firms commonly use in managing their portfolios which can assure a better performance, in terms of output and/or volatility? We will address this question later in the paper. Let's first illustrate the details of our model.

IV.1 Problem Formulation

Our model follows closely the one specified by Banerjee and Hopp (2001), with some modifications that we describe in the following paragraphs. The model is designed so as to address the research questions that motivated this paper and allows for "controlled experimentation". That is, it allows us to study the impact of a certain policy with respect to a benchmark initially identified, holding other variables constant. In designing the model, we also tried to make assumptions that either reflect or are not too far away from the reality of firms and that specifically tackle the research questions we initially asked.

A limited budget is initially specified (B) and it remains constant over time.[12] At each point in time, the firm allocates the budget to a portfolio of $M = \{1,...,m\}$ projects. The set of candidate projects that are available at each point in time is infinite, i.e., whenever the firm wants to initiate a new project, a new project is generated. Each project i must progress through n stages sequentially (in our specification of the model $n = 3$ since we model a 3-stage process).

For each phase j and project i there is an expected lead time $lt_{i,j}$ expressed in terms of months required to complete the phase. Hence, the total lead time of project i is

[12] As mentioned earlier, assuming a constant budget allows us to focus on the decision variables for the firm, i.e. the policies it can use to manage its portfolio and that might vary across firms.

given by $lt_i = \sum_{j=1}^{n} lt_{i,j}$. The time required to complete a phase is a random variable following a Weibull distribution. The Weibull distribution is often used to model "time until failure" (for instance in engineering work) because of the many shapes it attains for various values of its parameters. In our simulation, we chose the parameters so that the probability of obtaining a lead time longer than the average one observes in the pharma industry is higher than the probability of obtaining a shorter lead time, for each project and in each phase. Thus, there is uncertainty around lead times.[13]

Phase j of project i has a cost $c_{i,j}$ expressed in terms of dollars spent to complete the phase. Hence, the total cost of project i is given by $c_i = \sum_{j=1}^{n} c_{i,j}$. At each point in time (in the simulation, time is expressed in terms of months), the firm sustains a cost equal to the average monthly cost for that phase. Since the phase lasts for a number of periods equal to the lead time, the overall cost the firm sustains to work on a project in a certain phase is the product between the average monthly cost and the lead time. Thus, over each phase, there is also uncertainty around costs.

There is also an expected attrition rate for each phase j of project i, $(1 - p_{i,j})$, which represents the percentage of projects that are expected not to progress from one phase to the next. In our basic model, the attrition rate is independent of all other parameters used in the model. Whether a project succeeds or fails becomes known only when the project completes a certain phase. When a project completes a stage, it can be

[13] Different from the model developed by Banerjee and Hopp (2001) the lead time is not a function of the amount of resources allocated to a certain phase. Indeed, we wanted to abstract from such a further complication in our model. Moreover, our assumption is closer to the reality of pharmaceuticals, the industry we use as a benchmark.

terminated or moved forward to the next stage. However, it can only move into the next stage if there are resources available to start and complete the phase. If resources are not available, the project waits in a queue. Once resources become available to the project, it can then enter the next phase. A project is successfully completed when it makes it through all three phases. This implies that the probability of project success is given by $p_i = \prod_{j=1}^{n} p_{i,j}$.

Upon completion, each project i yields an expected revenue of α_i. To simplify the problem, we assume that each project has a constant return (i.e., projects are characterized by the same distribution of returns). This assumption is also justified by the fact that our interest lies in the impact of firms' behavior on performance, measured as number of projects that successfully complete the 3-stage R&D process and not as monetary return on investment.

The objective is to find a policy that lets the firm manage its portfolio so as to maximize the R&D output. Because we assume expected revenue does not vary across successful projects, this is equivalent to saying that the firm's objective is to maximize the expected net present value. For instance, if the focus is on resource allocation, the firm's objective is to find a policy that allocates the available budget to projects so as to maximize the expected net present value.

Let's consider the case in which the number of candidate projects is finite. In this case the formulation of our problem follows closely the one specified by Banerjee and Hopp (2001), except for the cost variable. Using their framework, let $t_i(\gamma)$ be a random variable indicating the time when project i completes all phases successfully under

policy γ. Then the net present value for the policy can be expressed as

$$\pi = \sum_{i=1}^{m} p_i(\alpha_i - c_i)E\left[e^{-\rho t_i(\gamma)}\right] \text{ subject to } B \geq \sum_{i=1}^{m} c_i \ .$$

In the equation, the parameter ρ is used as discount rate and the expectation is taken over all possible combinations of phase completion events. The formulation of the problem becomes more complicated if we allow the possibility of an infinite number of projects (i.e., whenever budget is available, the firm can invest in new projects).

A dynamic programming formulation for the general portfolio management model specified above could be developed. However, since the resulting problem would be too difficult to solve or analytically intractable even in the case in which a set of candidate projects is initially specified (the process we model is indeed non-Markovian[14]), we use computer simulations to explore the impact of different policies and firm's behaviors on R&D output and volatility.

IV.2 Model Specification

We model a 3-stage R&D process (see Figure 2). Such a process characterizes most of the modern companies if we think of phase I as research, of phase II as development, and of phase III as commercialization. In the rest of the paper, however, we refer specifically to pharmaceuticals as a benchmark. We chose pharmaceuticals because there is ample published data available on lead times, development costs, and attrition rates for well-defined stages of the R&D process (Phase I, Phase II, and Phase III

[14] A stochastic process (i.e., random function) has the Markov property if the future depends only on the present, not on the past, i.e., if the probability distribution of future states of the process depends only upon the current state, and conditionally independent of the past states (the path of the process) given the present state (for details see, for instance, Simon and Blume, 1994). A process with the Markov property is usually called a Markov process, and may be described as Markovian. A system without such a property instead is called non-Markovian.

clinical trials). In addition, there is also ample and highly reliable data available on R&D outputs (drugs approved) that can be used as a comparison to the outputs predicted by our simulation model. Moreover, we were familiar with existing practice in the industry, and this enabled us to use the most common behavior observed as a benchmark for our simulations.

The basic parameters of our R&D process are specified in Table 1. Note, by altering the parameters, one can essentially model a broad array of R&D processes. For instance, if interest lied in modeling a two stage process, it would only be necessary to set the lead time, variance, and attrition rate for a particular stage to 0.

Basic Parameters	
T	Total number of time periods (t=1,2,...,T); unit of time = month
M	Number of projects in progress that the firm can have active at time t=1
Total_RandD_budget	Equal to B = x_1*cost_R + x_2*cost_ED + x_3*cost_LD at any one point in time, with x_j = number of projects in progress in phase j.
x_1, x_2, x_3	Number of selected projects that are in phase 1, 2 and 3 respectively
cost_R, cost_ED, cost_LD	Cost in the Research stage (R), in the Early Development stage (ED), and in the Late Development stage (LD) respectively
LT_R, LT_ED, LT_LT	Lead Time in stage R, Lead Time in stage ED, and Lead Time in stage LD
prob_R, prob_ED, prob_LD	Attrition rate in stage R, ED and LD respectively

Table 1. Basic parameters.

