CREATING MORE EFFECTIVE RESEARCH AND DEVELOPMENT PORTFOLIOS[1]

Lawrence Phillips
London School of Economics and Political Science

A new approach to managing the research or development portfolio, which combines social processes with technical modelling, provides insight for difficult research and development decisions, gains buy-in from those affected and creates substantial added value through better use of available resources.

TWO KEY QUESTIONS are exercising the minds of pharmaceutical R&D directors at the moment: how to make better use of the available resources, and where to invest for growth. My experience of working with several major pharmaceutical companies over the past ten years shows that sound answers to those questions can be obtained by identifying the areas of greater and lesser opportunity in the portfolio, then allocating resources to the areas of higher priority. Finding the areas of greater and lesser opportunity can be accomplished by:

1. identifying and making explicit the *core values* underlying R&D decisions,
2. taking a *collective view of the portfolio* rather than seeing it as the sum of its parts,
3. generating information about the elements of the portfolio from *teams of key players*,
4. using *multi-attribute modelling* to bring together the values and information, and
5. applying a *socio-technical process* to examine trade-offs between areas, and to explore priorities under different assumptions about the balance of costs, benefits and risks.

These five issues form the outline of this paper.

Core values

Decisions about research and development are direct expressions of an organisation's core values. Making those values explicit is a first step to managing the portfolio. Spend a moment to think about what values your discovery research scientists and managers would say are important to them in guiding their work. Express the values as criteria against which decisions could be judged.

There are, of course, no right or wrong answers. Many pharmaceutical companies would see their Research Division as the engine that fuels the future of the company, so they would call up criteria such as those in the left column of the table (below).

---

[1] This is based partly on a presentation in London on 13 November 1997 at an IIR Conference, "Strategic Portfolio Management" and partly on "Value-for-money portfolio analysis" that appeared in *CMR News*, September 1995. I am grateful to Professor Stuart Walker for his support and encouragement of this work., though full responsibility for the content is mine.

Discovery Criteria	Late Development Criteria
• Potential for commercial return	• Commercial value, financial value, shareholder value, net present value
• Meeting an unmet medical need	• Meeting an unmet medical need
• Innovativeness	• Innovativeness, product superiority
• Patent position	• Protecting the base business and maintaining existing franchises
• Enhancing the science base	• Potential to evergreen or grow the product
• Number of candidates likely to be generated	• Establishing a new franchise
• Likely cost of developing the candidate	• Geographic extension
• Time to create a development candidate	• Time to market
• Discovery feasibility or doability in achieving a development candidate	• Risk, confidence, probability of registration

Ask this same question of people working on compounds in early development and the answer depends on whether sufficient resources are available to test every candidate received from discovery research. If so, then the only criterion is to reduce risk as quickly as possible, discarding unpromising candidates before they begin to eat up resources. If not, then the criteria will begin to look more like those used in late development, as shown in the right column of the table.

Comparing these two lists shows some overlap, but considerable differences. Even the first, relating to commercial attractiveness, is usually realised for late development candidates as a net present value figure, whereas for discovery work it is much too early to make any informed calculations.

But don't all these criteria eventually reduce to some measure of financial value? For example, doesn't an innovative compound that meets a medical need lead to financial return? Why, then, bother to consider these other values? I would like to propose two arguments for considering a mix of values.

The first begins by asking you to consider a hypothetical comparison for your company. Suppose you are comparing the further development of two compounds, A and B. Both will require exactly the same amount of development resource, and current estimates of their commercial value show their net present values are identical. However, compound A meets an unmet medical need to a greater extent than B. If resources were limited so that you had to prioritise the compounds in the portfolio, would you have any preference for A or B? If your answer is "yes", then medical need must contribute value to the organisation other than its financial value. Perhaps the value is in the satisfaction of knowing that the quality of people's lives is being improved, or in the image it creates for the organisation, or in the demonstration of being a scientifically-advanced company that can therefore attract top-rank scientists. Whatever the reason, if you have a preference for A over B, then medical need is contributing non-monetary value. It can contribute monetary value as well, and so the monetary and non-monetary values may be correlated, yet they are still separate values. Suppose, now, that A and B are identical for development cost, net present value *and* medical need. Are there other considerations that might lead you to prefer one over the other? By answering that question, particularly with reference to past decisions, you can begin to make explicit the non-monetary values that underlie decisions in your organisation.

