The Turbojet Engine Model of R&D Management

M.E.Glinsky, D. DePledge, J. Straccia, G. Davis and D. Nevelsteen

I. Executive Summary

A model for R&D management based on the analogy of a turbojet engine is presented. The goal is to use a system to manage our R&D that is as advanced as the R&D that we do. There are four guiding principles to the design of this system: simplicity, freedom, responsibility, and auditability.

The engine is divided into five stages: the intake, the compressor, the combustion chamber, the nozzle and turbo-charger, and the thruster and feedback controller. The intake, through a managed and coordinated process identifies technologies from internal R&D organization, business units, universities, marketing organizations, competitive intelligence, national laboratories, and project spin-offs. The compressor is a map of business units and business drivers to levers that technology projects can pull. This will allow for flexibility and quality control in determining the business value of technology. The combustion chamber is where technology meets business need creating business value, the ultimate goal. Technology projects are associated with levers, targets for those technology projects are set, costs and timing specified, deliverables stated, probability of success determined, and options for R&D outlined. The nozzle increases the efficiency of the engine by forcing a strict tollgating procedure on the system. The tollgates are reached by meeting the deliverables within the time frame and budget specified. Projects are allowed to continue on to the next tollgate according to a ranking based on an optimization of expected monetary value, EMV, constrained by risk, cash flow and resources. There is also a mechanism via which strategic alignment can also be incorporated in a quantitative way. The final stage is the thruster. Here the technology project will be implemented and/or commercialized – the business value will start to be extracted.

This process will be governed by a Research Committee that will control the process and arbitrate disputes. There will be an alignment of behavior and the process goal (i.e., optimizing EMV with constraints) by a system of incentive compensation.

There are several things that this management model adds. The first is an auditable and refinable valuation of each technology project based on the business technology maps and technology project targets – a valuation that includes realistic financial models and risked option valuation. The second is a quantified way to determine the level of R&D funding under risk, cash flow, and resource constraints. The third is a rational way to allocate R&D funding. The fourth is quantified value to present to host governments to justify the use and funding of technology. The fifth is an alternative method of R&D cost allocation to the business units. The sixth, yet most important, is an improvement in the efficiency of R&D in creating business value and the alignment of R&D with the corporate strategy.

II. Overview of the Turbojet Engine

The analogy that we use to develop the model of R&D management is that of a high-performance turbojet engine (see Fig. 1). The goal is to use a system to manage our R&D that is as advanced as the R&D that we do.

There are several guiding principles to the design of this engine. The first is simplicity. The process should be easy to understand. The simplicity will also prevent manipulation of the process that defeats the goal of maximizing the efficiency of the money spent on R&D. This simplicity should also be manifest in as little process overhead and burden on the researchers as possible.

The second is freedom. Teams and managers should be trusted to make the right decisions. Only at well-defined tollgates will there be an audit. This freedom to operate between tollgates will have an ancillary benefit of allowing a further optimization of the process.

The third is responsibility. Researchers and the people that oversee the process should be held responsible for their actions and decisions. This starts by designing a process where the value added by teams and individuals are quantified. It continues by linking the compensation of the teams and individuals by how much value they create. By aligning the motivation of teams and individuals with the goals of the R&D organization, the process controls can be simplified and freedom to optimize the process maintained.

The fourth is auditability. By requiring that targets are set and predictions of success are made, not only can the expected value of the R&D be calculated and used to control the process, a measurement can be made of the quality of those targets and predictions. This measurement can be used to improve the quality of the targets and predictions in the future. The result will be a more efficient process.

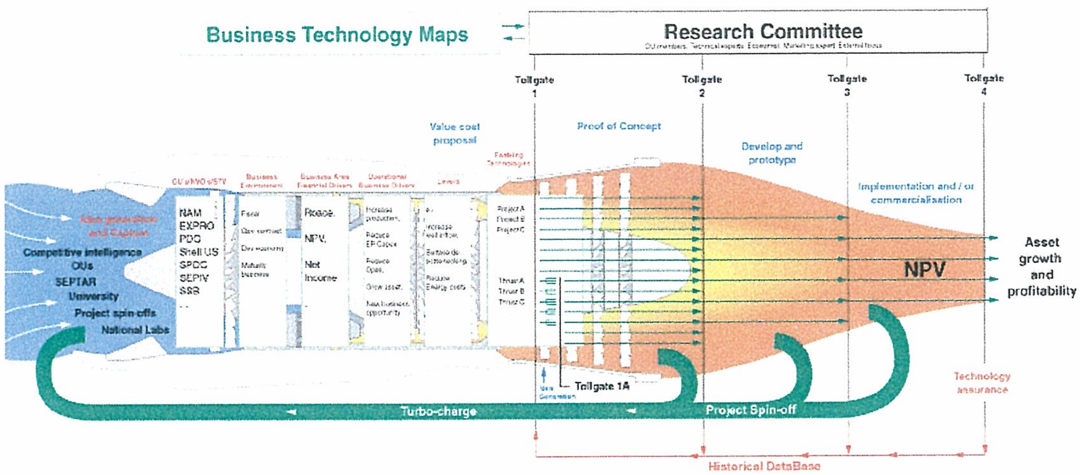

Figure 1. Schematic of the Turbojet Engine Model of R&D Management.

III. The Individual Stages of the Turbojet Engine

There are five stages to the turbojet engine: the intake, the compressor, the combustion chamber, the nozzle and turbo-charger, the thruster and feedback controller. Each one of these stages will be described and details of how they function are now presented.

A. Intake

The intake is designed to be an efficient collector of quality ideas - the air without which the engine would cease to function. This is a managed and coordinated process that has many sources of ideas. It is powered, as all of the stages, by an incentive compensation scheme (see Appendix 2).

The researchers will be empowered to search for ideas, with a modest budget (approximately one man month). This budget can be used to cover their time, fund external work, or pay for expenses. They will be compensated for each qualified idea that they generate. These ideas will be databased and quality controlled to eliminate duplication and frivolous ideas. The rules attached to spending this money should be the absolute minimum. The alignment of motivation and behavior will be relied upon to govern the behavior. This alignment of motivation and behavior will be a reoccurring theme in the overall process.