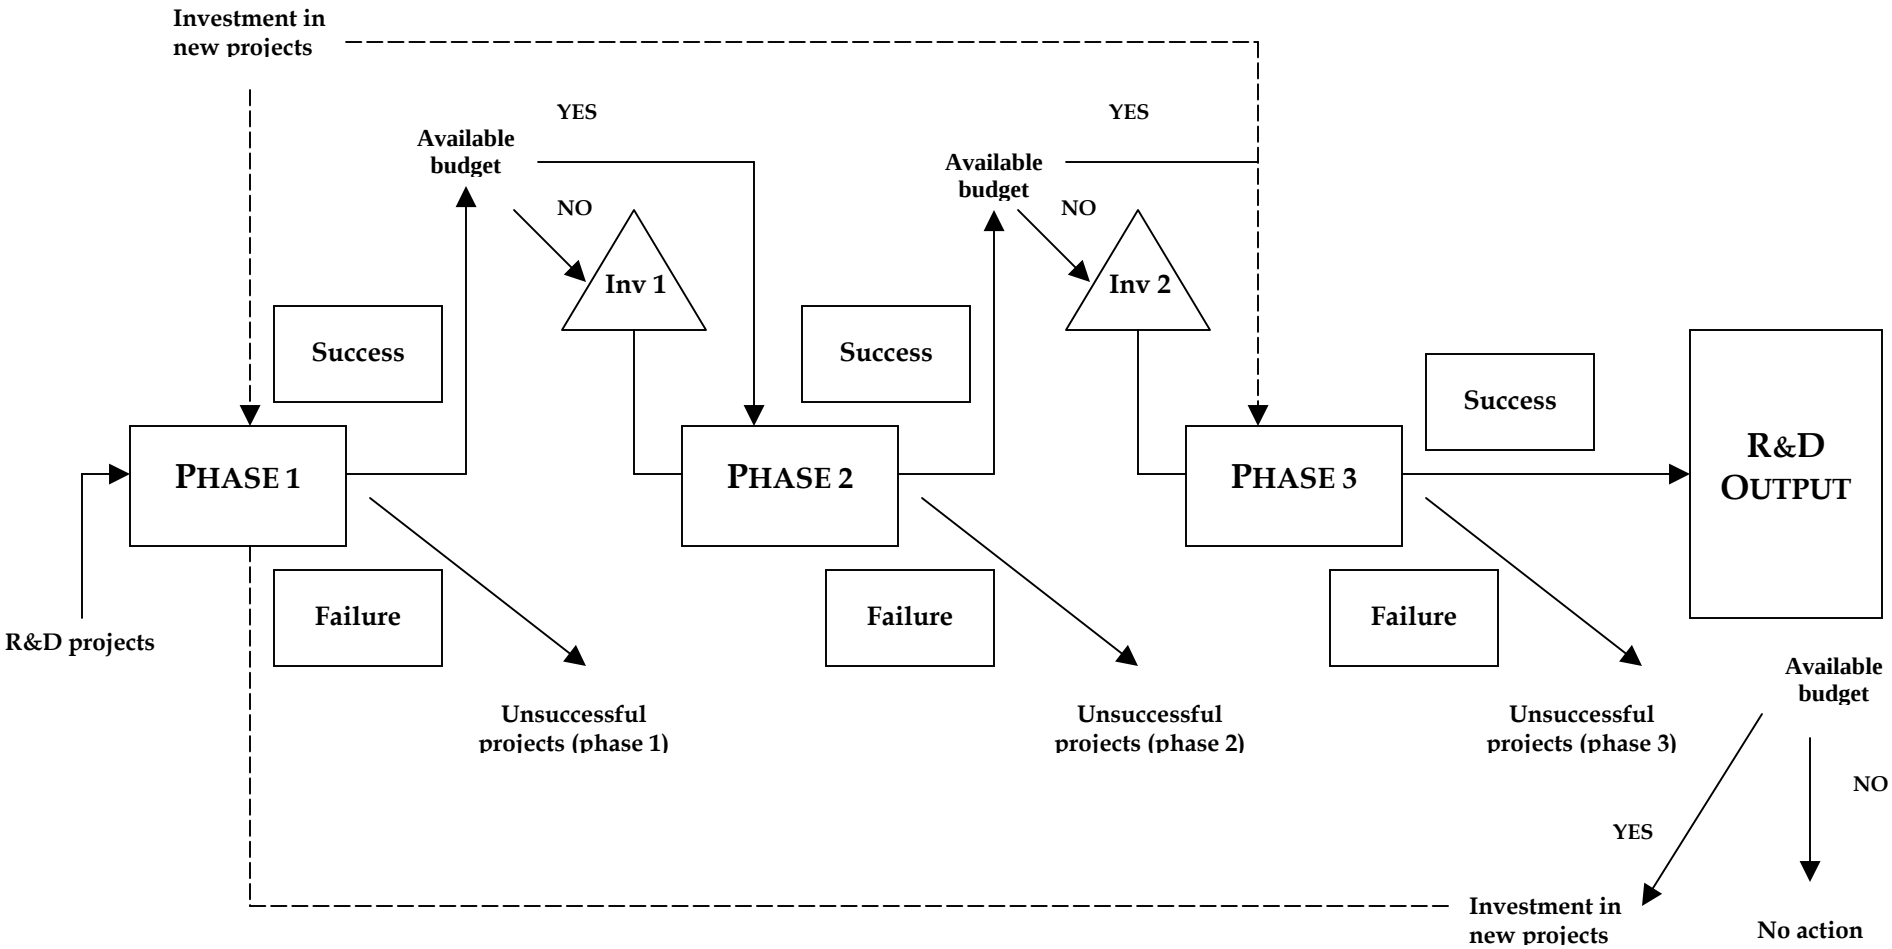

Figure 2. A 3-stage R&D process.

Settings with different levels of risk and uncertainty can be modeled by altering the attrition parameters and variance parameters.[15] And, settings with different cost and lead time structures can be explored by varying those parameters.

The data from the pharmaceutical industry we used to set the parameters in the simulation are presented in Table 2.

	Average Values (in each phase)		
	Phase I	**Phase II**	**Phase III**
Attrition Rate	60 %	50 %	15 %
Lead Time per Phase	21.6 months [st dev = 6.6]	25.7 months [st dev = 11]	30.5 months [st dev = 9.9]
Cost per Phase	MM $ 15.2	MM $ 23.5	MM $ 86.3

Table 2. Data used to set the parameters in the simulation.

* *Source*: J.A. Di Masi et al., "The Price of Innovation: New Estimates of Drug Development Costs", *Journal of Health Care Economics*, 22, 2003.

As shown in the table, phases have specific features. In particular, phase III is the most expensive phase, with the longest lead times but, at the same time, with the highest projects' probability of success. Instead, phase I is the cheapest phase, with the shortest lead times and the lowest probability of success. All these features are taken into consideration when a firm has to decide which projects to move forward and which to terminate, or which projects to fund first. We will come back to the characteristics of the phases of the R&D process when analyzing our results. How the basic parameters presented in Table 1 are used in our simulations is described in detail in the next section.

[15] Thus, differences among firms or industries could be represented by differences in the stochastic process used to "draw" projects and move them forward.

IV.3 Methodology

This section describes the simulation model we developed with varying parameters. The model allows one to set the policies and heuristics a firm might use to manage its portfolio and to face market uncertainty (e.g., in allocating resources among phases and projects, in dealing with risk and uncertainty, in using information rendered available during projects development), and compare their impact on R&D performance (measured in terms of R&D output and R&D volatility).

In the basic simulation model, two parameters, T and M, govern the behavior of a modeled firm. The total time T was set to 1920 months (i.e., 160 years) in each simulation and 100 iterations were run in each simulation set (for instance, we ran 100 times the simulation in which M was set to 5 and T to 1920). Results were then averaged over the 100 runs. Indeed, results of a single run may differ substantially from the "true" measures of the system being simulated: thus, several experiments are run and then a point estimate is computed. Thus, since realizations of the process are subject to stochastic variability, repeated simulations using the same initial conditions and parameters are used to estimate the distributions of outcomes.

We set T equal to 1920 in order to allow a "warm-up period" and thus comment our results after the system has reached a steady state. A steady state is usually defined as a condition in which some specified characteristics exhibit only negligible change over an arbitrarily long period. To put it differently, it is the condition the system reaches after all initial transient or fluctuating conditions have damped out. Using long simulation runs and deleting the data corresponding to the warm up period in the analysis of the results are two possible ways to reduce the impact of the warm-up period on the accuracy of the

measured data. We thus used both these methods to assure proper analyses of our simulation results.

The type of simulation we implemented is a *nonterminating steady state simulation*: there is no natural event that specifies the length of each run (think, for instance, about the performance of a hospital emergency room). 160 years seemed to be a reasonable long time to be used as a stopping rule. In pilot simulations we also checked that the use of longer time periods for T did not produce significantly different results.

At the beginning of each simulation, M projects are randomly generated based on the fixed budget. Each project is modeled as a vector, with each element representing a certain feature of the project. In particular, each project is characterized by the following features: a project ID; its cost in each phase (Research, Early Development and Late Development); its lead time in each phase (i.e., the time required for a project to complete that stage); its attrition rate in each phase (i.e., the percentage of projects which do not complete successfully the stage and thus do not move to the next); and an index indicating the current stage a project is in. Both lead times and attrition rates are randomly assigned. As soon as projects are generated, they enter the 3-stage process described above (and represented in Figure 2). During a simulation run, whenever the firm's budget allows for additional investments, new projects are generated.

Projects are independent identical distributed random variables since some of their features are generated randomly. That is, we are implicitly assuming a firm's portfolio is diversified, with each project assuring a fixed return if successful. This is so since the objective of the simulation is to explore how different firm's behaviors and policies used to manage R&D portfolios affect R&D performance *typically*, and not for a

specific project.[16] A relevant assumption is reflected in role of chance in the modeling. The values of some of the features of the projects firms invest in are not known a priori. For instance, a firm does not know a priori whether a certain project will be either a success (what we call later a "winner") or a failure ("loser"), or when it will fail.

The details about how each heuristic or behavior was modeled in the simulation are provided in the next section, in which we also discuss our results.

V. Analysis

Our analysis proceeded in 3 phases. Our first analysis was to explore the relationship between scale (R&D budget size) and R&D performance. We then examined how resource priority rules (holding constant scale) across stages of the development impacted R&D performance. And, finally, holding constant resource priority rules and scale, we examined the impact of risk preferences on R&D performance. The benchmark behavior in each set of simulations (and consequent analyses) is the policy commonly observed in the industry. Moreover, in each simulation run the policy a certain firm uses remains the same in each time period.

V.1. Scale Effects

There is a long literature in economics and the management of innovation research about the impact of scale on R&D performance. The issue is also one of great practical concern. Industry mergers in R&D industries are often justified, at least

[16] A similar argument is used by Rivkin (2000) in his discussion about randomly generated decisions and the role of chance in the simulations he uses.

publicly, on the benefits of scale for R&D. Empirical research on the impact of scale on R&D performance has shown mixed results.[17]

The issue of scale is multi-faceted. Scale might affect performance through its impact on control structures, incentives, organizational processes, and culture. Much of the debate about the impact of scale (economies vs. diseconomies) revolves around these issues. Far less explored is the impact of scale on variability of output over time. That is, *does scale reduce R&D risks?* Our simulation is not designed to explore all facets of scale, but instead focuses on the latter. In so doing, we explicitly assume constant returns to scale in the conversion of R&D inputs (dollars) to R&D outputs (successfully completed products). We make the further assumption that resources are allocated according to the following priority rule: projects in the 3rd (final) stage of development receive the highest priority. They get "first shot" at the resources. Projects in the second phase get the second highest priority. Projects in the first (earliest) phase get third priority. We start with this rule because based on prior field observations; we believe this one most closely approximates current practice. We will provide an explanation of why firms might use this heuristics later in the paper. At that point, we will also explore the impact of alternative resource priority rules.

Figure 3 shows the impact of scale (budget size) on output. The constant returns to scale assumption is clearly evident, as R&D output increases linearly with R&D budget size. Thus, these results suggest that as one adds more capacity, one gets more output. How "frequent" is such output over time? The frequency of mean R&D output is measured by mean time between launch. Our analysis of the effect of scale on mean time

[17] See Henderson and Cockburn (1996) for both a comprehensive analysis of scale effects in pharmaceuticals as well as an excellent overview of the issues.

between launch is contained in Figure 4. As one might expect, mean time between launch decreases exponentially with R&D budget size.