My second argument starts with the observation that shareholder value has become the sole criterion pursued by many companies, usually because they are concerned about their status in the financial markets of the world. The single-minded pursuit of shareholder value is, of course, an active decision made at the top of the organisation. However, a case can be made that shareholder value cannot in the long run provide the driving force for an organisation, that the creation of shareholder value is the *result* of managing the organisation effectively in pursuit of criteria that reflect the core values. In short, shareholder value is the cart, not the horse.

Several years ago one of my PhD students examined the records of many strategy models I had developed for a number of organisations at different levels of management. In all cases options had been evaluated against a mix of hard and soft criteria. Hard criteria included profit, market share, revenue growth, etc., while soft criteria covered future potential, image, synergy, confidence, etc. He found that in weighting the relative importance of those hard and soft objectives middle managers placed about 60% of the total weight on the hard criteria, and 40% on soft. At the top of the organisation, those percentages reversed: 40% on hard and 60% on soft criteria. This led me to formulate the hypothesis that senior executives steer the organisation against the soft objectives, and they use the hard criteria to keep score on how well they are doing.

A recent study by Collins and Porras (1996) confirms this view. They compared 18 visionary US companies with 18 companies that were just successful or enduring, to discover how they differed. Merck was paired with Pfizer, Johnson & Johnson with Bristol-Myers Squibb, 3M with Norton, General Electric with Westinghouse, Hewlett-Packard with Texas Instruments, Walt Disney with Columbia, etc.

Their research found that:-

> "Contrary to business school doctrine, "maximizing shareholder wealth" or "profit maximization" has not been the dominant driving force or primary objective through the history of the visionary companies. Visionary companies pursue a cluster of objectives, of which making money is only one—and not necessarily the primary one." p. 8.

If not money, then what force aligns resource allocation in these visionary companies? Collins and Porras found that it is their core values or ideologies. Merck pursues advances in medical science and help to humankind; it gave away Mectizan, its drug that cures river blindness. Johnson and Johnson is explicit about its values in a Credo that recognises it is responsible to the users of its products, to its workers, to its managers, to the surrounding communities, and finally to its stockholders, to whom a "fair return" is promised. Among other things, 3M values innovation; General Electric, improving the quality of life through technology and innovation; Hewlett-Packard, technical contributions in their fields of involvement; Walt Disney, continuous progress. Such statements are either absent in the merely successful companies, or are not a living part of the culture of the organisations.

In short, identifying areas of greater and lesser opportunity in the portfolio requires not just a single monetary criterion against which to make that judgement, but several criteria, most of which will be considered 'soft'. These soft objectives are the *means* to achieving hard results, and it is the pursuit of these means objectives that motivates people, so is more likely to lead in the long run to the end objective of monetary return. To deal explicitly with soft criteria it is not necessary to develop models that translate outcomes on soft objectives into monetary ones. Later I will show how multi-attribute decision analysis can handle means objectives along with end objectives, as well as deal with uncertainty about the outcomes, all in the service of identifying the inequalities of opportunity in the portfolio.

A collective view

Why are inequalities of opportunity the key? The answer can be found in the 'commons dilemma', named after the overgrazing centuries ago in England caused by too many animals for the available pasture land, which was then held in common. It was in the best interest of each farmer to add more animals, for his individual profit more than offset the loss of grazing lands, but collectively this worked to the disadvantage of all the farmers as the shared land was over-grazed.

