HOW TO MANAGE RISK BETTER

Apply the Jarrow and Turnbull model for credit-risky bond pricing to the calculation of an appropriate hurdle rate for each stage of a research project.

David R. Walwyn, David Taylor and Glenn Brickhill

OVERVIEW: *Risk management practices in the R&D departments of many chemical and pharmaceutical companies lack much of the rigor and sophistication of the equivalent corporations in the financial sector. For instance, investment decisions on research projects are guided by techno-economic indicators that do not reflect changes in the financial framework of the project such as the prevailing interest rate structure, or the risk of "project default" as a result of termination. Although much has been achieved over the last ten years in improving the returns from R&D through the implementation of quality assurance processes such as the stage–gate methodology, considerable potential still exists for improving the management of investment risk and ensuring a return for shareholders. The Jarrow and Turnbull model for credit-risky bond pricing can help to calculate the internal rate of return (IRR) spread or hurdle rate that should be applied at each stage of a research project.*

David Walwyn is a senior manager at the CSIR Division of Food, Biological and Chemical Technologies, Modderfontein, South Africa, where he is responsible for new business development, marketing and information technology strategies. He has previously managed both research portfolios and individual research projects aimed at developing new process technology for chemical and biotechnology products. He is a member of the South African Institution of Chemical Engineers, and regularly acts as an advisor to the South African government on science and technology policy. His current interests include stimulating biotechnology industries in developing countries. He has a Ph.D. in bio-organic chemistry from Cambridge University.
dwalwyn@csir.co.za

David Taylor is an associate professor in the School of Computational & Applied Mathematics at the University of the Witwatersrand, Johannesburg. He holds a Ph.D. in applied mathematics and is the coordinator of the program in advanced mathematics of finance, where he is currently supervising six research Masters students.

Glenn Brickhill is a quantitative analyst at Standard Corporate and Merchant Bank, Johannesburg. He worked as a research scientist at De Beer's Diamond Research Laboratory after obtaining his B. Sc. Hons in physics at the University of the Witwatersrand. More recently, he completed his Masters degree on credit risk and the pricing of credit derivatives in the advanced mathematics of finance program at the University of the Witwatersrand.

Research is, by its very nature, a speculative activity with an uncertain outcome. For instance, historical information has shown that fewer than 50 percent of companies in the chemical industry realize an acceptable return on their R&D portfolios (*1,2*). Nevertheless certain methodologies do exist, which, if correctly applied, can militate against failure. One of these is the stage–gate approach, which is a quality assurance system imposed on a development project to ensure that certain key questions are answered within the relevant project stages. Examples of such questions include: what is the market size, what are the technical risks, what is the return on investment (as measured by the internal rate of return, IRR)?

The outcome of a gate review is a binary stop/go decision, and IRR is frequently used as a hurdle factor to guide the decision. If a project return exceeds the company's working average cost of capital (WACC) plus a nominal amount, the project can proceed; if not, it is terminated. The standard practice is to use a flat hurdle rate, across all the stages of a project. This practice is based on an erroneous and misleading assumption because a number of factors are not constant during the life of a project. The technical and commercial risk, the stage period, the cost of capital (relevant borrowing rate) and the phasing and size of investment all change as the project proceeds from concept to commercialization. In particular, this practice may bias the decision against projects that have a good chance of success.

In this article we describe a new method for adjusting the IRR spread over each stage of a project, thereby ensuring that the overall capital reward targets are met.

0895-6308/02/$5.00 © 2002 Industrial Research Institute, Inc.

Project Profiles

As noted, research projects vary in investment and risk over the course of the development work. If we apply a number of standard definitions for each phase of a project, we can plot average values of investment and dropout rate as a function of stage (see Figures 1 and 2, and "Data and Definitions," page 39). The data in these figures have been extracted from the chemical industry and may not be typical of other industries (*2*).

The trajectories in Figures 1 and 2 are simplifications of the real situation. Every project will be different, and companies themselves will have different track records and experience levels. However, it is still useful to look at the average information, since it emphasizes the four-dimensional nature of the problem of setting hurdle rates as a function of stage, namely variations in stage time, magnitude of investment, changing interest rate structure or cost of capital, and finally risk profile.

Other authors have argued that the hurdle rate, or IRR, should be simply weighted according to the probability of success in each stage (*3,4*). For instance, if the chance of success of a project from stage 3 onward is 20 percent (one-in-five), and the usual hurdle rate is 15 percent, then the hurdle rate at gate 2 should be 75 percent. Although this approach is an improvement on its alternative, which is to not consider probability of success at all, and hence to fall short on return on investment, unfortunately it is not generally applicable because of the inaccuracy of its underlying assumptions. Thus, we propose an alternative approach for the calculation for a risk-adjusted IRR (or IRR spread) based on the theory of defaultable bond pricing.