There will also be focused workshops and thrust areas. These efforts will be determined by the identified dangling levers, which will be discussed more fully in Sec. IIIB. Alignment with dangling levers will also be used as a criteria to fund external work. Funding of external work and interaction with the external world is an important part of this intake mechanism. This will allow us to identify promising scientific advancements before other R&D organizations and to find out the commercial value of the advancements, as well as file patents, before our competitors. Some external work will be funded through the budget made available to individual researchers, but there needs to be a larger source that is centrally managed.

The source of ideas will be the internal R&D organization, business units, universities, marketing organizations, competitive intelligence, national laboratories, and project spin-offs.

The funding for this effort should be a healthy fixed amount for the first few years, then allowed to float to its equilibrium level. This equilibrium level will be approached as a capital allocation (portfolio) decision.

The cost value proposition will determine this equilibrium level. A major input to the proposition will be the probability of success of these efforts. Because of how important this number is to the proposition and how important the intake is to the functioning of the tubo-jet engine, we have recommended the healthy fixed level until enough historical data can be collected to accurately determine the probability of success.

There will need to be a small group that will manage this process. Their function will be to determine which ideas are new and profound (i.e., worthy of compensation), database those ideas, organize thrust areas, and manage funding of external efforts.

This process places large emphasis on external sources of technology project ideas and puts in place a process to determine the optimum funding level for those external sources.

B. Compressor
The compressor concentrates the air supplied by the intake (ideas) by identifying the ways in which it can add value to our business. This is done by constructing the business part of the business technology maps.

The Business Technology Maps (BTMs) are a mapping of business units to business environments to financial drivers to operational drivers to levers to enabling technology projects. This mapping allows a flexible analysis and quality control of the process. An example of the flexibility is subtracting or adding business units as our business portfolio changes and examining the implication of the changes on our technology project portfolio. An example of the quality control of the process will be summing potential impacts of all technologies to ensure that the total impact of technology will be consistent with historical norms.

The process via which the business part of the BTMs are constructed starts by collecting information from the group economists on the business environment, business areas, financial drivers, operational drivers, barrel breakdowns, productions targets, cost targets, tax regime, and equity ownership. The business units will be welcome to supply this information, but as much use of the centrally located economists, as the business units feel comfortable, will be made. This is to reduce the data supply burden on the business units. A group of Technology Marketing Managers (TMMs) will then visit the business units once a year to define the business levers and their dependencies. Another group called the Research Committee will have the ability to challenge and modify the levers and their dependencies. The TMMs will have ownership of the BTMs. They will be posted on the web for challenge and for use as the process continues.

A few more words need to be said about the Research Committee. It is a group of 15 to 20 people. They are the body controlling the entire process and are the supreme court, arbitrating disputes in the system. In addition, they will have many more responsibilities outlined in later sections. They will meet once a month and be made up of business experts (senior business unit representatives), marketing experts (commercialization and implementation specialists), technical experts, economists, and external focus specialists. All of the members can revolve in and out as needs dictate. There will be a chair of this committee who will manage and plan the meetings, and will have responsibility for making sure that all necessary information is collected for making the decisions. This committee will be crucial to the process functioning well, and will meet for a week each month.

The compressor provides an auditable way to identify levers with business drivers and business value.

C. Combustion Chamber
The combustion chamber is where the air (technical ideas) meets the fuel (business levers) and ignite and generate the energy (business value). There is a two step process to constructing this value cost proposal. First, the researchers propose technology projects that will pull the business levers. They also set technology project targets for how much the levers will be pulled. Second, the TMMs go back to the business units to determine the applicability of these technologies.

The preparation of the research proposal will be owned by a research team conductor (RTC)[1]. Whenever a person or group of people have come up with a promising notional idea, the Research Committee will assign a RTC from an available pool. A RTC will be assigned multiple projects. It will be the responsibility of the RTC to arbitrate team disputes, to make sure that the proper resources and personnel are available to accomplish the project objectives, to prepare the research proposal, and to be the advocate for the project. The RTC will receive incentive compensation based on whether the projects they conduct meet their expected probability of success.

The research proposal will contain: an executive summary, a one page technology project description, a list of business levers which it could pull, the technology project targets of how much it will be able to pull those levers, buyup/maintain/buydown option, description of tollgate deliverables, cost estimates by tollgate, time estimates by tollgate, resource estimates by tollgate, possibilities for outsourcing at each phase of the research, description of competitive landscape, dependencies upon other projects, and test area if the project makes it to prototyping phase. The estimates of costs, resources, times will have minimum, maximum, and most likely. Deviation of targets, costs, resources, and times for historical average values will have to be justified. The structure of the tollgates will be described in the next section.

A project can pull more than one lever. It can also create its own branch of levers starting at any level of the BTM. This includes creating its own business unit if it will lead to a new business opportunity. This creation of new branches should be welcomed since it indicates innovation, that is, out-of-the-box thinking.

Since a lever may be pulled by more than one project, it should be determined by the TMMs, during their visit to the business units, whether the projects will have a complimentary effect or an exclusive effect.

The buyup option will consist of an alternate research plan that will increase the applicability of the technology project and/or decrease the time to implementation if more money is spent on the project. The buydown option is the opposite. The farmout option means different things at different tollgates. At the earlier tollgates, it could be contracting university, private research laboratories, petroleum service organizations, or national laboratories to do the proof-of-concept and/or development of a prototype. At the latter tollgate, it is commercialization. The farmin option is doing it internally. Things that should be captured in the definition of the farmout option should be the direct costs, cost of supervision, cost of loss of core competency, cost of internalization of results, time to next tollgate, and the availability of needed infrastructure. Detailed information only needs to be given for the next tollgate. Only an indication of whether the farmout option will exist for later tollgates should be given. The RTC will facilitate the discussion leading to the creation of these options.