Today, the same phenomenon is responsible for the decline in fish stocks around the world; each nation is fishing in its own best interests, with collectively-adverse consequences. The commons dilemma makes clear that individually optimal solutions are not collectively optimal. Applied to research or development projects, this implies that any approach which optimises decisions on an area-by-area, a compound-by-compound or a project-by-project basis is bound to create a sub-optimal portfolio.

What is the solution? The dilemma can be resolved by looking at trade-offs between areas. This can be done only if options are generated for each area, some requiring more resource than at present, and others requiring less. If each area presents only a single strategy or plan, then there is no scope for trade-offs, and senior managers can do little more than tinker with the areas, squeezing out a little here to fund modest growth there. But this approach is far from solving the commons dilemma, and has led to trickle funding in many R&D divisions when opportunities exceed resources. Only if heads are lifted above individual silos so the territory can be surveyed and the more promising patches identified, can resources be allocated to better effect. The key, here, is two-fold: creating options within areas and judging trade-offs between areas.

Teams

Creating options is better done by teams than individuals, particularly where the teams are composed of people representing many perspectives on the issues. A discovery research team might include the team leader, chemists, biologists, biochemists, technology specialists, disease state experts, franchise managers, clinicians and planners. Perspectives in a development team could represent the relevant scientific and technology disciplines, financial and commercial views, regulatory concerns, clinical experience, and so forth. The group can pool information, challenge each other's judgements, and carry out peer review on the spot. When the work of many teams is complete, then another group, composed of the team leaders and senior managers, can best judge the trade-offs between the areas.

The reason for working in multi-functional teams is simply this: the outputs of a group of people working together can be better than outputs created by even the best individual in the group. Many heads truly can be better than one, but only under certain conditions. Recent research shows that a group can out-perform its members if three conditions are met: impartial facilitation of the group process, modelling, and instant playback of model results so the group can learn and modify its work (Regan-Cirincione, 1994). These are all aspects of the socio-technical process I describe below.

Multi-attribute decision analysis

So far I have been arguing for consideration of multiple values and for many perspectives to be brought to bear on issues of portfolio prioritisation and resource allocation. Does this not complicate analysis of the portfolio? Since 1976 the answer has been 'no'; a theory and workable technology exist for dealing with such complexity. It is called multi-criteria decision analysis (MCDA), first described by Keeney and Raiffa (1976), though Goodwin and Wright (1998, Chapters 2 and 12) is a briefer introduction.

The key idea in MCDA is to break down a complex problem involving multiple objectives into four elements: objectives expressed as several criteria, options, evaluations of the options on the criteria, and weights to reflect the relative importance of the criteria. Applying simple weighted averaging mathematics allows single overall evaluations to be determined for the options, which shows their relative standing on a scale of preference.

For portfolio work in the pharmaceutical industry, the criteria are drawn from the core values of the organisation and from objectives set by the research and development divisions. In this way, the process links to corporate objectives. One criterion that is common to all pharmaceutical companies is some indication of monetary value. In discovery research, this is often expressed as the potential for financial return. For late development, including additional development of marketed products, the criterion is usually net present value. Additional criteria cover medical need and the probability of reaching a clearly-defined goal. No more than six criteria are usually necessary.

Options are defined strategically, as statements of what could be done in an area and why, rather than how to do it and by when. For discovery research, the areas are usually defined as therapeutic areas, and the options map out whole sequences of activities, assuming all are successful, until a development candidate is produced. Additional areas might cover new technologies and feasibility studies. For development, where the goal is registration of a compound, the areas are compounds or projects, perhaps with the addition of areas covering technology, studies or other major uses of development resource.

Evaluations of the options on the criteria, whether they are based on net present value calculations, other numerical measures or directly-expressed judgements, are all converted to value scales, the common metric used in MCDA. These might be interval scales (where the zero-point and unit of measurement are arbitrary, as in Fahrenheit or Celsius temperature) or ratio scales (only the unit of measurement is arbitrary, as in length or mass), or any mixture of the two. Thus, both hard data, where it is available, and judgement can be accommodated in MCDA.