Defaultable Bond-Pricing Theory

The process whereby R&D project viability is determined suggests that a technique from derivative pricing theory may be implemented, namely the model used in determining arbitrage-free prices for credit-risky bonds and their derivatives. This model can be modified to yield the required IRR for each risk-category according to the historical data associated with each category and can incorporate variable lending rates available (corresponding to different project lengths), as well as projects of different size and time span. We draw a parallel between the default of a bond issuer and the termination of a project at a particular stage. These data can be updated to reflect a change in the default/termination rates within an organization, or to reflect a desired success rate.

To understand the application of bond pricing theory in calculating IRR spread, the following background information is needed.

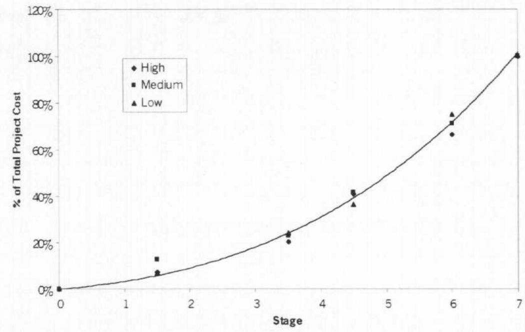

Figure 1.—R&D expenditure is not spread evenly over the course of a project; the early stages require comparatively less funding than the final stages of commercialization.

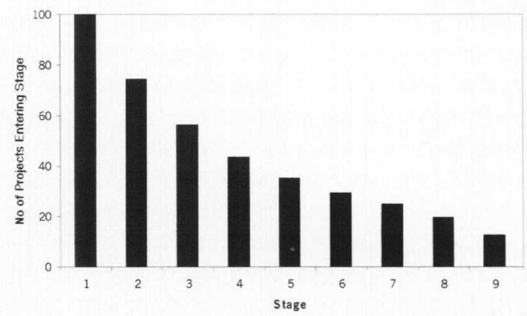

Figure 2.—Out of 100 ideas entering idea evaluation (stage 1), typically only 13 will be successfully commercialized (stage 9). Nevertheless, the cumulative "risk of failure" of a research project decreases as the project nears completion.

Credit-risky financial instruments

Debt of various maturities is issued by both the government and corporations in most countries. It is a tacit assumption by markets that the government debt has no possibility of default. This is dependent on the debt being issued in the sovereign currency. Any other issuer of debt is open to market scrutiny. If the market feels that the debt is of a lower investment grade than that of the government, it will demand a higher premium for the purchase. This translates simply into a higher rate of interest or return offered by the issuer and a consequentially lower price for the debt.

For example, if a government bond that promises R100 in a year's time is currently worth R90.91, this represents a simple rate of interest of 10% (R90.91 times 1.10 equals R100). A corresponding corporate bond, also promising R100 in a year's time, may trade at R89.29, which represents a simple interest rate of 12%. There is thus an inverse relationship between bond prices and the interest rate (or yield) that they represent. The higher the price, the more valuable the bond is regarded by the market and consequently, the lower the yield implied.

Pricing of credit-risky debt (bonds) begins with the assumption that the difference in price between bonds of other identical characteristics (face value, term-to-maturity) can be attributed solely to credit risk. Simply put, this is the possibility that the issuer of the bond may default on repaying it. Because there is a possibility that you might buy a "promise" of a future payment (or payments in the case of coupon-bearing bonds), which will not be redeemed, the issuer must offer a discount on the price. This translates into a higher yield.

A number of models for pricing this debt exist, of which the easiest to use is the discrete-time binomial model of Jarrow and Turnbull (5). This model assumes that over each discrete time step in the life of the bond, there exists a probability that the bond may go into default. Because the bond's value decreases dramatically (in extreme cases to zero) if default occurs, no matter how slight the probability, the value now must be less than it would be if there were no possibility at all. This binary default structure is then overlaid on an existing discrete-time model for pricing government debt and a value for the bond can be calculated. (Alternatively, the inverse problem of determining the implied probability of default given the observed value of the bond can be solved.)

The model has some complexity, as it is necessary to estimate from historical data what the effect of default will be on the bond's value. It is unfortunately part of the nature of debt regulations that bonds are not automatically worthless if the issuer goes insolvent. Their post-default value can be anything from 90 precent of their pre-default value to nothing.

