European Journal of Operational Research 242 (2015) 465–476

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Stochastics and Statistics

Project planning with alternative technologies in uncertain environments

Stefan Creemers [a], Bert De Reyck [b], Roel Leus [c,*]

[a] Management Department, IESEG School of Management, Lille, France
[b] Department of Management Science & Innovation, University College London, United Kingdom
[c] Research Group ORSTAT, Faculty of Economics and Business, KU Leuven, Belgium

ARTICLE INFO

Article history:
Received 26 August 2013
Accepted 5 November 2014
Available online 21 November 2014

Keywords:
Project scheduling
Uncertainty
Net present value
Alternative technologies
Stochastic activity durations

ABSTRACT

We investigate project scheduling with stochastic activity durations to maximize the expected net present value. Individual activities also carry a risk of failure, which can cause the overall project to fail. In the project planning literature, such technological uncertainty is typically ignored and project plans are developed only for scenarios in which the project succeeds. To mitigate the risk that an activity's failure jeopardizes the entire project, more than one alternative may exist for reaching the project's objectives. We propose a model that incorporates both the risk of activity failure and the possible pursuit of alternative technologies. We find optimal solutions to the scheduling problem by means of stochastic dynamic programming. Our algorithms prescribe which alternatives need to be explored, and how they should be scheduled. We also examine the impact of the variability of the activity durations on the project's value.

1. Introduction

Projects in many industries are subject to considerable uncertainty, due to many possible causes. Factors influencing the completion date of a project include activities that are required but that were not identified beforehand, activities taking longer than expected, activities that need to be redone, resources being unavailable when required, late deliveries, etc. In research and development (R&D) projects, there is also the risk that activities may fail altogether, requiring the project to be halted completely. This risk is often referred to as *technical risk*. In this text, we focus on two main sources of uncertainty in R&D projects, namely uncertain activity durations and the possibility of activity failure: we incorporate the concept of *activity success or failure* into the analysis of projects with stochastic activity durations, where the successful completion of an activity can correspond to a technological discovery or scientific breakthrough. We examine the impact of these two factors on optimal planning strategies that maximize the project's value, and on its value itself.

This work is a continuation of De Reyck and Leus (2008), where an algorithm is developed for project scheduling with uncertain activity outcomes and where project success is achieved only if all individual activities succeed. Reference De Reyck and Leus (2008) constituted the first description of an optimal approach for handling activity failures in project scheduling, but neither stochastic activity durations nor the possibility of pursuing multiple alternatives for the same result, and the inherent possibility of activity selection, were accounted for. Earlier work studied optimal procedures for special cases; see Chun (1994), for instance. Other references relevant to this text stem from the discipline of chemical engineering, mainly the work by Grossmann and his colleagues (e.g., Jain & Grossmann, 1999; Schmidt & Grossmann, 1996), who studied the scheduling of failure-prone new-product development (NPD) testing tasks when non-sequential testing is admitted. They point out that in industries such as chemicals and pharmaceuticals, the failure of a single required environmental or safety test may prevent a potential product from reaching the marketplace, which has inspired our modeling of possible activity and project failure. Therefore, our models are also of particular interest to drug-development projects, in which stringent scientific procedures have to be followed in distinct stages to ensure patient safety, before a medicine can be approved for production. Such projects may need to be terminated in any of these stages, either because the product is revealed not to have the desired properties or because of harmful side effects. Illustrations of modeling pharmaceutical projects, with a focus on resource allocation, can be found in Gittins and Yu (2007) and Yu and Gittins (2008).