The business units will be able to challenge the technology project targets, costs and timing. The researchers and RTCs will be able to challenge the applicability and the complimentary/exclusive determination of the TMMs. These challenges will be arbitrated by the research committee. The research proposals along with the business technology maps will be kept on the web.

The probability of success for the next phase of the research will be determined by a secret ballot of the research committee. They will vote on which bin of probability (0-10%,10-25%, 25-50%,50-75%,75-100%) the project falls. Probability of success for subsequent phases will be historical averages.

The targets that are set in the combustion chamber form the basis for calculating the asset based value add. There is also a structure for a fair and proper challenge of the value add. The R&D options outlined will allow for more flexibility in the system and increase the efficiency of R&D in creating business value.

D. Nozzle and Turbo Charger

The nozzle increases the momentum of the gas (technologies) thereby increasing the thrust (EMV) generated by the turbojet engine. The momentum of the gas is increased by a strict tollgating procedure controlled by the research committee. The turbo charger also increases the thrust of the engine by recirculating spinoff ideas and premature technologies back into the air intake.

The tollgating procedure starts by ranking the projects according to a maximization of the expected monetary value (EMV) with constraints on risk, cash flow, and resources (discussed in Appendix 1)[2,3]. A further modification to the ranking can be made to align it with the E&P strategy (discussed in Appendix 3). An overall funding level is determined by the point at which the risked value investment ratio (RVIRT) of adding the next project is equal to that of investments other than technology. The subscript on RVIRT is because the investment is the R&D cost and any additional capital needed to implement the technology. Once this optimum funding level is determined, projects are funded starting at the top of the ranked list until the optimum funding level is reached. This will be applied to all projects when the turbojet engine is first started. Thereafter, it will only be applied when a project reaches a tollgate.

The project development will be broken up into five stages: initial proposal, proof of concept, develop and prototype, implementation and/or commercialization, and exit. Each phase will be divided by a tollgate where a decision will be made. Generally, there will be seven options: farmin/farmout/minimal and buyup/maintain/buydown options for the farmin and farmout options. Upon reaching a tollgate, a research proposal for funding will be made. If the project is ranked above the line for optimum funding it will be granted its next round of funding. When the project has spent the most likely funding level or the most likely time for that phase is reached, it will come under review by the research committee. Only for good cause will approximately 50% of the projects be allowed to continue. Once the maximum limits have been reached, the project will be given minimal funding for abandonment. Only in very rare cases will the project be allowed to continue at more than the minimal funding level.

Researchers will have an incentive compensation based on having project on which they are working successfully pass a tollgate. The research committee will determine when a project has successful reached a tollgate by secret ballot. The criteria will be delivery of all deliverables associated with that tollgate as specified in the research proposal. Passing the tollgate is only associated with meeting the technical deliverables. Whether or not a project will be granted its next round of funding, once it has passed the tollgate, will be a strictly financial decision described in Appendix 1.

As the technologies move into the prototype phase, the research committee should appoint a user group to oversee development. This user group should be technical people who are representative of the people who will be using the technology. The research committee should rely on their judgment of whether or not a successful prototype is made. The researchers should be encouraged to have regular contact with the user group during the development.

An additional tollgate is required for research thrusts. Research thrusts are associated with dangling levers which have no technology projects associated with them. They are identified and proposed by the group that oversees the intake process. They will be assigned a RTC and will compete for funding along with the technology projects. The difference is that they will be trying to identify technologies that could pull the dangling lever. Any technology that they identify will have to have a proposal for research funding made for it. The consideration of this proposal is the additional tollgate.

The nozzle is a procedure for strict tollgating. This tollgating ensures that the option value of the process is realized.

E. Thruster and Feedback Controller

The thruster is where the air reaches it highest velocity at the end of the nozzle and is where the thrust (EMV) is realized. This happens as the technology is applied to our asset base or is commercialized. We consider that the technology has left our turbojet engine when it reaches full implementation. Information on the performance of the engine (historical data) should be collected and input to the feedback controller (research committee) to optimize the performance of the engine.

Centrally resourced funding should cover establishment of infrastructure for implementation and demonstration of widespread value. This will not cover ongoing costs once the infrastructure and value is established. This funding of the ongoing costs should come from fees associated with use of the technology. The centrally resourced funding should be viewed as the last round of venture capital funding. A marketing organization, either internal or external, should be the group driving the implementation. This marketing organization will be responsible for training, sales and marketing. As the technology exits the nozzle (reaches full implementation) all control should be turned over to the marketing organization and all centrally resourced funding should be discontinued. They should directly fund cosmetic changes to the product through user fees. Major new enhancements should be input as a proposal to the turbojet engine.

As the technology starts to be applied it will be the responsibility of the TMMs to do a technical assurance. This assurance is to ascertain via interviews with the field development teams, after the technology has been applied, the marginal value of the technology. This marginal value should be compared to the technical target, and used to normalize future technical targets. The timing of this technical assurance is critical. There is a limited window when the technology will be applied after all previous technology has tried. The technology will then become part of the base technology and separating its value from other technologies will become problematical.

A historical database should also be maintained on the probability of project passing tollgates (i.e., probability of success), number of spinoff projects, time to pass the tollgates, and costs to pass the tollgates. The purpose of this historic database is to make this an intelligent process that learns from itself and uses that information to optimize its future performance. Ownership of the historical database shall rest with the TMMs.

The definition of what is meant by full implementation is determined at tollgate 3. It should be agreed upon by the researchers, Research Committee and business units. The TMMs will collect the data that will verify the implementation. The business units and/or user groups must certify this data is accurate and meets the definition of full implementation before any incentive based compensation is made.

Costs of TMM/business unit interaction, research proposal preparation, RTC oversight per project and phase, the research committee should be captured so that the costs of this process can be tracked and managed.

The salesmen of the technology should be compensated according to how much they sell, that is, how much implementation they stimulate. The motivation of the thruster by incentive compensation of the sales force and a strong emphasis on self funding of the internal marketing organization will drive the process to adding the maximum business value. The feedback control given by the technical assurance will give additional business value.