Finally, weights on the criteria usually reflect higher-level concerns about the portfolio. For example, gaps in the development pipeline might lead to more weight being placed on the discovery research criterion of time to create a development candidate. Thus, weights solve the organisational issue that not all criteria are equally important for all areas, while also addressing the technical issue of equating the units of measurement from one scale to the next to make all scales comparable.

The model puts all this information together, thereby using the teams' judgements and assessments, to construct the best portfolios of strategic options for a given level of total resource. The model used is implemented with the Equity software, developed at the London School of Economics in the 1980s for prioritisation, resource allocation and budgeting (ELSE, 2002). Inputs to the model consist of separate multi-attribute models, of the form described by Edwards and Barron (1994), for each area, and the trade-off weights. Avoiding the complexity of decision trees, yet accomplishing the same result, the MCDA model is simple but not simplistic, transparent and easily understood.

The basis for prioritising the options is deceptively simple, though it could be carried out differently for the development and discovery portfolios. For development portfolios, the usual approach is to calculate the total cost and total risk-adjusted benefit associated with each option, determine the ratio of the benefit to the cost, and orders the options from the one with the largest ratio to the one with the smallest. This value-for-money ratio establishes, with a single figure, the overall priority of each option. For the discovery portfolio, it might be more appropriate to calculate the ratio of the overall benefit to the risk, where risk may itself be a complex combination of several factors.

The most important output of this modelling is shown in the figure below, illustrated for one pharmaceutical company's development portfolio. The horizontal axis shows increasing total resource, in this case forward research spend, for all projects across 21 compounds. The vertical axis shows risk-adjusted total benefit, where benefit is a combination of several criteria. Any point on the curve represents the best portfolio for that level of total resource, where a portfolio consists of a collection of projects drawn from across the compounds, shown by all points to the left and down. The portfolio of projects that are currently underway is shown at point P; that portfolio is clearly not best, an example of the commons dilemma. A better portfolio, costing about the same but showing greater benefit, is indicated at point B. Equity showed that to move from the current portfolio to the better one would involve cutting six currently-funded projects and using the resource to fund 45 new projects. The benefit lost from dropping the six projects would be recovered and added to by the 45 projects.

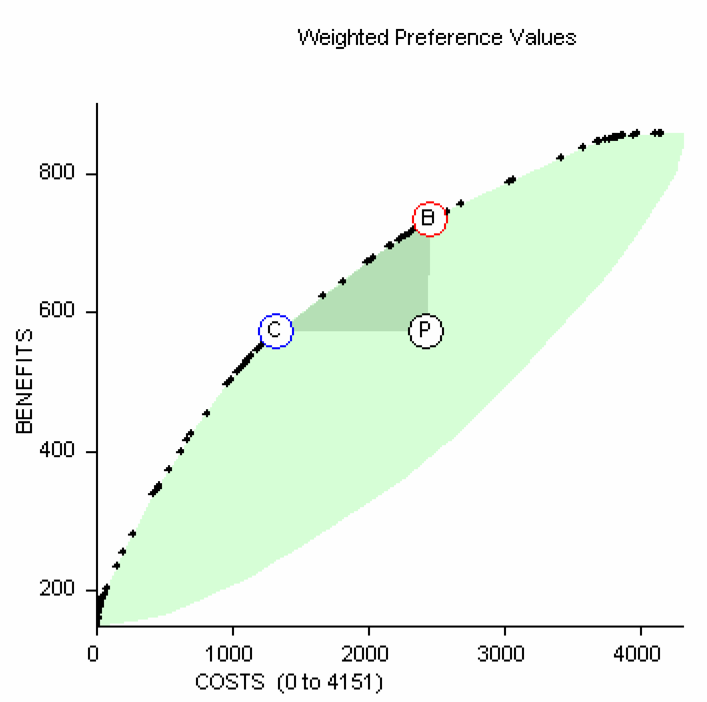

The increase in expected net present value by moving from the P to B position for this case was more than $2 billion in expected (probability weighted) net present value. *Experience shows that the difference between the P and B portfolios is on average 30%, and can be as large for a major pharmaceutical company as the equivalent of a blockbuster drug!*

The computer also moved horizontally to the left of point P to find a cheaper portfolio, C, at no loss of benefit, but this company was not interested in investing less. Note that any point on the curve between C and B is expected to yield more benefit at less cost than the current portfolio.