Due to the risk of activity failure resulting in overall project failure, it has been suggested that R&D projects should explore multiple alternative ways for developing new products (Sommer & Loch, 2004). To mitigate the risk that an individual activity's failure jeopardizes the

* Corresponding author. Tel.: +32 16 32 69 67; fax: +32 16 32 66 24.
E-mail addresses: s.creemers@ieseg.fr (S. Creemers), bdereyck@ucl.ac.uk (B. De Reyck), Roel.Leus@kuleuven.be (R. Leus).

http://dx.doi.org/10.1016/j.ejor.2014.11.014
0377-2217/

entire project, we model projects in which the same (intermediate or final) outcome can be pursued in several different ways, where one success allows the project to continue. The different attempts can be multiple trials of the same procedure or the pursuit of different alternative ways to achieve the same outcome, e.g., the exploration of alternative technologies. Following Baldwin and Clark (2000), a unit of alternative interdependent tasks with a distinguished deliverable will be called a *module*.

Project profitability is often measured by the project's net present value (NPV), the discounted value of the project's cash flows. This NPV is affected by the project schedule and therefore, the timing of expenditures and cash inflows has a major impact on the project's financial performance, especially in capital-intensive industries. The goal of this article is to find optimal scheduling strategies that maximize the expected NPV (eNPV) of the project while taking into account the activity costs, the cash flows generated by a successful project, the variability in the activity durations, the precedence constraints, the likelihood of activity failure and the option to pursue multiple trials or technologies. Thus, this article extends the work of Buss and Rosenblatt (1997), Benati (2006), Sobel, Szmerekovsky, and Tilson (2009) and Creemers, Leus, and Lambrecht (2010), who focus on duration risk only, and of Schmidt and Grossmann (1996), Jain and Grossmann (1999) and De Reyck and Leus (2008), who look into technical risk only (although Schmidt and Grossmann (1996) also explore the possibility of introducing multiple discrete duration scenarios).

Our contributions are fourfold: (1) we introduce and formulate a generic model for optimal scheduling of R&D activities with stochastic durations, non-zero failure probabilities and modular completion subject to precedence constraints; to the best of our knowledge, such a model has never been studied before; (2) we develop a dynamic-programming recursion to determine an optimal policy for executing the project while maximizing the project's eNPV, extending the algorithm of Creemers et al. (2010) with activity failures, multiple trials and phase-type (PH) distributed activity durations instead of exponentials; (3) we conduct numerical experiments to demonstrate the computational capabilities of the algorithm; and (4) we examine the impact of activity duration risk on the optimal scheduling policy and project values. Interestingly, our findings indicate that higher operational variability does not always lead to lower project values, meaning that (sometimes costly) variance reduction strategies are not always advisable. To the best of our knowledge, this is the first article to numerically support such a recommendation.

The remainder of this text is organized as follows. In Section 2, we provide the necessary definitions and a detailed problem statement. We produce solutions by means of a backward dynamic-programming recursion for a Markov decision process, which is discussed in Section 3. Section 4 reports on our computational performance on a representative set of test instances. In Section 5, a computational experiment is described in which we examine the effect of activity duration variability on the eNPV of a project and Section 6 evaluates two different choices for the policy class to be considered. Section 7 contains a brief summary of the text.

2. Definitions and problem statement

2.1. Stochastic project scheduling

A project consists of a set of activities $N = \{0, \ldots, n\}$. The execution of a project with stochastic components (in our case, stochastic activity outcomes and durations) is a dynamic decision process. A solution, therefore, cannot be represented by a schedule but takes the form of a *policy*: a set of decision rules defining *actions* at *decision times*, which may depend on the prior outcomes. Decision times are typically the start of the project and the completion times of activities; a tentative next decision time can also be specified by the decision maker. An action entails the start of a precedence-feasible set of activities (see Section 2.2 for a statement of the precedence constraints). In this way, a schedule is constructed gradually as time progresses. Next to the information available at the start of the project, a decision at time t can only use information on duration realizations and activity outcomes that has become available before or at time t; this is the so-called *non-anticipativity constraint*. Activities should be executed without interruption.