IV. What is New about the Turbojet Engine

There are several things that this management model adds. The first is an auditable and refinable valuation of each technology project based on the BTMs and technology project targets – a valuation that includes realistic financial models and risked option valuation. The second is a quantified way to determine the level of R&D funding under risk, cash flow, and resource constraints. The third is a rational way to allocate R&D funding. The fourth is quantified value to present to host governments to justify the use and funding of technology. The fifth is an alternative method of R&D cost allocation to the business units.[♠] The sixth, yet most important, is an improvement in the efficiency of R&D in creating business value and the alignment of R&D with the corporate strategy.

Appendix 1. The Financial Model

The financial model consists of three main parts: an economic model with country and field specific contractual terms, a stochastic evaluation of the incremental value and costs of technology, and a constrained optimization of the EMV.

The economic model is packaged as a "black-box" module that is constructed and maintained by central office economic staff. This ensures the confidentiality regarding the contractual terms.

There are four basic types of fiscal models covering Royalty/Tax, Production Sharing, Margin and Technical Service agreements. Each model has, as far as practicable, been built to reflect the specific contractual arrangements for each field (or fields covered by the BTM) to account for incremental Shell, government, national oil company (NOC) and partner takes relating to the application of the various new

---

[♠] The model also gives a method for distribution of the R&D costs to the business units. See Appendix 5 for more information.

technologies. This approach has been taken in preference to the simplified "after tax $ spent" or "after tax $ earned" approach to properly account for government take, be it in cash (e.g. tax) or kind (e.g. profit oil). This is particularly important with Production Sharing Agreements (PSAs) with creaming mechanisms where increased production and/or reduced expenditures may not necessarily lead to an increase in Shell share NPV.

The inputs to each model have been standardized to accept baseline and incremental Production profiles, CAPEX profiles, OPEX profiles and R&D cost profiles. Standard outputs are baseline and incremental NPVs, PV CAPEX, OPEX and R&D costs split in to Shell share, partner share, NOC share and government take.

A stochastic evaluation is done for the incremental value and costs of the technology projects. All values and costs are PV with a discount rate of 7% (the Royal Dutch Shell Group WACC). The stochastic variables underlying this evaluation are: the oil and gas price, the conditional passing of a tollgate, the time to pass the tollgate, the cost of the research required to reach the tollgate, the additional CAPEX needed for implementation, the impact on CAPEX of the application of the technology, the impact on OPEX, the impact on recoverable hydrocarbon reserves, and the impact on the production rate. Stochastic variables associated with tollgates are specific to each technology project. Stochastic variables associated with the application and impact are specific to each technology project and each BTM. The distribution used for the oil and gas price is a triangular distribution whose minimum is $7/BBL (Brent), most likely is $14/BBL, and maximum is $21/BBL. Triangular distributions are also used for the time to pass the tollgates and the cost of the research. The minimums, most likely values and maximums are those supplied by the researchers on the research proposals. The distributions of the financial and production impact of the technologies are also triangular distributions whose most likely values are specified by the targets set by the business units. The minimum values are set to one half of the most likely and the maximum is set to twice the most likely. This is an attempt to capture the historical under estimation of the impact of technology. We have tended to have unanticipated success. The realized impacts of technologies are summed across all technologies and applied to the cost and production profiles input to the economic models. The realized Shell take and costs can then be summed on a company wide or operating unit basis. The realized incremental values or costs can be averaged over all the stochastic realizations leading to risked or expected values.

Three major classes of explicit dependencies are also captured in the stochastic evaluation: lever-project, project-project, and lever-lever. The lever-project dependencies capture the relationship of the projects to business value. Several projects could pull the same lever, but the lever can only be pulled one time. That is, the value can only be realized once. We therefore create an OR structure. The value of the lever is realized if project A OR project B is successful. This keeps us from claiming the value of the lever multiple times and grossly inflating the value of technology. The project-project dependencies capture the value of supporting or basic research (e.g., seismic inversion can only be successful if the rock physics is understood). A project can be dependent on several other projects in either an AND and/or OR fashion. Project A can be successful if Project B AND/OR Project C are successful. The lever-lever dependencies capture operational dependencies (e.g., the value of increased production rate can only be realized if the additional production can be processed) and business value augmentation (e.g., the value added by multilateral wells could be doubled if a better reservoir model is constructed). We use an AND structure. Lever A can be pulled if Lever B AND Lever C are pulled. Business value augmentation is modeled by creating a second lever which is dependent on both the base lever and the lever which augments the value. Both the project-project and lever-lever dependencies are recursive.

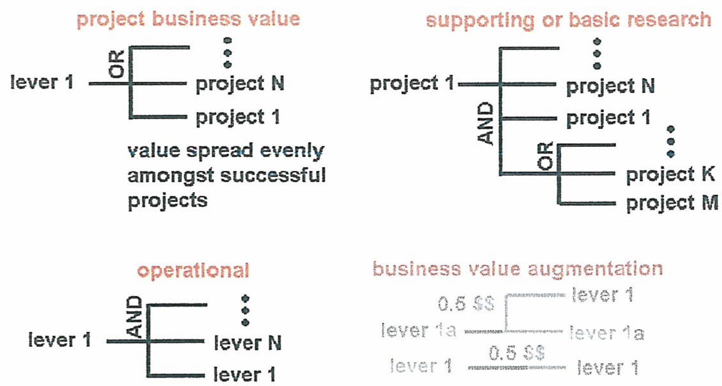

Figure 2. Explicit dependencies taken into account in the financial model: lever-project, project-project and lever-lever. Operational and supporting research dependencies are recursive.

The optimization can be broken down into two parts. The first is the definition of the objective function to be optimized and the constraints to be applied. The second is methodology used to do the constrained stochastic optimization.

The fundamental objective function that we wish to maximize is the risked incremental value (EMV). The primary constraint is a capital constraint so we would like to hold the risked Shell investment below a certain level. The Shell investment is the R&D cost plus any additional CAPEX that is needed to implement the technologies. There could also be additional constraints put on cash flow (maximum expected payback and/or maximum cash sink) and resources (maximum manpower). We have decided to apply only the capital constraint, then examine the cash flow profiles and resource needs to determine if any additional constraints need to be applied.