The MCDA approach is much to be preferred to the scoring systems that many R&D organisations have invented to prioritise projects and programmes of work. The main problem with these home-grown systems is that they are theoretically unsound, so they usually fall into disuse as their failings become evident. There are many ways to evaluate projects and programmes, but only a few ways to do it correctly. This is an area where consulting experts is essential. Home grown evaluation systems created by chemists, biologists and other scientists are no more likely to work than a drug developed by a specialist in decision analysis.

A socio-technical process

These four elements, values, a collective view, teams and MCDA, all come together in an approach that is both social and technical. The social element brings appropriate multi-functional teams of people together to work on that part of the portfolio which is relevant to their expertise. A Steering Committee, which includes the R&D Director, provides overall strategic guidance for the prioritisation process. A Project Group implements the process. Working Teams for discovery research are usually formed around therapeutic areas, and around projects for development, though other disciplines, such as technology, might be included.

Implementing this socio-technical approach starts with briefings of senior executives to explain the approach, and to engage them in enough hands-on modelling for them to see how the technical aspects work, to experience the social benefit from working in teams, to identify the cost, benefit and risk criteria, and to achieve their commitment to the approach. Next, the bounds of the target portfolio are defined, and Working Team members for each therapeutic area or compound are identified, with careful attention to ensuring breadth of perspective on all aspects of the therapeutic area or compound through its discovery, development, registration, marketing and post-marketing phases.

A Kick-off meeting then launches the process, shows the team leaders what their responsibilities are and what preparation is required, and engages the whole group in further development of the cost, benefit and risk criteria. Following the Kick-off, separate meetings are held with each team to generate the options and assess them against the objectives. The teams create the separate MCDA models that form one input to Equity. The completed work of the teams is then subjected to peer review by senior managers and, where possible, by outside, impartial experts. Revisions then go forward to a Merge meeting, where trade-offs between the areas are assessed, and relative weights to the criteria are assigned. Here the role of senior staff becomes particularly important, for they are in a better position than the project teams to see the 'big picture' and to recognise the differences between areas. This weighting process completes the technical work, and provides sufficient information for Equity to combine the work of the teams and identify the best combinations of strategic options at different levels of total resource.

Numerous sensitivity analyses are conducted, allowing participants to explore the priorities under different assumptions about the relative importance of the criteria and of the areas themselves, and to test the effects of differences in judgement. This process shows that many differences of opinion and much inadequacy in the available data makes no difference whatsoever to priorities about the strategies. All this is done on-the-spot in the Merge meeting. Computer output is projected so everyone can see the results immediately, and any unease about the results is explored to determine why. The purpose of the explorations is to help the group develop a sense of common purpose about the way forward. After the Merge meeting, time is required for reflection, leading to refinements that capture new insights about the portfolio. Decision makers meet, using this work to inform decisions about revising priorities, agreeing the composition of the portfolio and allocating resources.

The process is designed to ensure active participation of all those involved. All meetings are working meetings, conducted in the "here and now"; prepared presentations and fixed agendas are kept to a minimum. Team meetings and the Merge meeting are facilitated by outside specialists in decision technology and group processes. They bring an impartial view, which helps the groups to maintain a strategic perspective on their work and to stay focused on the task at hand (Phillips and Phillips, 1993). People are encouraged to express their views, however contrary they may seem, and to test the emerging group understanding against their own experience and intuition. The sense of unease expressed by individuals helps to fuel discussion and to stimulate creativity. New ideas and thinking "out of the box" are valued. With the exception of net present value calculations, no back-room modelling is done—all computer modelling is conducted on-line, live, as the group works, with no "black box" calculations. Opportunity is given to revise and improve the teams' work before the merge meeting. As I mentioned earlier, the goal of this process is to create a shared understanding of the portfolio without requiring consensus, a sense of common purpose about the areas of greater and lesser opportunity that leaves scope for disagreement, and a commitment to the way forward that encourages individual initiative.