In summary then, in order to determine a fair value for a defaultable bond, we need to know a structure of default probabilities (since the probability of default is not necessarily constant throughout the life of the debt) and what the effect of default would be (this is called the recovery value.)

Evaluating hurdle rates for R&D projects

The above methodology fits fairly neatly into the modelling problem concerning evaluation of hurdle rates for R&D projects. Evaluation takes place at set points: the hurdles. There are a fixed number of hurdles, and they are, in general, equally spaced in time. The project's progress can then be considered as a discrete-time process. At each review, projects either default (are ter-

Data and Definitions

Definition of Project Types

A number of definitions are used for distinguishing different project types, including distinctions based on size and duration, complexity and risk. Although each approach has its own advantage, the problem is that the data available are not consistent and it is necessary to make certain assumptions about the link between size/duration and risk.

The following definitions of project types (by size and duration) are assumed:

- A minor project is a project that results in a product extension, enhancement or minor process modification.
- A major project, on the other hand, is a project that results in the implementation of a new product or process that requires significant development in either chemistry or manufacturing.
- A platform project is a project that develops a new chemical entity, material or process that enables many individual product offerings.

Definition of Stages

Each project, regardless of its category is divided into the same number of stages as follows:

1—idea generation (or opportunity identification).

2—idea development (or idea growing and enabling science identification).

3—preliminary assessment (laboratory-scale investigation).

4—detailed assessment (laboratory and bench-scale studies).

5—development and validation (pilot-scale development).

6—basic engineering (leading to project report for capital sanction).

7—project execution and commissioning.

8—business operation and product on market (optional stage).

9—profitable business (optional stage).

However, the duration and cost of each stage are different, and dependent on the size of the project, as shown in Table 1. In practice, the actual numbers may be different from those listed here. Adjustments can be made using a spreadsheet calculation.

Dropout Rate

The dropout rate as a function of project type will be dependent on the experience within individual companies or institutions. For instance, the success rate of an electronics company may be far higher than that of a chemical company. The information provided in this article is obtained from both literature sources and personal communications, and is typical of the chemical industry. It is stated in three categories, as a function of project risk.

minated) or continue through to the next hurdle. A "default-free" project is one that makes it through all hurdles and into production. In this sense, it is akin to government debt in that it presents no possibility of loss of value through default during its life.

The process we are then trying to model is the inverse of what is modelled with a credit-risky bond. We do not know *a priori* which project will be rejected, but we do have historical information about the number of rejections at each hurdle and we can convert this data into default probabilities. We also know that at each stage the project will have spent a "known" fraction of its total cost. In the event of rejection (default), its "value" decreases from the amount already spent on it to zero. In other words, we know the probability of default at each hurdle, and we know the recovery rate.

The main assumption of the model is that a project can be valued at any stage as the expected value at that time of all of its possible future component cash flows. On the balance sheet, any project is a sequence of projected cash flows that may or may not occur. It is useful to consider a project guaranteed of success (risk free). This means that the project will not be rejected at any of the hurdles and will successfully reach production. At inception, the various developmental stages of this project will have expenditures associated with them. Given the prevailing interest rates, it is possible to invest at this time sufficient cash to cover all of these expenditures when they occur. This is a direct result of using familiar time-value-of-money equations. Consequently, it is possible to determine a value today for the project, which is the present value of all future cash flows. This calculation can be done at any stage (time) to value the project throughout its life.

A more useful model needs to incorporate the possibility of projects being terminated at any stage. The more risky a project is, the greater the probability that it will be terminated before its completion. It is possible to define a unique characterization of a risky project by its historical "default" structure. This is the set of probabilities defining the chance of a project being discontinued at any stage prior to completion.

This information is obtained statistically from historical data. This data mining can be carried out in such a way that "risk categories" can be determined across various project types. In this way, statistically relevant "default" probabilities can be obtained for high-capital-expenditure (high risk) projects through to low-capital-expenditure (low risk) projects. Once these default probabilities have been established, it is possible to value a risky project in a rigorous manner.

The risky project is valued in a similar manner to the guaranteed project, except that the probabilities of default are taken into account. This means that the present value of each cash-flow is adjusted by the probability of its occurring, conditional on its present state. If the cash-flow has already occurred, then the probability attached to it is one, whereas future cash flows have an element of uncertainty associated with each.

The above valuation implies that the value of a risky project will be less than that of a riskless project (as in the case of a risky bond). This means that the characteristic of a risky bond delivering a higher rate of interest than a riskless bond translates in project evaluations into a higher required IRR for riskier projects. This is an intuitively agreeable result. The valuation model described here determines the IRR that should be delivered by a project such that its "value" is a fair reflection of the associated risk. This valuation can be carried out at any stage and so a "fair" IRR can be determined at each project stage.