One of the shortcomings of Mean Time Between Launch as a measure of volatility is that it is computed as the ratio between the cumulative R&D output over a certain number of years and the number of years itself. Thus, such measure does not capture how volatile the output is in the considered time range.[18] An alternative and preferable measure of volatility (that avoids the above shortcomings) is the standard deviation of output over time. This is the measure that we use below and in all subsequent analysis of volatility.

Figure 5 shows the relationship between volatility (measured in terms of descaled variance[19]) and scale. This relationship is not straightforward: Increasing the R&D budget size does not necessarily lead to lower R&D volatility over time. Thus, when increasing the budget size, a firm is initiating a higher number of projects and it is diversifying its portfolio even more. In so doing, a higher number of initiated projects at each point in time is likely to translate into a higher number of projects in the pipeline, progressing from a phase to the next one. As a result, as the budget size increases, a higher R&D output becomes more likely over time. However, the effect on volatility is not as clear since a larger portfolio over time does not imply lower uncertainty in the features of the projects in the portfolio itself, such as lead times or attrition rates. Thus, if

[18] For instance, suppose firm A registered the following R&D output in the last 10 years: [0 1 0 1 0 1 0 1 0 1]. Firm B, instead, registered the following sequence: [0 0 0 0 5 0 0 0 0 0]. Overall, both firms had an output equal to 5 projects launched over a 10-year period. As a consequence, the mean time between launch is the same in both cases, i.e. 6 months. However, the volatility of the output – as we define it –is different in the two instances described. Firm B's output is more volatile than firm A's one.

[19] We use descaled variance instead of variance in order to be able to compare the results for different budget sizes.

one increases scale (i.e, the number of projects) there is no effect on lead time uncertainty or attrition rate, and, as a result, there is little effect on volatility.

In essence, this result suggests that there is room for firms to affect R&D performance in terms of both output and volatility by *changing* either the policies they implement or their *behavior*.

V.2. The Effects of Resource Prioritization

Our second analysis was to explore how, holding constant scale, changing the heuristic used for prioritizing R&D investments across stages of the development process affect R&D performance.

There is a long literature on resource allocation processes going back to Bower (1970). In general, this literature emphasizes that heuristics (policies) play a key role in resource allocation decisions. Policies like "we will not invest in projects that do not have an expected return about our weighted average cost of capital" would be an example of a heuristic that plays an important role in capital budgeting decisions. In R&D, there are various decision rules that influence allocation of resources to specific projects. Our field work (largely conducted in the pharmaceutical industry) suggests that one of the issues causing the most tension within companies is the prioritization of resources across stages of the development cycle (e.g. research vs. development vs. commercialization). This appears to be most pronounced when specific stages of the R&D process are the responsibility of distinct sub-groups within the organization ("Research", "Development", "Sales & Marketing"). Given that organizations tend to pursue a larger number of projects than they have the resources to fund (Wheelwright and Clark, 1992),

prioritization across stages can become a source of intra-organizational conflict with each group vying to have their projects funded first.

There are various organizational factors that influence the implicit or explicit resource prioritization rules in place. Organizational culture and history may play a critical role. For instance, a company founded by academic scientists may have a bias toward funding earlier stage research projects. At one biotechnology company we interviewed, one of the senior managers noted that for many years the company had difficult advancing projects further into the development process because the organizational culture (heavily influenced by the academic origins of the company) emphasized early stage research. For companies similar to this one, new projects look always better than ongoing projects, based on the fact they seem promising.[20] In contrast, at the other end of the spectrum, there are companies with very strong sales and marketing organizations that tend to exert a strong pull on resources devoted to advertising or incremental improvements in existing products.

The other factors that can influence resource prioritization have to do with the firm's time horizon. A company in a difficult financial position or facing serious short-term cash issues might choose to focus its resources on products closer to market. In contrast, a firm in a stronger financial position might be able to afford to initiate more new projects that will not be ready for the market until several years out.

The most common resource prioritization rule we have observed in practice is to fund later stage development projects first. That is, projects closer to the market get first dibs on resources, and prioritization declines as projects move further back in the

[20] A similar tendency can sometimes be found in Academia, when researchers show their excitement for new projects and then find it more difficult to take them to an end.

"pipeline". Such a rule appears to make quite a bit of sense in many regards. It enables the firm to capitalize on its prior R&D investments as quickly as possible, and is likely to lead to the fastest payback of investments. In addition, in many contexts, customers are waiting for new products, and delaying those projects can have deleterious effects on market position. However, in multi-period setting, it is not clear what impact such a rule might have on long term R&D performance, because the decisions made today have an impact on future options. For example, a decision to give greater priority to late stage projects (vs. early stage projects) today means that at some point in the future the firm will have fewer late stage options, and this will cause a gap in product launches.

In our simulation model we implement the commonly observed prioritization rule as follows: the firm always spends resources on the latest phase projects first, the second to latest phase projects second, and so on. Moreover, it can only start a new project if it has resources left over that have no already been committed to phase III, II and I projects. We term this rule as *321*.[21] The impact of such a policy for advancing projects on both R&D output and R&D volatility is compared to the following rules: *rule 231*, according to which the firm spends resources on projects waiting in the first buffer (the one between phase I and II), then on projects waiting in the second buffer (the one between phase II and III), and finally, if budget is still available, it invests in new projects; *rule 123*, according to which the firm spends resources on new projects first,[22] then on projects waiting in the first buffer (the one between phase I and II), and finally on projects waiting

[21] Thus, "3" refers to the buffer between phase II and phase III, "2" refers to the buffer between phase I and phase II, while "1" refers to investments in new project (see Figure 2). There is no cost associated with projects waiting in buffers. Moreover, we assume there are no capacity constraints for buffers.

[22] To be precise, a firm can invest in new projects every three months. If this was not the case, the firm would not have enough resources to move projects forward in the process and thus no output would be observed.

in the second buffer (the one between phase II and III); *rule rnd*, according to which the firm allocates its resources randomly at each point in time when budget is available.

How do such different resource prioritization rules impact R&D performance, over long periods of time? Intuitively, how resources are prioritized affects the way projects progress through the phases of the R&D process. That is, resource allocation rules affect where the firm adds capacity in the system, i.e. they determine the firm's capacity to work on a higher number of projects in a certain phase. Different rules thus result in a more or less "unbalanced" pipeline, i.e., the number of projects in each phase or waiting in the buffers between two subsequent phases greatly varies, according to how capacity is allocated. The system is "fed" differently based on the rule used and thus either starving or blocking might be observed as soon as the system reaches a steady state.

For instance, with the rule 321, the distribution of projects in the pipeline is such that the highest number of projects is in phase III or waiting in the buffer between phase II and phase III. The rule gives priority to projects in the latest phase, i.e., projects with the highest probability of success. This means that when projects reach phase III they are likely to succeed. As soon as the buffer between phase II and phase III becomes empty, however, the firm has to wait for quite a long time before other projects reach that stage. That is, after late stage projects get completed, there is a gap in the project pipeline. In giving priority to the latest phase, indeed, the firm is not worried about filling the pipeline in the early stages of the development process (i.e. it does not feed the system while it is working on projects in the latest phase). So, the early stage projects get "starved" for resources while late stage projects get what they need. When the firm starts feeding the

pipeline again, it invests in projects with the lowest probability of success (indeed, phase I is the most uncertain). And this results in high volatility of output.

A different dynamics is at work with rule 123. In this case, the firm feeds the system on a quite regular base over time (as mentioned earlier, indeed, investments in new project can be undertaken every 3 months), since capacity is allocated to new projects and early phase first. Thus, even if projects progress at a lower pace within the pipeline and result in a lower output over time, such rule avoids starving to be in place. Thus, if one compares rule 123 vs. rule 321, one would expect rule 321 to be the one leading to the highest output over time but at a cost of higher volatility.

An important aspect the reader should keep in mind when looking at our results is that the rule a firm uses in a simulation does *not* change over time. For instance, if the rule under study is 321 then it remains the same from period $t = 0$ to period T. This allows the possibility for feedback loops to emerge and to reinforce the impact on R&D performance a certain policy has. The only rule that differs in this respect is the random rule (i.e., rnd rule). With this rule, indeed, at each point in time available funds are allocated randomly to projects waiting in the buffer between phase I and phase II, projects waiting in the buffer between phase II and phase III, and to investments in new projects. So, basically, what happens with rnd rule is that the rule used to allocate resources changes from period to period. Because of the uncertainty inherent in the R&D process, one might expect such random rule to perform well in terms of both output and volatility. Indeed, the randomness in the allocation of available resources at each point in time makes it more likely to happen that firms have a more balanced pipeline, which, in turn, assures a steadier stream of output over time.

Figure 6 shows the effect of the different priority rules on output. Over time, for most portfolios, the industry standard rule (i.e., rule 321) leads to the highest output. Results seem to be consistent across different portfolio sizes. However, the rule 321 also leads to the most volatility (for smaller portfolio), as shown in Figure 7. The explanation of these findings goes back to the above discussion.

The analysis of the impact of resource prioritization rules on R&D performance highlight another interesting finding. The random rule performs quite well, particularly for larger portfolios in terms of R&D output and, at the same time, if compared to other rules (and for most portfolio sizes) it gives lower volatility. Again, a possible explanation for such results has been provided earlier in this section, while discussing the intuitions behind the dynamics in place based on the rule used to prioritize resources.

Taken together, these findings suggest an important tradeoff firms face when managing their R&D portfolio by using the same type of rule to prioritize resources to projects over time, i.e., the tradeoff between output and volatility. Thus, given a fixed budget, priority rules that assure higher outputs come with a "burden": higher volatility.

In the next paragraphs, we will show that the same tradeoff between output and volatility exists when the changing policy regards firms' risk preferences instead of prioritization rules.