Each activity $i \in N\backslash\{n\}$ has a probability of technical success p_i; we assume that $p_0 = 1$. We do not consider (renewable or other) resource constraints and assume the outcomes of the different tasks to be independent. We define a *success (state) vector* as an n-component binary vector $\mathbf{x} = (x_0, x_1, \ldots, x_{n-1})$, with one component associated with each activity in $N \setminus \{n\}$. We let X_i represent the Bernoulli random variable with parameter p_i as success probability for each activity i, and we write $\mathbf{X} = (X_0, X_1, \ldots, X_{n-1})$. Information on an activity's success (the realization of X_i) becomes available only at the end of that activity. We say that $\mathbf{x}$ is a *realization* or *scenario* of $\mathbf{X}$. The duration $D_j \geq 0$ of each activity j is also a stochastic variable; the vector $(D_0, D_1, \ldots, D_n)$ is denoted by $\mathbf{D}$. We use lowercase vector $\mathbf{d} = (d_0, \ldots, d_n)$ to represent one particular realization of $\mathbf{D}$, and we assume $Pr[D_0 = 0] = Pr[D_n = 0] = 1$.

We assume that all activity cash flows during the development phase are non-positive, which is typical for R&D projects: the (known) cash flow associated with the execution of activity $i \in N\backslash\{n\}$ is represented by the integer value $c_i \leq 0$ and is incurred at the start of the activity. We choose $c_0 = 0$. If the project is successful (see Section 2.2 for the specific conditions under which this is true) then the final activity n can be executed. This corresponds with obtaining an end-of-project payoff $C \geq 0$, which is received at the start of activity n (which is also its completion time). The value $s_i \geq 0$ represents the starting time of activity i; we call the $(n+1)$-vector $\mathbf{s} = (s_0, s_1, \ldots, s_n)$ a schedule, with $s_i \geq 0$ for all $i \in N$. We assume $s_0 = 0$ in what follows: the project starts at time zero. The value $s_i = +\infty$ means that activity i will not be executed at all.

We follow Igelmund and Radermacher (1983), Möhring (2000) and Stork (2001), who study project scheduling with resource constraints and stochastic activity durations, in interpreting every scheduling policy Π as a function $\mathbb{R}_{\geq}^{n-1} \times \mathbb{B}^n \rightarrow \mathbb{R}_{\geq}^{n+1}$, with $\mathbb{R}_{\geq}$ the set of non-negative reals and $\mathbb{B} = \{0, 1\}$. The function Π maps given samples $(\mathbf{d}, \mathbf{x})$ of activity durations and success vectors to vectors $\mathbf{s}(\mathbf{d}, \mathbf{x}; \Pi)$ of feasible activity starting times (schedules). For a given duration scenario $\mathbf{d}$, success vector $\mathbf{x}$ and policy Π, $s_n(\mathbf{d}, \mathbf{x}; \Pi)$ denotes the makespan of the schedule, which coincides with project completion. Note that not all activities need to be completed (or even started) by s_n, nor that the realization of all X_i's needs to be known.

We compute the NPV for schedule $\mathbf{s}$ as

$$f(\mathbf{s}) = Ce^{-rs_n} + \sum_{\substack{i=1 \\ s_i \neq \infty}}^{n-1} c_i e^{-rs_i}, \qquad (1)$$

with r a continuous discount rate chosen to represent the time value of money: the present value of a cash flow c incurred at time t equals ce^{-rt}, where e is the base of the natural logarithm. Our goal in this article is to select a policy Π^* that maximizes $\mathbb{E}[f(\mathbf{s}(\mathbf{D}, \mathbf{X}; \Pi))]$, with $\mathbb{E}[\cdot]$ the expectation operator with respect to $\mathbf{D}$ and $\mathbf{X}$; we write $\mathbb{E}[f(\Pi)]$, for short. The generality of this problem statement suggests that optimization over the class of all policies is probably computationally impractical. We therefore restrict our optimization to a subclass that has a simple combinatorial representation and where decision points are limited in number: our solution space $\mathcal{P}$ consists of all policies that start activities only at the end of other activities (activity 0 is started at time 0). The solution space also contains policy Π_0, which corresponds with immediate abandonment of the project (formally, all starting times apart from s_0 are set to infinity), which will be preferable when C is not large enough compared to the costs of the