The process, which is known to decision analysts as *decision conferencing* (Phillips, 1989), goes a long way to resolving the tension between individual and organisational views about the best way forward. Research at the State University of New York at Albany has shown that the approach is particularly valued in organisations that are open and receptive to change (McCartt & Rohbaugh,1995).

They have also shown that decision conferencing helps to resolve the competing values implicit in people's differing views about effective decision making, i.e., whether a decision is based on participation of all interested parties, or the result of a logical, rational process, or based on a thorough analysis of all relevant data, or the result of a flexible creative process (McCartt and Rohrbaugh, 1989). Decision conferencing can serve all these values.

Conclusions

One major pharmaceutical company observes that socio-technical analysis of the portfolio has brought about substantial improvement in communications between divisions that historically had little to do with each other. Even more importantly, they report that the approach has had a profound effect on the culture of the company, providing a new way to work together to common objectives. Another company sees the socio-technical approach as a means of encouraging teams to get out of their individual 'silos', to work with greater awareness of the larger picture. Senior managers are enthusiastic about the broad scope of options considered for all the areas; for the first time they gain a holistic view of the possibilities inherent in the portfolio. Companies with more limited resources are enabled to focus, stop the trickle funding and make more efficient use of the available resources. Companies that are growing their portfolios report that the process gives them clear guidance on how to justify additional spend to top management. Most importantly, decisions made by senior management about the portfolio will be implemented quickly and with the understanding of those affected, for they have contributed to the process and so understand the reasoning behind the decisions. For pharmaceutical companies wishing to deepen understanding of their research and development portfolios, and to develop new ways of looking at prioritisation and resource allocation, tackling the commons dilemma head on can release substantial hidden value in the current portfolio and show the way forward for growth.

References

Collins, J.C., & Porras, J.I. (1996). *Built to Last: Successful Habits of Visionary Companies*. London: Century Limited.

Edwards, W., & Barron, F.H. (1994). SMARTS and SMARTER: Improved simple methods for multi-attribute utility measurement. *Organizational Behaviour and Human Decision Processes, 60*, 306-325.

Goodwin, P., & Wright, G. (1998). *Decision Analysis for Management Judgment, 2nd Edition*. Chichester: John Wiley.

Keeney, R.L., & Raiffa, H. (1976). *Decisions With Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley.

ELSE (2002). Equity for Windows. London: Enterprise LSE, Houghton Street, London WC2A 2AE, telephone +44 (0)207 955 7128.

McCartt, A.T., & Rohrbaugh, J. (1989). Evaluating group decision support system effectiveness: A performance study of decision conferencing. *Decision Support Systems, 5*, 243-253.

McCartt, A.T., & Rohrbaugh, J. (1995). Managerial openness to change and the introduction of GDSS: Explaining initial success and failure in decision conferencing. *Organization Science, 6*(5), 569 -584.

Phillips, L.D. (1989). People-centred group decision support. In G. Doukidis, F. Land, & G. Miller (Eds.), *Knowledge-based Management Support Systems*. Chichester: Ellis Horwood.

Phillips, L.D., & Phillips, M.C. (1993). Facilitated work groups: Theory and practice. *Journal of the Operational Research Society, 44*(6), 533-49.

Regan-Cirincione, P. (1994). Improving the accuracy of group judgment: A process intervention combining group facilitation, social judgment analysis, and information technology. *Organizational Behaviour and Human Decision Processes, 58*, 246-70.