V.3. The Effects of Risk Preference

Our third analysis focuses on firms' behavior towards uncertainty and risk. Given that uncertainty is inherent to R&D, it is surprising that there is no more research on the influence of risk preferences on R&D performance. Any serious behavioral theory of

R&D should include risk preferences. In particular, different organizations are likely to have different attitudes and preferences for risk. These may be tempered by similar factors discussed in the previous section such as organizational culture, history, financial conditions, etc.

One of the chief ways that risk preferences are revealed is through the criteria used to decide which projects move forward in development and which ones are terminated. We can think of each stage of the R&D processes as an investment in acquiring information about the technical and/or market feasibility of a project. Thus, for instance, at an initial "concept development" phase of a project, preliminary data might be collected on the potential market attractiveness of the proposed new product. Or, in pharmaceuticals, early experiments of a new compound in mice are designed to provide some information about the likely effect of the drug in humans. It is important to recognize that the experiments conducted during the R&D process produce only raw data. Inferences about what those data mean, and, more importantly, decisions about what to do next, are made by managers in the organization. And here is where risk preferences of various sorts are introduced into the process.[23]

The issue comes down to how much risk is the organization willing to take on the program *given the incompleteness of the information.* Provided the exact same data (distribution of outcomes from an experiment), two different organizations may well

[23] In some cases, the results of an experiment yield unambiguous data. For instance, thinking about pharmaceuticals, a drug that is tested on 100 patients that yielded very beneficial therapeutic effects in 99 patients with no apparent side effects clearly warrants additional development (and conversely, the drug that had no benefit for anyone in the study and produced undesirable side effects is probably one that will be shelved). Depending on the context and the circumstances, the results from experimentation are often not clear-cut. The study of the drug in 100 patients might yield very beneficial effects in 20 patients, moderate effects in another 30, and slight effects in 10; but it might also have had some other undesirable side effects in a significant portion of those patients. The initial market study to determine the commercial attractiveness of a proposed new product is more likely than not to yield a forecast with a very broad standard error.

draw very different conclusions about whether the project is worth considering. Guedj and Scharfstein (2004), for example, have documented significant variation in the rates at which biotechnology firms terminate cancer drug programs at early phases in the process. Some firms appear to have a preference for "weeding out" projects at an earlier phase, while others appear to be more willing to take comparable projects deeper into the process. Interestingly, Guedj and Scharfstein (1994) demonstrate that these differences are related to differences in the firm's financial prospects and alternative development opportunities. Firms with weaker financial prospects and more limited development options appear to be willing to advance their programs into later stages of the development process.

The issue of risk preference comes down to the preferred tradeoff between Type I and Type II errors. Firms which are "tough" at the early stages (i.e. only advancing projects with very strong ex-ante prospects for success) are more likely to kill "winners". In general, they tend to kill projects early in the development process. We label them as *risk averse* firms. In contrast, firms which are willing to take projects with questionable prospects forward are more likely to spend money on projects that are actually losers. Thus, firms that set low bars on evidence for advancing projects forward early in the development process and higher thresholds later are said to be *risk lover*. A firm is said to be *risk neutral* when its behavior in setting thresholds for progressing projects is similar across phases, i.e., the firm uses almost the same threshold in earlier and later phases. Finally, we also model an industry standard behavior by setting the values for threshold equal to the ones commonly observed in pharmaceuticals.[24]

[24] Such values are similar to the ones used to model risk neutrality.

The optimal tradeoff may depend on the cost structure of the development process, the signal-to-noise ratio from experimental results, the returns from "winners", and the overall scale of the R&D budget. We explore these tradeoffs and the impact on R&D performance under various conditions in the simulations presented below.[25]

In particular, we compare the impact of different behavioral rules regarding firms' risk preferences to face market and technological uncertainty. In the basic model (the one used to determine the impact of scale and the effect of priority rules on R&D performance), projects move through the 3-stage process "passively". They have a certain probability of success for each phase and they either succeed or not. But, in reality, firms make *judgments* about which projects move forward based *on incomplete data*.

Formally, the general problem we are modeling in this set of simulations is as follows:[26] Assume a firm chooses a policy to maximize the expected (net present) value of its profit function. Thus, the firm's returns π depend both on its policy γ and the state of the world w regarding whether projects succeed or fail, i.e., $\pi(\gamma, w)$. The firm does not observe w, i.e., it does not know a priori whether the project will be a success or a failure. Indeed, it considers w to be a realization of a random variable W. However, the firm observes the realization of a signal β, which is also a random variable and which conveys information about the true realization of w. In our simulation, we assume that the firm can get an accurate signal *without* a cost.[27]

[25] Also in this set of simulations we hold constant returns from successful projects.
[26] The general model described above is related to prior studies on the value of signals that convey information about parameters unknown to the decision maker (see, for instance, Persico, 2000; Jewitt, 1989; Athey and Levin, 2000; Siggelkow, 2002). Such prior work has built on Blackwell's research (1951, 1953) studying the conditions under which the value of a signal is higher than the value of another one.
[27] Persico (2000) examines analytically the incentive to acquire information, based on the assumption that information can be made more informative at a cost.

Consider a firm choosing a policy γ to maximize the expected (net present) value of its profit function $\pi(\gamma, w)$. Let $\gamma^*(w)$ denote the policy γ that maximizes $\pi(\gamma, w)$ for a given w. Thus, if the project is a winner, then the optimal decision would be to advance the project further, while if it is a loser the optimal decision would be to terminate the project (at the end of the phase). For a given w, as the policy γ the firm uses moves away from the optimal one (e.g., with γ it is more likely that the firm will terminate a project when it is a winner or, similarly, with γ it is more likely that the firm will advance a project when it is a loser), the higher the costs the firm will sustain if the project is a loser or the higher the missed revenue if the project is a winner.

Before discussing our findings, we describe how we modeled risk preference. In our simulation, when a project is generated, it is labeled randomly as either a "winner" or a "loser" (by the simulation). The probability of obtaining a winner in each random draw is equal to 17%, which is the percentage of the projects entering clinical development that eventually result in a marketed product.[28] At each phase, and without knowing if the project is a winner or a loser, the firm must decide whether to advance or kill the project, based on data generated during the phase. Data take on the value of either 0 or 1, to indicate failure and success respectively. If the project was originally a winner, then the firm receives a signal (which is equal to an increment in the probability of getting a winner in the data generated during the phase). Risk preferences are then modeled as *threshold for advancing projects*. Thus, the probability of technical success is used as a measure of risk.

[28] See J.A. Di Masi et al., "The Price of Innovation: New Estimates of Drug Development Costs", Journal of Health Care Economics, 22, 2003.

Table 3 provides the values for threshold that were used in the simulation for each type of firms' risk preferences.

	Values for (ex-ante) threshold used in each phase *(a project progress only if its probability of success is higher than the threshold, the threshold being the attrition rate)*		
	Phase I	**Phase II**	**Phase III**
Risk Lover Firm	20 %	50 %	58.5 %
Risk Neutral Firm	41.69 %	50 %	41.69 %
Risk Averse Firm	65.2 %	50 %	2.25 %
Industry Standard	60 %	50 %	15 %

Table 3. Different types of firm's behavior toward risk.

Thus, for instance, a risk averse firm might advance a project from phase I to phase II only if the probability of success is higher than 65.2%. A risk lover firm, instead, might advance a project from phase I to phase II only if the probability of success is higher than 20%. Finally, a risk neutral firm might advance a project from phase I to phase II only if the probability of success is higher than 41.69%.

How do different behaviors toward risk and uncertainty affect R&D performance, over long periods of time? Intuitively, and similar to what discussed for resource prioritization rules, how firms set up bars based on evidence for moving projects forward has an impact on the way projects move from phase to phase in the multi-stage R&D process we modeled. Thus, the "balance" in the firm's pipeline differs based on risk preferences. For instance, a rule that tends to set high bars for stopping projects moving

from phase I to phase II might result in a front-loaded pipeline. Instead, a rule that tends to set high bars for stopping projects moving from phase II to phase III might result in a back-loaded pipeline. And finally, a rule that tends to use similar thresholds for moving projects from phase I to phase II and from phase II to phase III might result in a evenly-weighted pipeline.

What is affected using different risk preferences in this case is not capacity, but screening. That is, if resource allocation was about where the firm expands capacity, risk preferences is ultimately about where the firm keeps projects in progress and how it screens projects from phase to phase.

For instance, with risk aversion the firm lets through fewer projects and this leads to projects being more "spaced out" in time. The firm is actually taking more initial "draws" because it is killing projects faster and then gets to pick again. What happens however when it does let a project go through? Chances are it goes deep in the development process because it has a lower chance of picking a loser. As that project goes further into development, it begins to consume more resources, and this then starts to constrain the number of initial draws. So, the firm begins to offset some of the advantages of the risk aversion behavior. A risk averse firm tend to kill projects early but it can replenish its initial development phase since budget becomes available. Thus, if one compares a risk averse firm with a risk lover, one might expect a higher output for the first over time but also a higher volatility.

Figure 8 shows that this expectation is correct. Risk aversion does lead to higher output, but the effect is not really significant if compared to pharma industry standard practice. However, risk aversion also leads to higher volatility, as Figure 9 illustrates.

Thus, similar to the results we obtained when changing priority rules, also in the case of risk preferences, firms face a tradeoff between output and volatility. That is, by changing its preferences toward risk and uncertainty, a firm can assure a higher output but at the cost of higher volatility.

Figure 10 shows the effect of different risk preferences on R&D costs per drug. As the figure illustrates, the output of the rules is different but the costs change in a linear fashion. Thus, what is interesting is the complete lack of difference. This might be explained by the fact that at each point in time the available budget is entirely spent on projects at different stages of development. Rarely, available budget remain not invested.