The results of the model are reported in terms of the IRR spread, or premium above the WACC that is required for each project within a particular stage. It is noted that the primary objective of the application of this model is to ensure that a portfolio of research projects achieves across time the WACC. Some projects will fail, and hence will write off the investment in that project to date. The IRR spread, if correctly applied, accounts for these investment losses in a thorough methodology and determines the intermediate IRR (which equals IRR spread plus WACC) for project evaluation.

Applying the Model to Historical Data

The utility of the model is illustrated by considering its application to the development of a novel chemical product X by chemical company A. It is assumed that X is a product that can be developed from an existing platform of technology, and therefore corresponds to a major project with a medium-risk profile, for which typical data can be extracted from the information provided in Tables 1 and 2 (*1,2*).

The model requires a number of inputs as follows:

- Number of stages over which the IRR spread is to be evaluated (this will depend on that version of the stage–gate methodology in operation within a specific company).
- Success rate as a function of stage (typical data are given in Table 2).
- Expected project cost or payment structure as a function of stage (typical data in Table 1).
- Expected duration of each stage (usually project-specific information, but in this case assumed to be 7 stages of equal duration of 7 months).
- Interest rate structure to be applied (in this case assumed to be flat at 5%).

Table 1.—Suggested Values for the Key Characteristics of Various Types of Projects

Characteristic	Project Type: Low Risk (Minor Project)	Medium Risk (Major Project)	High Risk (Platform Project)
Time to Market (months)[1]	16	36	62
Time to Profitability (months)[2]	32	42	82
Total Cost Excluding Capital ($)	266,129	1,693,548	14,516,129
Maximum People Employed (FTE)	4	7	26
Average Duration per Stage (months)	1.5	7.1	9.1
Cost Stage 1 ($)	8,000	50,806	435,484
Stage 2	14,933	94,839	812,903
Stage 3	23,653	150,218	1,287,581
Stage 4	33,867	215,081	1,843,548
Stage 5	46,667	296,371	2,540,323
Stage 6[3]	61,333	389,516	3,338,710
Stage 7[4]	78,133	496,210	4,253,226

(1) Stages 1 to 8, where stage 8 is commissioning; (2) Stages 1 to 9 where stage 9 is profitable operation; (3) Includes design charges; project approval is normally given at this point; (4) Excludes any capital expenses for plant equipment.

Table 2.—Previous Experience from the Chemical Industry Suggests That Out of 100 Project Ideas Entering Idea Evaluation, Only 32, 14 and 8 for High, Medium and Low-Risk Projects Respectively, Will Be Successfully Commercialized

Number of Projects Entering Stage	Low Risk	Medium Risk	High Risk
1	100	100	100
2	70	68	67
3	61	56	55
4	53	46	44
5	46	37	34
6	41	30	26
7	36	23	18
8	33	18	12
9	32	14	8

The results from the model are shown in Figure 3. For each stage of the project, the hurdle rate to be applied in the onward evaluation of the techno-economic feasibility of the project is calculated. For instance, at the end of stage 2, the project should be expected to achieve a return of at least 20 percent above WACC, whereas this figure will have fallen to 10 percent at the completion of construction and prior to commissioning.

The results from the model are most interesting. IRR spread is seen to vary significantly as a function of stage, and it is important that this indicator be adjusted throughout a development project in order to ensure that the total R&D portfolio achieves the targeted return on investment (such as WACC). If the variation of success rate with stage is ignored, it is inevitable that the latter will not be met and the net return on R&D expenditure will be below company targets.

A second example is drawn from published information for the cost of developing a novel pharmaceutical (6). In this case, the project runs over a ten-year period, and has an overall chance of success of 1 in 5000. For the purposes of this analysis, the project was divided into five two-year intervals, corresponding to initial screening, pre-clinical, clinical trials phase 1, phase 2

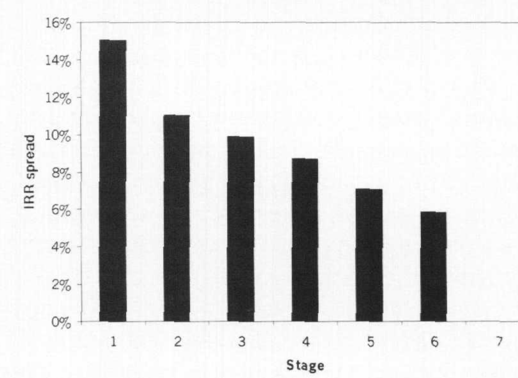

Figure 3.—The IRR spread varies from 15 percent to zero as a function of project stage; the closer the project gets to final commercialization, the closer the hurdle rate approaches the company's working average cost of capital (WACC).