If the value of the technology projects were independent we could rank them according to their individual risked value investment ratio of technology (RVIRT), where the RVIRT is the EMV of applying the technology divided by the risked Shell investment in that technology project. Unfortunately, the explicit project-lever, project-project and lever-lever dependencies and the economic models that are not simple marginal contracts destroy the independency of the technologies' values. Therefore, we have had to adopt a more complicated approach. Ideally, we would like to form the portfolio of N technology projects by first determining the portfolio of 1 technology project to be pursued that maximizes the RVIRT. Then using this portfolio as the reference portfolio determining the portfolio of 2 projects that maximizes the RVIRT, where the incremental value and Shell investment from the reference portfolio is used. This process would be continued until the portfolio of N projects is reached. Obviously, N stochastic optimizations would need to be done. We have used an approximation of this approach. We have only done M optimizations. We have started by determining the portfolio of N/M technology projects that maximizes the RVIRT. A marginal analysis is done around this portfolio by adding a technology project if it is not in the portfolio, and subtracting it if it is in the portfolio, then calculating the marginal RVIRT of this change. A small number of manual changes are done based on these results, keeping the total number of technology projects in the portfolio constant. The financial performance of the portfolio with these manual changes is compared to the portfolio without the changes and the best is kept. The portfolio of 2N/M technology projects that maximizes the RVIRT is then determined using the optimum portfolio of N/M technology projects as a reference. This process is continued until the portfolio of N technology projects is determined. This has two advantages. It only takes M optimizations. It also involves a human intervention that will increase the confidence in the results and remove much of the "black box" nature of the optimization. It will be a reasonable approximation to the ideal solution as long as the RVIRT does not change significantly over N/M technology projects.

There is an additional constraint put on the optimization to increase its efficiency. Up to this point a constraint on the total number of projects is put on the optimization. Some of these choices of projects may not make sense because of the project-project dependencies. If the projects upon which a project is dependent are not chosen, that project should not be chosen. An additional set of constraints is put on the optimization to guarantee this.

What would be presented to senior E&P management in order to determine the funding level would be a creaming curve showing the M options (EMV vs. risked Shell investment, including P20, P50, and P80 values and RVIRT). The senior management would then decide between these options largely decided by when the RVIRT for the next N/M projects equals that for alternative uses for the capital such as new field developments. After the base funding level is established by senior management, a detailed list of all projects is generated along with the marginal RVIRT for each project. A detailed look is taken at each project, especially those near the funding cutoff. A manual revision can be made by the research management to the portfolio about this base case, as was done for the original optimization. The revised portfolio should be evaluated to make sure that some unwanted effects of the dependencies are not caused by the portfolio revisions.

This process has been piloted using a BTM constructed for Syria. The projects are those prepared for the BAA 2000. The results are shown in the following series of figures.

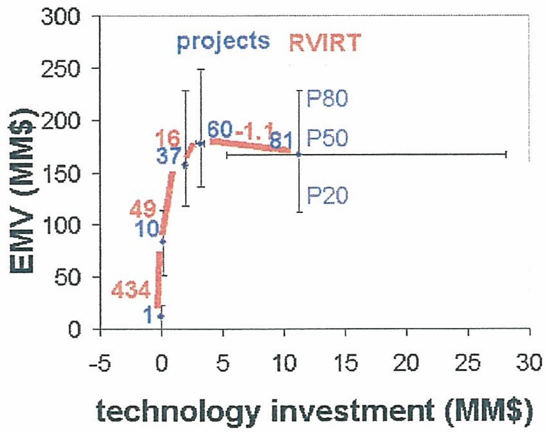

Figure 3. Creaming curve showing expected monetary value (risked Shell incremental value) vs. risked Shell investment. The RVIRT is shown in red. The number of projects in the portfolio is shown in blue. The P20, P50 and P80 values are shown.

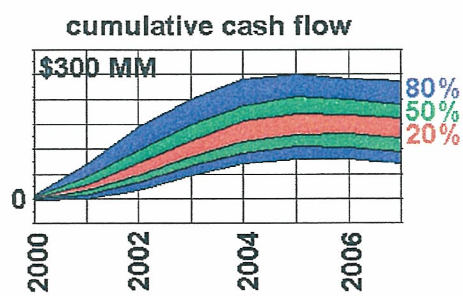

Figure 4. Cumulative case flow curves for the 81 project portfolio. The P20, P50 and P80 ranges are shown.

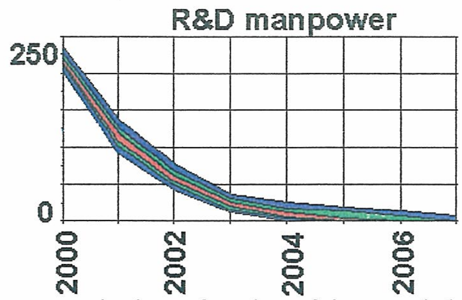

Figure 5. R&D manpower required as a function of time needed to pursue the 81 project portfolio. The P20, P50 and P80 ranges are shown. The small uncertainty is a result of a mature technology portfolio.