Through tough screening in the early stage of the development process, a firm is more likely to kill winners. Figure 11 and 12 report the number of killed winners and non-killed losers. two actions are costly for a firm: to kill winners (because of the missed future revenues) and not to kill losers (since the firm sustains costs until it actually decides not to progress the project further in the development process). As one would expect, with the risk averse policy, since firms tend to kill projects early, a higher number of winners is likely to be killed (see Figure 11). Implications might be relevant, since there is a cost associated to killing winners. Think for instance about drugs: a firm kills drugs that might have been winners, and the winners in the pharma industry not only can be very important (e.g., in curing diseases) but they also have high returns. The same result is shown in Figure 12, but from the "opposite" perspective, i.e., by looking at non-killed losers instead of killed winners. Again, it is worth noting that continuing developing losers is costly for a firm.

V.3.1. Robustness of the results for the effects of risk preferences

In the simulations described in the previous section, every time a project was generated it was randomly attached a label, i.e., "winner" or "loser". Whenever the project was a winner, the firm received a signal (β). In particular, the signal was modeled as an increment of 0.20 in the probability of getting a winner in the data generated during the phase. We conducted further simulations to check the robustness of the results presented in the previous section. In particular, we used different values for signal (0, 0.2, 0.4, 0.6 and 0.8). Overall, the analyses conducted show robustness in our findings, since results do not change if different values for signal are used.

VI. Conclusion

In this paper we have explored the underlying causes of volatility in R&D performance over time at the firm level.

Our behavioral approach was motivated by the fact that in most contexts, R&D portfolio is an extraordinarily complex endeavor. Managers must not only select projects they think have the best chance for success (*portfolio selection problem*), but they must make a sequence of decisions through multiple development stages, e.g. to continue the investment or kill the project (*portfolio management problem*). All of these decisions must be made amidst high levels of uncertainty about a project's prospect for success (technical and commercial), the times it will require to be developed, and the resources required. Moreover, while textbook examples often focus on choices about individual projects, a large corporation may have to make these decisions simultaneously for a large

number of development projects.[29] Moreover, the decisions are inter-related. Given fixed budgets for R&D, a decision to invest heavily in one project means investing less in others. Managers have at their disposal various analytical techniques for dealing with these problems (e.g. real option valuation). However, given the sheer complexity of the problem to be managed, it is virtually impossible to develop optimized solutions to the problems of R&D portfolio resource allocation and portfolio management. Instead, organizations rely on various heuristics and strategies for allocating resources and making portfolio decisions. Such heuristics might lead to suboptimal outcomes.

Using computer simulations, we have explored the impact of changes in different heuristics on various dimensions of R&D performance. In particular, we have focused on three types of behavior. First, we explored the relationship between scale and R&D performance. Second, we examined how resource priority rules (holding constant scale) across stages of the development impacted R&D performance. In the simulations, the firm is given the possibility to choose how to prioritize R&D investments across stages of the development process and the impact of several priority rules is considered. Finally, holding constant resource priority rules and scale, we examined the impact of risk preferences on R&D performance. In this case, the firm is given the choice of how to move projects forward in the process based on their probability of success.

Overall, our findings suggest that there is a certain amount of "native" uncertainty inherent to R&D processes. However, firms with their behaviors can exacerbate the volatility of R&D output. That is, policies firms use to manage their portfolio do affect volatility. Specifically, firms face a tradeoff between two performance measures: in most cases, indeed, the use of a certain policy leads to a better outcome in terms of R&D

[29] A large pharmaceutical company, for instance, might have 100 projects in various stages of development.

output, but at the cost of higher R&D volatility. Both the policies we modeled (i.e., resource allocation and risk preferences) are mechanisms for influencing the balance of a firm's pipeline. However, the way they affect it is different. Indeed, while resource allocation affects where (more) capacity is added in the process, risk preferences affects the screening process a firm uses to move projects forward. In both cases, in order to "gain" in terms of volatility, and not "lose" too much in terms of output, a firm should try to keep its pipeline balanced over time. If one thinks of each of the phases of the development process as independent, this is equivalent to saying that a firm should diversify its focus (in terms of either resources to allocate or thresholds used to progress projects) among the phases in order to face the uncertainty and unpredictability of R&D performance over time.

Our findings contribute to both theory and practice. As for the contribution to practice, our results provide some insights about how managerial policies might influence portfolio performance. Such insights set the stage for further research and discussion on the topic of the role of firms' behavior in R&D portfolio management. As for the theoretical side, there are two main contributions. First, we underlie the drivers of volatility and highlight the importance of the interaction of a set of "environmental" or exogenous characteristics of the R&D process (lead time uncertainty, technical uncertainty, information availability) and managerial policies (resource allocation rules and risk preferences). As for the second, it stands from the development of a general model of multi-stage R&D processes that can be easily implemented to study the impact of different rules on outcome measures, such as output, total costs, and volatility.

The present paper suggests several lines of future research. First, we need to validate our model empirically. This might be done by using data on R&D performance of pharmaceutical companies over time to explore the validity of the predictions of the model. Further research could also look at the interaction effects of policies and parameters of the process.

Another interesting direction for future research is the study of licensing impacts. Indeed, a common strategy in the pharma industry is to license-in products to fill in "gaps" in the pipeline. For instance, if a company is weak on phase III projects, then it might license those in. If we allow the possibility of licensing in projects at any stage during each period, how do licensing decisions affect R&D performance? Do they make the problem worse? One licenses in a phase III project that consumes resources so no earlier projects are moved up or no new projects are started. This leads to a future gap in phase III, so a firm has to license in even more products.

At the current stage, the model does not include the possibility of projects to be licensed in. Neither it included the possibility for firms to learn from failures. Indeed, there is no learning in the model. Many of the biases studied in psychology tend to limit people's ability to learn. For instance, individuals are prone to attribute good outcomes to their own skills, and bad outcomes to the luck of the draw. This is *self attribution bias*. When people suffer such bias, they do not learn from their mistakes simply because they do not seem them as their mistakes. Including learning from failure and the impact of behavioral biases in learning is thus another line of research it might be worth pursuing in the future.

Finally, it might be interesting to explore the impact of clusters of projects. Both correlated and competitive projects could be considered. In the case of correlated projects, we might model a situation in which if project A fails then the probability of success of project B decreases. Instead, in the case of competitive projects, we might explore a situation in which if project A fails, then the probability of success of the competitive project B increases.

Figure 3

The Effect of Scale on R&D Output
(assuming Constant Returns to Scale)