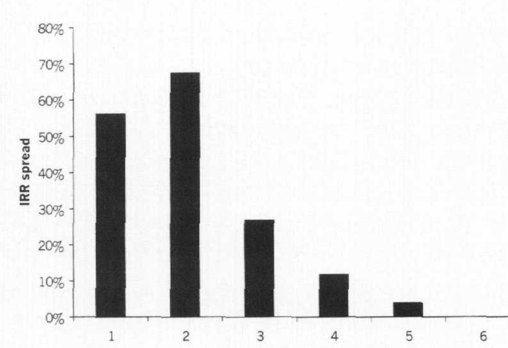

Figure 4.—A pharmaceutical project with a 1-in-5,000 chance of success should be subjected to a high initial hurdle rate as a result of the low probability of return from the initial investment in screening for new chemical entities.

and phase 3, respectively, and the required investment taken directly from the literature reference. A flat interest rate structure was assumed and the data generated as shown in Figure 4. It is noted that the initial hurdle rates are extremely high due to the high risk of failure in the early stages. Fortunately, the total value at risk is low because the bulk of the investment occurs in the clinical trials.

In order to generate an acceptable return from R&D investment, it is strongly recommended that this model be applied in all companies that follow stage–gate methodologies. For each application, the historical data set will have to be adjusted to reflect the company profile of research projects in terms of time, expenditure and outcome. The model itself has been encoded in an Excel program (called VentureSum™) which can be accessed directly from the Internet (*7*).

Another useful application of the model is to the valuation of intellectual property that has already been created.

Valuing Intellectual Property

Another useful application of the model is to the valuation of intellectual property that has already been created. This is a contentious and difficult problem, and one that weighs heavily on industry analysts and mergers and acquisitions consultants. For instance, valuations are often required of technologies that have not yet been fully commercialized. These technologies might be simply up for sale, or part of a company's intellectual property where the company itself is on the market. The value of an item on sale can usually be established by the buyer in one of two ways; namely, what someone else will pay for it (its going market rate), or as the net present value of its future cash flows. The establishment of technology exchanges such as yet2.com goes some way toward quantifying the former; the latter is usually highly speculative and subjective but can be estimated by applying our bond-pricing model to calculate the cost of completion of the technology, and hence its value as fully commercialized intellectual property.

For the seller, the deal price should at least exceed its cost of development or manufacture, and in the case of technology, the sum of any historical investment, inflated at an appropriate "discount rate." It is the latter quantity that is difficult to determine. The application of the bond-pricing model is a more rigorous method of determining the value of an incomplete technology. The assumption is made that its value can be calculated from the cost of its development, weighted for the statistical chance of success (WACC plus IRR spread). The model will supply the discount rate by which prior expenditure can be inflated to calculate the present value of the historical investment.

Summing Up

The model we have developed during this initial investigation appears to satisfy the criteria of rigor and applicability. In essence, it attempts to apply a financial model to the solution of a number of common problems for managers in chemical and pharmaceutical companies, namely how to achieve a satisfactory return from a portfolio of projects, and how to value an incomplete technology as either the seller or buyer. We are confident that it adds value to the process of portfolio management and project evaluation, and should be used in conjunction with other accepted techniques.

References

1. Pittiglio, Rabin, Todd and McGrath (PRTM) survey. "1997 Product Development Benchmarking Study—Chemicals and Plastics Industry Segment Report." November 1997.
2. Walwyn, D. and White T. "The challenge for innovators in the South African chemical industry." *Chemical World*, October 1998, pp. 26–31.
3. Nelson, T. R. and Mukherji, A. "Valuing biotechnology assets." *Nature Biotechnology*, 1998, Vol. 16, pp. 525–529.
4. Boer, F. P. "How do you assess the value of your technology?" *Chemtech*, January 1999, pp. 14–21.
5. Jarrow, R. A. and Turnbull S. M. "Pricing Options on Derivative Securities Subject to Credit Risk." *J. Finance* 1995, Vol. 50, No. 1, pp. 53–85.
6. Hansen, Kent F., Weiss, Malcolm A. and Kwak, Sangman. "Allocating R&D Resources: A Quantitative Aid to Management Insight." *Research · Technology Management*, July–August 1999, pp. 44–50.
7. See http://users.iafrica.com/g/gl/glenbrik/venturesum.htm