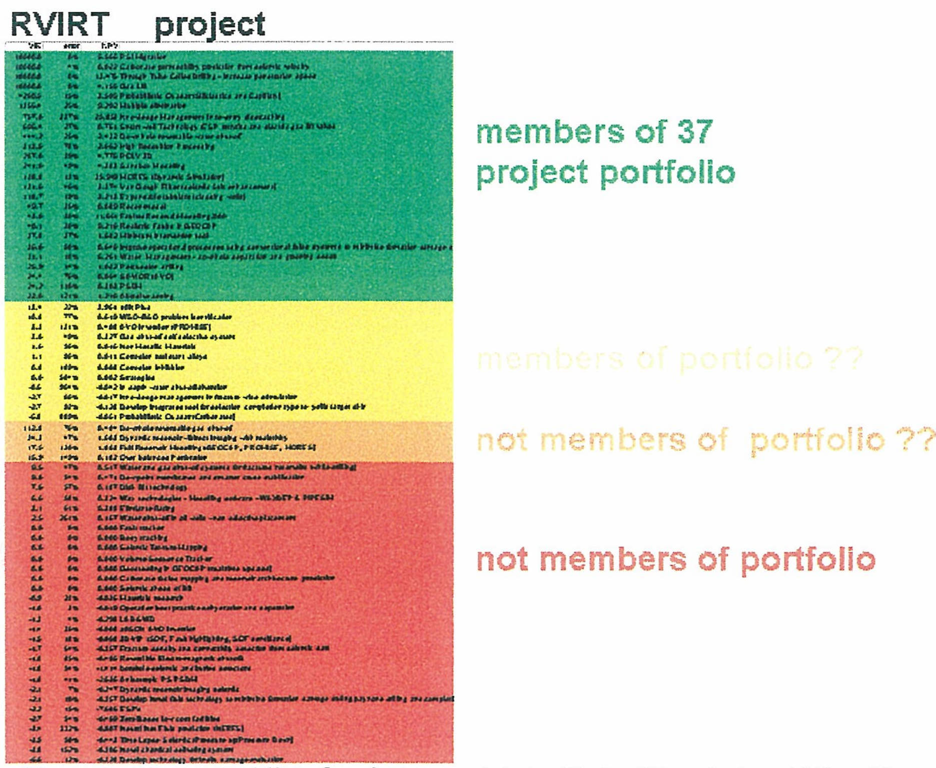

Figure 6. The ranked list of projects associated with the 37 project portfolio. They are ranked according to the marginal RVIRT. This ranking is done separately for the projects in the portfolio and those not in the portfolio. This is the list that would be presented to the research management after senior management had decided upon the 37 project portfolio.

The dependencies, both indirect (non-marginal contracts) and direct, have profound effects on the results. For the Syrian BTM pilot dataset an interesting effect of the contract structure is shown in Fig. 7. The Shell EMV increases as the price of oil decreases, while the government EMV decreases, as expected. This is a result of the PSA, which has a lower effective marginal government "tax rate" when our total profit is less. This is a hedge against a low oil price. The government is the party providing the hedge through the PSA. The explicit effect of the dependencies is shown in Fig. 8. Here the portfolio has been chosen taking into account dependencies and not taking into account dependencies. The economic value for both portfolios is calculated using the dependencies. The value of the portfolio for the same risked Shell investment is over 50% more if the dependencies are taken into account.

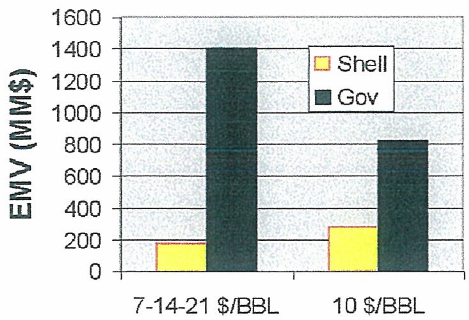

Figure 7. A plot demonstrating the effect of the PSA contract structure. The Shell and government EMV are shown for the 81 project portfolio, for the case of the triangular oil and gas price distribution and the case of a fixed $10/BBL case.

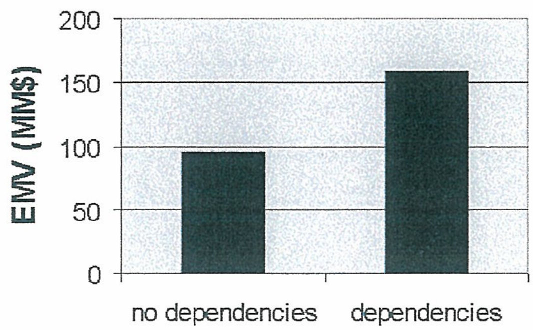

Figure 8. The EMV for a portfolio of 37 projects chosen neglecting the explicit dependencies, and for a portfolio of 37 projects chosen taking the explicit dependencies into account. The value for the no dependencies case is normalized so that the risked investment levels for the two cases are equivalent.

The economic models have been implemented in EXCEL using the Shell EFM functions. The stochastic evaluations have been done using the Crystal Ball EXCEL addin and some custom Visual Basic for Application (VBA) functions. The optimization has been done using OptQuest, an application that runs on top of Crystal Ball. The marginal tune-up is done with a custom VBA function. The additional set of constraints that honor the project-project dependencies is automatically done with another custom VBA function. The Crystal Ball Turbo product is used to distribute the calculation over 6 NT workstations with a corresponding 5 fold increase in speed. The optimization of an 80 project portfolio with one BTM takes 3 days on one 600 MHz Pentium III workstation. The time is expected to scale linearly with the number of BTMs. There will be a non-linear scaling with the number of projects leading to a practical limit of between 100 to 200 projects. If 5 optimizations need to be done using 20 BTMs and 40 workstations, the complete analysis would take one to two weeks.

Appendix 2. Structure and Implementation of Incentive Compensation

The incentive compensation scheme is divided into two parts. The first deals with the Research Committee members and the TMMs. Its emphasis is placed only on reaching the finish line – full implementation of a technology that results in significant business value. The second deals with the researchers, RTCs, and implementation personnel. This is a bit more complicated than the former and will be different for different stages of the nozzle.

The first group, TMMs and the Research Committee, are the people who monitor the total system and control it. The behavior that is desired of the system is a maximization of the EMV. The straightforward way to incentive this group is to give them bonuses whenever a technology reaches full implementation and a technical assurance process has established its realized business value and costs. They should receive a fraction of the realized incremental Shell value.