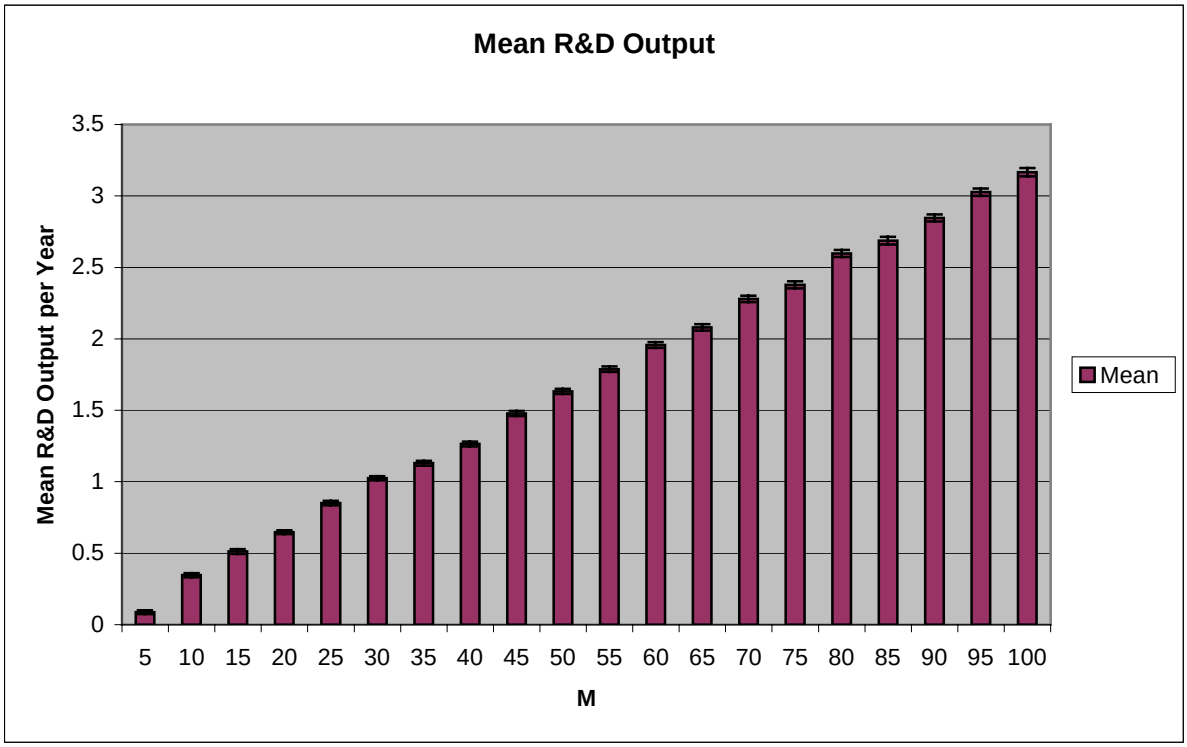

Figure 4

Impact of Scale on Mean Time Between Launch

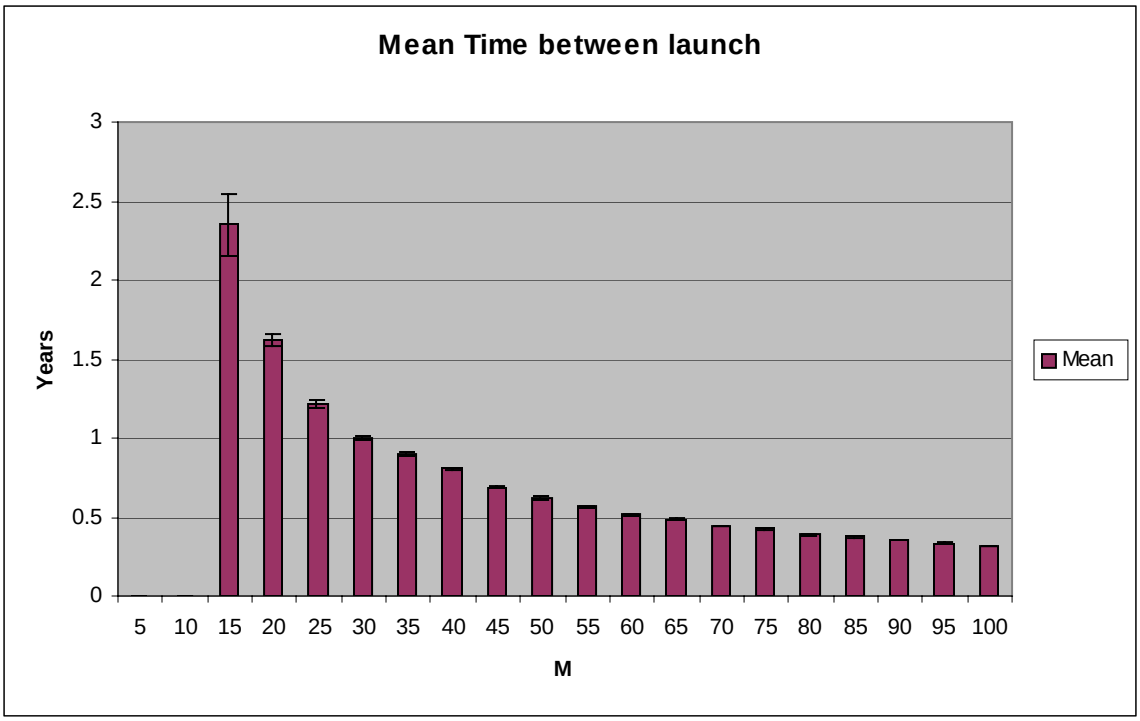

Figure 5

Impact of Scale on R&D Output Volatility

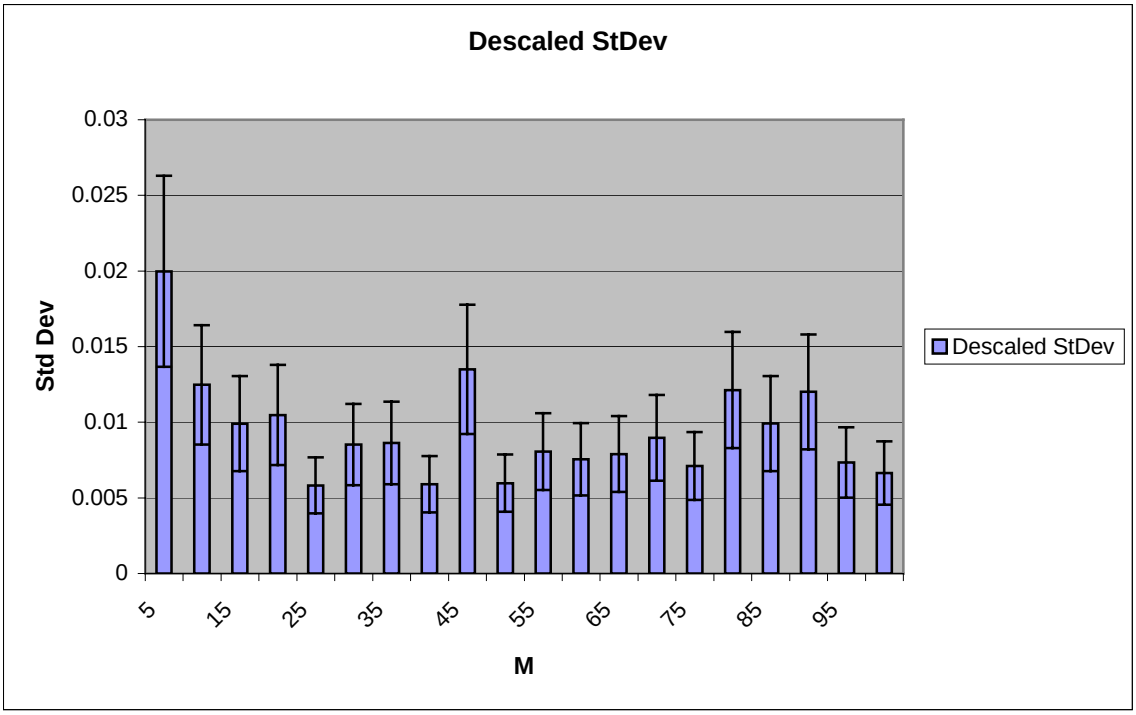

Figure 6

Impact of Resource Prioritization Rules o R&D Output

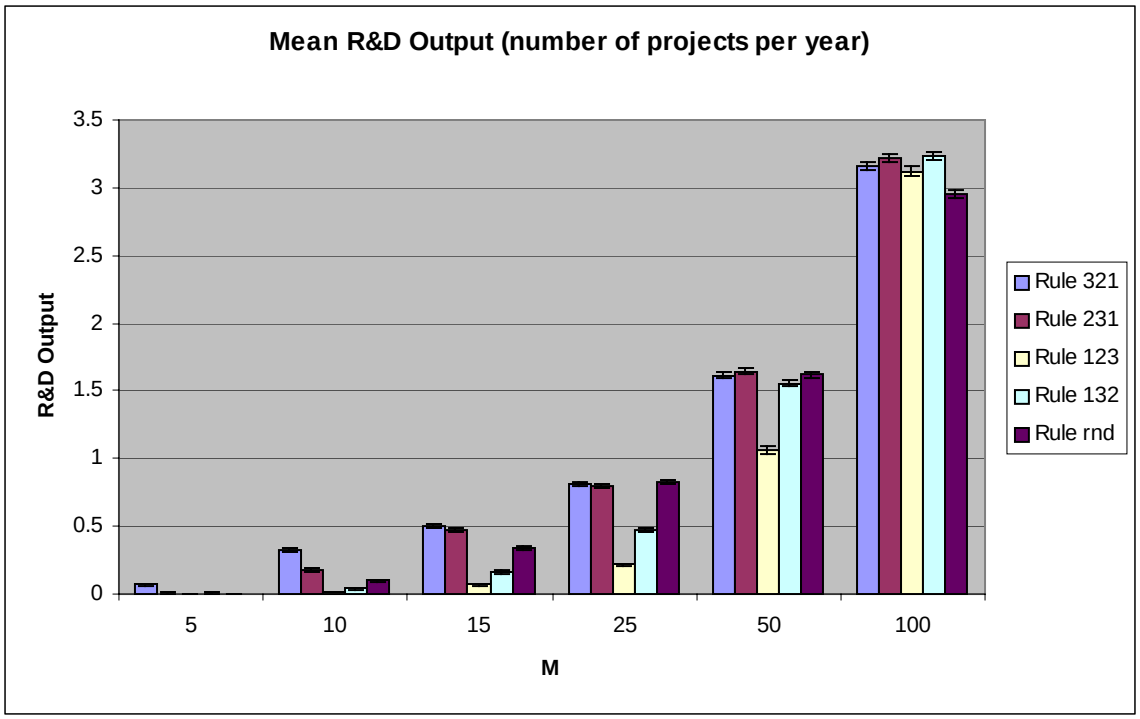

Figure 7

Impact of Resource Prioritization Rules on R&D Output Volatility

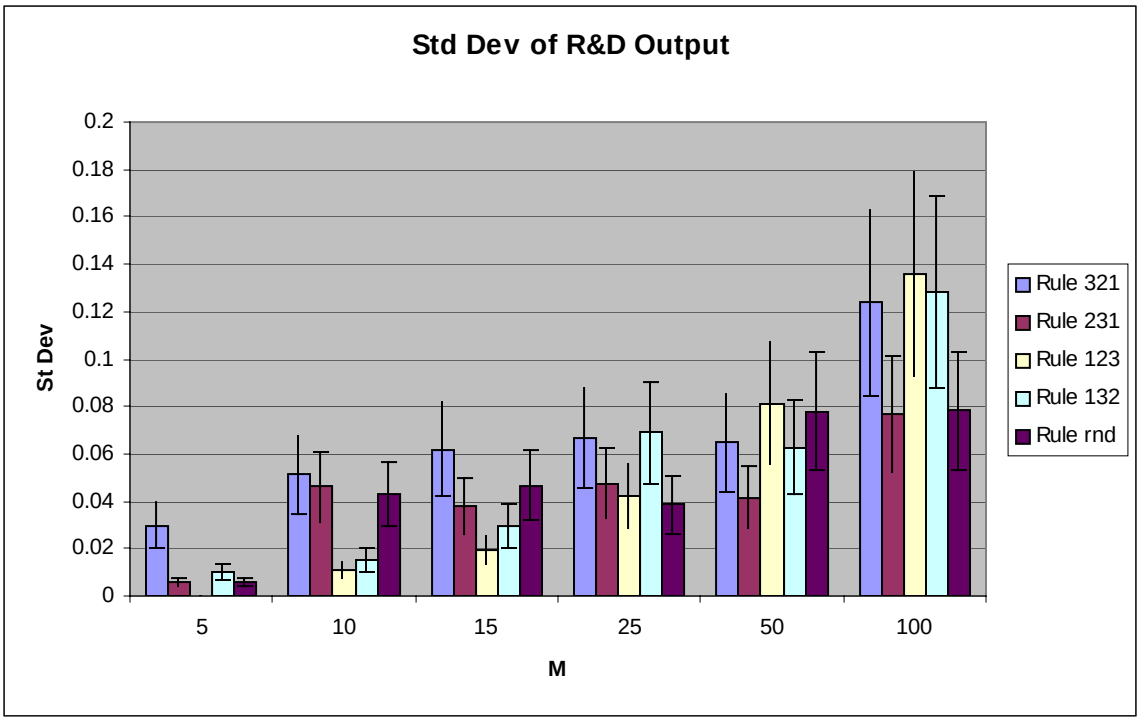

Figure 8

Impact of Risk Preferences on R&D Output[30]

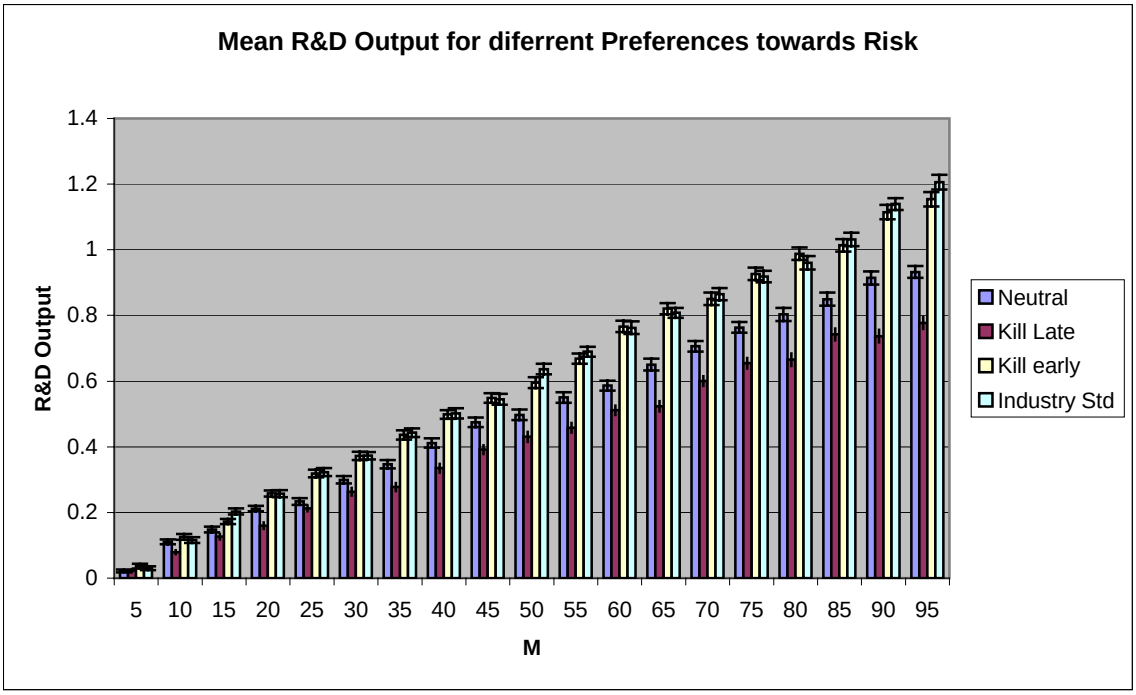

[30] In the figures used to illustrate the effects of risk preferences on R&D performance we use the term "kill early" to refer to risk averse firms, "kill late" to refer to risk lover firms, and "neutral" to refer to risk neutral ones.