The second group, researcher, RTCs, and implementation personnel should receive compensation based on crossing the next tollgate. The payout of this compensation should not all be done when the next tollgate is reached, since the ultimate goal is not to pass the next tollgate, it is to realize full business value by passing all of the tollgates. Therefore, a bonus of 50% will be given when the project reaches full implementation. The amount of the payout will be proportional to time / Ps; where time is the estimated time to reach this tollgate and Ps is the estimated probability of reaching this tollgate. The two factors have very specific goals. The time factor removes the bias to work on project with an estimated short time to pass the next tollgate. The Ps factor removes the bias to work on projects with a large Ps. An estimation of the time and Ps will have to be made at the first tollgate. Historical data will be a good basis on which to make these estimations. Idea generation should be rewarded on submission of an accepted idea. There should be a fixed amount awarded per accepted idea. The implementation personnel can be rewarded on a more individual basis via a commission. For each new application of a technology the implementation personnel will be rewarded. The magnitude of the compensation will be based on the two factor expression divided by the number of new applications projected per implementation person.

There will multipliers in front of each of the compensation formulas. These multipliers will determine the amount of incentive compensation drive given the system. It is recommended that these multipliers start out small and be gradually increased. Incentive compensation is a powerful driving force acting on a complicated system. Not all of the effects can be foreseen. Such a gradual increase in the drive will allow for changes to be made to correct the unwanted effects.

Appendix 3. Incorporation of Strategy into the Financial Model

A significant effort has been expended to create a ranking tool for E&P capital allocation that will help achieve the aspired portfolio. The change in portfolio needed to achieve this portfolio is not taken into account in our calculation of the value of technology. We calculate the value if the technology is applied to our current asset base, not the future asset base to which the technology will be applied.

To formulate a solution to this problem, we have re-examined the portfolio matrix developed by the central office economists. An alternative approach is proposed that is equivalent to the portfolio matrix, but can be used when the value of a technology is dependent on the other technologies that are pursued. The portfolio matrix approach places the projects in a two dimensional space where the axes are project attractiveness and strategic imperative. The matrix is then divided into nine sections and the projects ranked according to in what section the project is located. The RVIRT optimization takes into account the elements of financial attractiveness. It does not take into strategic imperative. The use of the matrix approach assumes that the project attractiveness is independent of the other projects that are done. Unfortunately, this is not the case with technology.

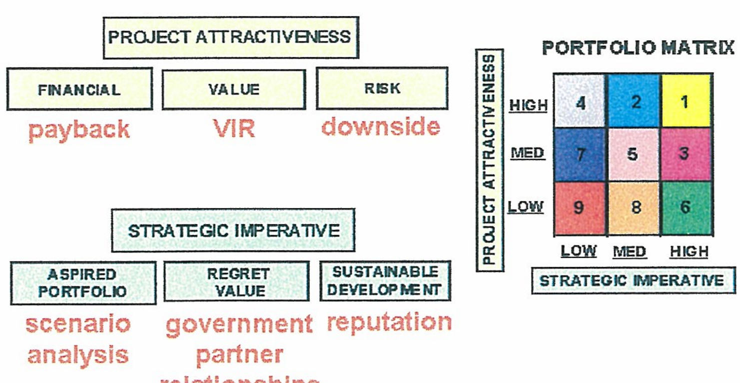

Figure 9. E&P portfolio matrix tool used for the year 2000 capital allocation process. (courtesy of Miles Biggins, EPT-B)

The proposed alternate approach is to consider the strategic imperative as a buy-sell-hold recommendation with an expected factional change in the number of projects that have been rated a sell (decrease) and a buy (increase) over the next 15 years (average field lifetime). What needs to be determined are the fractional changes. To this end, we now examine the ranking of the sections of portfolio matrix. If it is assumed that a project that has a high financial attractiveness will have twice the value of a project that has a medium attractiveness and a project that has a low financial attractiveness will have half the value of a project which has a medium financial attractiveness, it can be shown that we will increase our holdings by between 35% to 50% for projects in the buy bin and decrease projects in the sell bin by the same amount. This can be shown by correcting the EMV for the future size of the portfolio and examining the resulting ordering of the sections on the portfolio matrix. The correct ordering is only obtained for fractional changes in the specified range.

This approach can now be used to modify our ranking of technology projects. Each field or set of fields covered by a BTM will be classified as having low medium or high strategic imperative according to the same criteria used in E&P capital allocation process. The appropriate multiplier will then be applied to the net present value realized for the set of fields covered by that BTM. The constrained optimization can then be done as before.

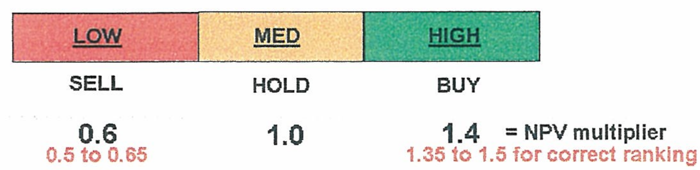

Figure 10. Interpretation of strategic imperative as a NPV multiplier.

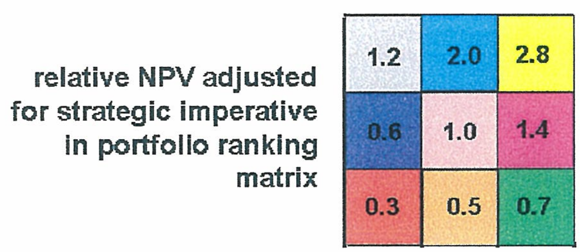

Figure 11. The relative value of the sections of the E&P strategy matrix using strategic imperative interpretation of Fig. 10. The horizontal axis is strategic imperative and the vertical axis is financial attractiveness. Note that the ordering is the same as that in Fig. 9.