Figure 9

Impact of Risk Preferences on R&D Output Volatility

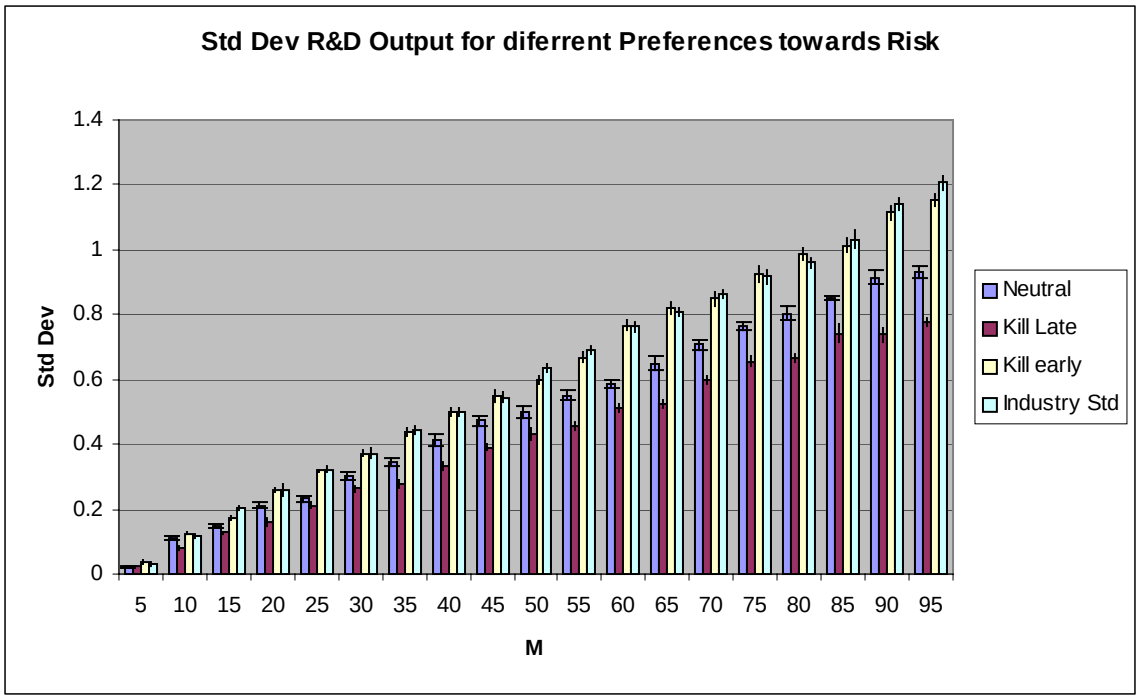

Figure 10

Impact of Risk Preferences on R&D Costs Per Successful Drug

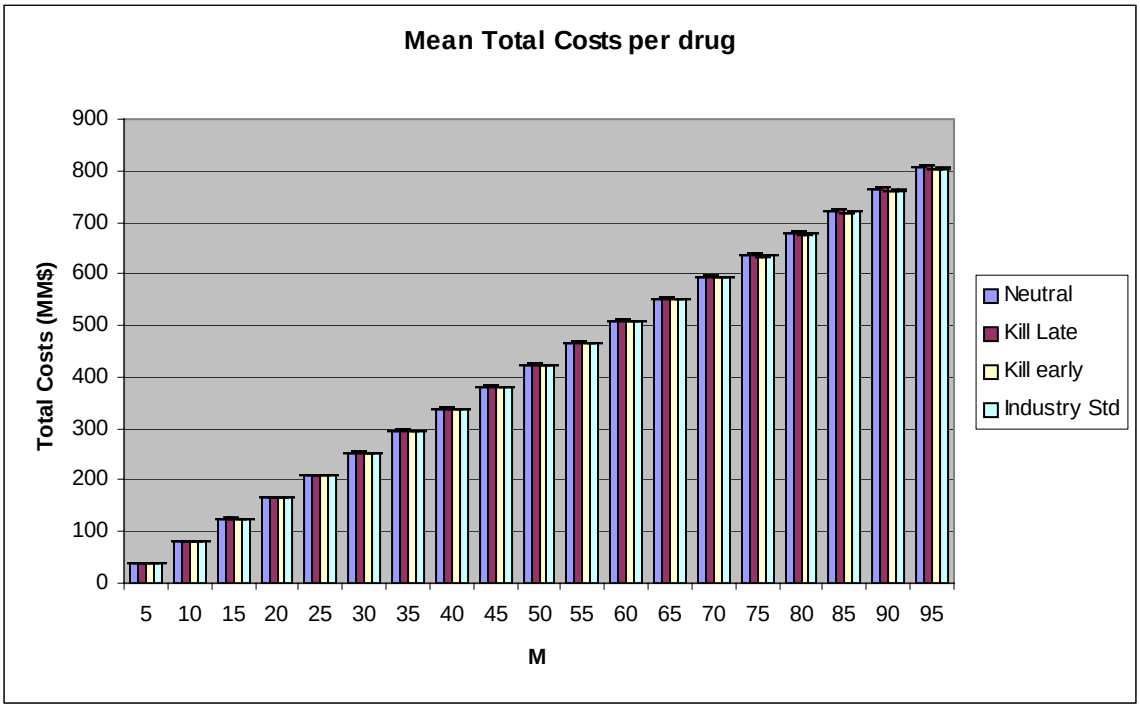

Figure 11

Risk Preferences and Killed Winners

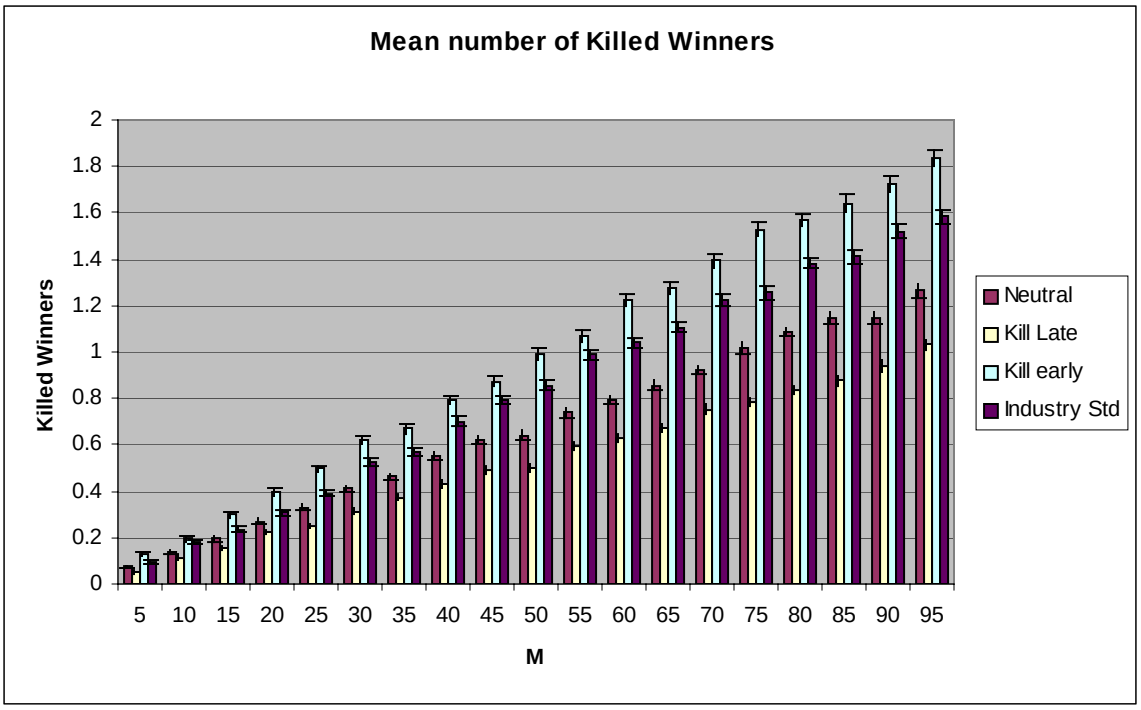

Figure 12

Risk Preferences and Non Killed Losers

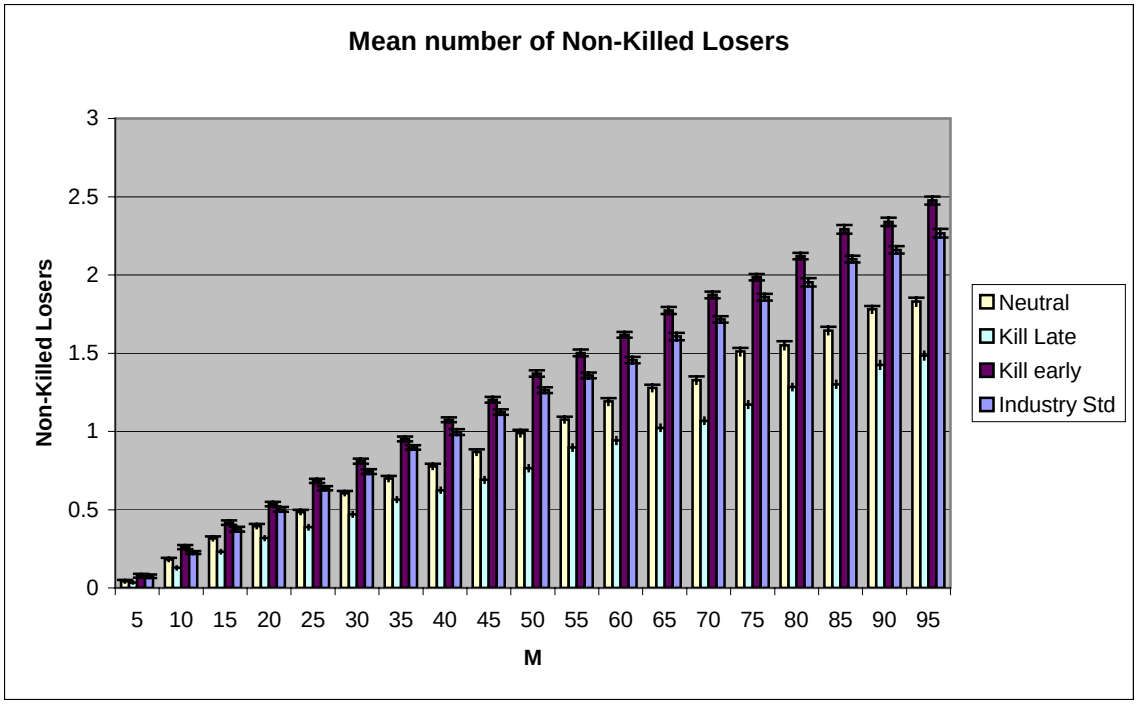

References

Ainslie, A., X. Dreze and Zufryden, F., "Competition in the Movie Industry," Working Paper at the University of California, 2002, Los Angeles, December.

Athey, S., J. Levin, "The value of information in monotone decision problems," Working paper, 2000, MIT, Cambridge, MA.

S. Banerjee, W.J. Hopp, "The Project Portfolio Management Problem," Department of Industrial Engineering and Management Sciences, 2001, Northwestern University, June.

Blackwell, D., "Comparisons of experiments," *Proc. Second Berkeley Sympos. Math. Statist.,* 1951: 93-102.

Blackwell, D., "Equivalent comparisons of experiments," *Ann. Math. Statist.,* 1953, 24: 265-272.

Bower, Joseph. L., *Managing the Resource Allocation Process: A Study of Corporate Planning and Investment*, 1970, Boston, MA: HBS Division of Research.

Clark, Kim, and Taka. Fujimoto, *Product Development Performance*, 1991, Boston, MA: Harvard Business School Press.

Christensen, Clayton M. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, 1997, Boston, Mass.: Harvard Business School Press.

Cooper, R. G., S. J. Edgett, E. J. Kleinschmidt, *Portfolio Management for New Products*, 1998, Perseus Books, New York.

Craig, C. Samuel, William H. Green and Susan P. Douglas, "Culture Matters: A Hierarchical Linear Random Parameters Model for Predicting Success of US Films in Foreign Markets", Working Paper, 2003, Stern School of Business, New York University, August.

Dixit, Robert and Robert Pindyck, *Investment Under Uncertainty*, Princeton, NJ, 1994: Princeton University Press.

Fleming, Lee, "Recombinant Uncertainty in Technological Search," *Management Science*: Volume 47:1, 2001: 117-132.

Froot, K., David S. Scharfstein, and J. Stein. "A Framework for Risk Management." *Journal of Applied Corporate Finance* (Fall 1994): 22-32.

Grabowski, Henry, and John Vernon, "Innovation and Structural Change in Pharmceuticals and Biotechnology," *Industrial and Corporate Change*, vol 3., no. 2, 1994: 435-449.

Guedj, Ilan and David Scharfstein, "Organizational Form and Investment: Evidence from Drug Development Strategies of Biopharmaceutical Firms," HBS working paper, 1994, March.

Henderson, Rebecca and Ian Cockburn, "Scale, Scope and Spillovers: The Determinants of Research Productivity in Drug Discovery." Iain Cockburn. *Rand Journal of Economics*, Spring 27(1), 1996: 32-59.

Henderson, Rebecca and Kim Clark, "Architectural Innovation: The Reconfiguration of Existing Product Technologies and The Failure of Established Firms." *Administrative Science Quarterly*, Vol 35, 1990: 9-30.

Hopp, W.J., "A Sequential Model of R&D Investment Over an Unbounded Time Horizon," *Management Science* 33 (4), 1987: 500-508.

Iansiti, Marco, *Technology Integration*, Boston, MA, 1997: Harvard Business School Press.

Jewitt, I., "Choosing between risky prospects: the characterization of comparative statics results, and location independent risk", *Management Science*, v.35 n.1, Jan. 1989: 60-70.

Liberatore, M., and Titus, G., "The Practice of Management Science in R&D Project Management," *Management Science*, Vol. 29, No. 8 (August), 1983: 962-974.

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

Loch, C. H., M. T. Pich, M. Urbschat, C. Terwiesch, "Selecting R&D projects at BMW: A case study of adopting mathematical programming models," *IEEE Trans. Engrg. Management* 48(1), 2001: 70-80.

MacCormack, A. and E. D'Angelo, "Activision: The 'Kelly Slater's Pro Surfer' Project", HBS case study, N9-605-020, 2004.

Nelson, Richard, and Sidney Winter, "In Search of Useful Theory of Innovation," *Research Policy*, 1977.

Nelson, Richard and Sidney Winter, *An Evolutionary Theory o Economic Change*, Cambridge, MA, 1982: Harvard University Press.

Paraxel, 2000

Persico, Nicola, "Information Acquisition in Auctions," *Econometrica*, vol. 68(1), January, 2000: 135-148.

Pisano, Gary, *The Development Factory: Unlocking the Potential of Process Innovation*. Boston, 1996: Harvard Business School Press.

Rivkin, Jan W., "Imitation of Complex Strategies", *Management Science*, 46 (6), June, 2000: 824-844.

Siggelkow, Nicolaj, "Misperceiving Interactions among Complements and Substitutes: Organizational Consequences." *Management Science*, 48, 2002: 900-916.

Simon, Carl P., and Lawrence Blume, *Mathematics for Economists.* New York, 1994: Norton and Company.

Subramanian, D., Pekny, JF and Reklaitis, GV, "A Simulation-Optimization Framework for Addressing Combinatorial and stochastic aspects of an R&D pipeline management problem," *Computer and Technical Engineering*, 24, July, 2000: 1005 - 1011.

Teece, David, Gary Pisano, and Amy Shuen, "Dynamic Capabilities and Strategic Management." *Strategic Management Journal* 18, no. 7 (August), 1997: 509-533.

Thomke, Stefan. *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*. Boston, Massachusetts: Harvard Business School Press, 2003.

Tushman, Michael, and Philip Anderson., "Technological Discontinuities and Organizational Environments." *Administrative Science Quarterly* 31, 1986: 439-465.

Tversky, A. and Kahneman, D., "Judgement under uncertainty: Heuristics and biases." *Science*, 185, 1974: 1124-1130.

Ulrich, Karl and Steven Eppinger, *Product Design and Development*, McGraw-Hill, 2003.

Vogel, Harold L., *Entertainment Industry Economics: A Guide for Financial Analysis*, Fifth Edition, 2001, Cambridge University Press.

Wheelwright S C, Clark K B, *Revolutionizing product development: quantum leaps in speed efficiency and quality*, 1992, Free Press, New York.