There is the possibility for a broader generalization and use of the methodology developed in this report in the E&P capital allocation process. We have concentrated up to this point on technology alone. We have only considered the E&P capital allocation process as an external process that requires, in order to maintain consistency, a modification of the net present value in our calculations. One can take a more general view and cover, not only the management of technology development, but the whole E&P capital allocation process including technology development, and new field developments. In this model, one not only needs to optimize the EMV over which technology projects to pursue, but also which new field developments to pursue. With this formalism one can now explicitly take into account dependencies amongst the new field developments (e.g., a field development with high financial attractiveness which is dependant on the infrastructure of a field development with low financial attractiveness), and dependencies between the new field developments and technology development (e.g., a field development which is not financially attractive with existing technology, but which would become financially attractive with the application of new technology). What would be optimized is total EMV (modified for strategic imperative) under financial constraints (e.g., maximum payback) and constraints on downside risk (e.g., minimum P20 incremental value). This would consider all of the factors included in the current portfolio matrix ranking tool, reduce to same result as the current portfolio matrix ranking tool if dependencies are neglected, but be able to take those dependencies into account.

Appendix 4. An Example of Technical Assurance

This appendix gives a concrete example of how technical assurance works. The example that we will use is that of PROMISE, an advanced seismic inversion tool, which is nearing completion of the implementation phase.

It took one year to successfully complete the proof of concept stage with a cost of $0.3 million, and 2 years to successfully complete the prototype stage with a cost of $2.4 million. An additional $2.7 million has already been spent in the implementation stage with an additional $1.8 budgeted for the year 2000. It is projected that it will reach full implementation after 4 years at the end of 2000. This information is an example of the historical data needed for the turbojet engine feedback system. For future projects managed according to the model being presented, these delivery times and costs will be compared to the predicted times and costs to determine the value of using the predicted times over the historical averages.

Now, the value side of the equation is examined. PROMISE has been applied to more than 2 dozen fields. Of these fields two have been studied in detail (interviews with the asset teams after the completion of the project) to determine the marginal business value of applying the technology. The values calculated from these interviews were also reviewed and certified by the asset teams.

The first field is Ram/Powell in the Gulf of Mexico. It is a deepwater turbidite field 90 miles south of Mobil Bay. It is in water depths of 2400-4000 ft., with estimated reserves of 250 MMBEQ. It is a roughly equal partnership with Exxon and Amoco that Shell operates. First production started in 1997 from a tension leg platform. An AVO PROMISE inversion was done in 1998. Most of the business value identified was in reference to the N sand, a channelized turbidite that had shown some surprising water updip of oil. This reservoir was being produced from two wells at the time that the inversion was done. A final development well was scheduled to be drilled in the spring of 1999. The PROMISE results showed a complex pattern of oil that was consistent with all well control (i.e., oil saturations, net sand, mass balance, and pressure transient analysis). The results showed that the location for the additional development well preferred by our partners had almost no probability of encountering significant hydrocarbons. We were not able to persuade them to move the location of the pilot hole based on these results, but we were able to convince them to deviate the well downdip if the pilot hole did not encounter significant hydrocarbons. The pilot hole encountered 17 ft of net sand with 5 ft of oil. The downdip deviation encountered 83 ft of net sand, pay to base. It is estimated that this well will add $5 million of marginal value to the project through accelerated production. The downdip deviation would not have been drilled without the PROMISE result. If the PROMISE result would have been done before the development of the N sand, it is estimated by the asset team that an additional $20 million in CAPEX would have been saved by not drilling a well and preventing a sidetrack. There is also possible value in the development of two other sands, the M and Q, but it is too early to quantify the value.

The second field is Gannet E in the North Sea. It is a four dip closure of a distal turbidite, Forties formation. The first phase of development, consisting of a single horizontal producer, has already been completed. We are currently defining the second phase of development with our partner, Esso. There is significant uncertainty in the N/G of the sand and Esso is requiring that the second phase of development be economic at the P90 STOIIP. This uncertainty in N/G has historically lead to a 70% probability of sidetracking wells, at a cost of $3.4 million. The P50 marginal NPV for the second phase of development is $34 million. The delay cost is $0.85 million/month. Before a PROMISE inversion was done this spring there had been a 4 month delay in obtaining approval for the phase II development plan because the P90 STOIIP was not economic. The AVO PROMISE inversion has significantly reduced the uncertainty in N/G and lead to a P90 STOIIP that is well above the economic cutoff. It is estimated that this has at least decreased the time to gain approval from Esso by 3 months. There is also a 5% probability that Esso would not have approved the project without the PROMISE result. The reason that this probability is so low, is a need to fill facilities. The asset team has also estimated that the improved N/G maps have decreased the probability of sidetrack to 35%. The business value that was realized is $5 million. Another $3.5 million in increased value would have been realized if the PROMISE technology would have been applied earlier. The additional cost of doing the AVO PROMISE inversion has been $60,000 for each of these projects.

If a similar interview would be done for the other 20 or so fields, an average business value per field application would be obtained. This information would then become part of the historical database. This information could be compared to the technology target predicted at the start of the research to establish the reliability of those predictions. If the initial predictions were found to be unreliable, this information would form the basis for a historical value to be used in place of a predicted target, or be used to determine a correction factor that would be applied to the predicted targets.

Appendix 5. A Method of R&D Cost Allocation to the Business Units

The model also gives a method for distribution of the R&D costs to the business units. Because of the structure of the BTMs, not only can the NPV of a technology be summed across business units; the NPV obtained by a business unit can be summed across technologies. The cost of the R&D portfolio can now be distributed across business units so that each business unit will obtain the same VIR. This method of distributing costs will create a constructive tension in the system. Business units will want to assign a high value to research projects that they want to see funded, but if that value is set high they will increase the amount of the R&D cost that they will have to bear.

There is one problem with this system, it does not allow for a strategic redistribution of capital from mature business units to immature ones with growth potential. A method for allowing for this would be to separate the business units into three groups: growth, maintenance and exploitation. The maintenance group would be assigned costs according to the group wide R&D VIR. The costs would be assigned to the growth group at a VIR above the group average, and to the exploit group at a VIR that conserves the overall VIR.

References

(1) "Hot Groups", Harvard Business Review, July-August 1995, p. 109.

(2) "How SmithKline Beecham makes better resource-allocation decisions", Harvard Business Review, March-April 1998, p. 45.

(3) "Scientific management at Merck: An interview with CFO Judy Lewent", Harvard Business Review, January-February 1994, p. 89.