

Chapter 19: Decision Conferencing

Lawrence D. Phillips

Operational Research Group, Department of Management
London School of Economics & Political Science
Houghton Street, London WC2A 2AE

First published in Great Britain in 2006
by the Operational Research Group, Department of Management
London School of Economics and Political Science

Copyright © The London School of Economics and Political Science, 2006

The contributors have asserted their moral rights.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior permission in writing of the publisher, nor be circulated in any form of binding or cover other than that in which it is published.

Typeset, printed and bound by:
The London School of Economics and Political Science
Houghton Street
London WC2A 2AE

Working Paper No: LSEOR 06.85

ISBN: 07530 2077 7

Chapter 19: Decision Conferencing

Lawrence D. Phillips

Operational Research Group, London School of Economics and Political Science

Abstract

This chapter presents the current status of the decision conference process, a way of helping a group of key players to resolve important issues in their organization by working together, under the guidance of an impartial facilitator, with the aid of a decision analysis model of participants' perspectives on the issues, developed on-the-spot over a period of two days. The facilitator serves as a process consultant, guiding the group through the stages of discussing the issues, developing a model and exploring the results, without contributing to the content of discussions. The model serves as a 'tool for thinking,' not as providing an optimal solution or 'the right answer.' Participants are encouraged to express their sense of unease at any stage in the process, for it is the discrepancy between model results and intuitive judgment that drives the dialectic in the group. Exploration generates new insights and stimulates creative thinking, resulting in changes to the model and to intuitions. As this process settles down, participants develop a shared understanding of the issues, generate a sense of common purpose, and gain commitment to the way forward. Two case studies illustrate a typical individual decision conference and how sustained engagement with a client, *decision conferencing*, can lead to committed alignment in a group. Research on decision conferences provides insights into why decision conferences work.

Key Words: decision conferences, decision conferencing, process consultancy, multi-criteria decision analysis, evaluation, prioritization, group processes, quality decisions, aligned commitment, requisite decision models, facilitation skills

Introduction

An unexpected event led to the development of decision conferences. The managing director of the Westinghouse Elevator Company brought 20 of his staff, rather than only the few key players expected, to a two-day 'contact meeting' at Decisions and Designs, Inc., in May 1979 to deal with issues about the design of a new factory. Cam Peterson, then the Technical Director of DDI, wanted to break away from the standard consultancy model, the 'doctor-patient' model described by Schein (1999), in which the consultant gathers information from the client, goes back to the office to analyze the problem and develops some answers, then returns to the client to sell the solution. Instead, the contact meeting brought together the few key people who knew about the problem, and after much discussion and exchange of views, they provided relevant data and judgments, which were input to a computer-based decision model, on-the-spot, displayed on large monitors. The central idea of the contact meeting was that good information and data are best obtained directly from the decision maker, and that perspectives shift and change as information is exchanged. However, this time a large group of key players showed up, which Peterson felt at the time was overkill. Still, in the role of an impartial facilitator and specialist in decision analysis, he led the group through many sensitivity analyses and changes to the model as participants deepened their understanding of the issues. At the end of two days, the group agreed a decision, which was implemented quickly.

Later follow-through revealed that the success of this contact meeting owed much to the alignment of the 20 participants created during the two days, though DDI remained skeptical that substantial issues could be handled satisfactorily in just two days. But subsequent experience showed it was indeed possible, largely because the information needed to resolve the issues already resided in the heads of the key players, not necessarily in printed papers and reports. In

addition, bringing the key players together, encouraging them to participate in the problem-solving process, led everyone to understand how results were obtained, even if not everyone was pleased with the outcome. The buy-in of the team to the results of the modeling was, according to the managing director, the key. And so, the 'decision conference' was born, with aligned commitment considered as important as a quality decision.

When I heard about this development early in 1981 I realized that my two major interests, decision analysis and group processes, could be merged. Within weeks, the first decision conference in Europe, facilitated by Clint Kelly, a director of DDI, was held for the UK's Post Office. I could see the great advantage of getting all the key players around the table, talking to each other, exchanging information, debating from different perspectives, arguing their value positions, and using the model as a neutral repository for all the information and value judgments. As the model fed back to the group the collective results of their inputs, it didn't argue or take a position. "Here are the logical results of what you have been telling me," it seemed to say, "and if you don't like them, then feel free to change whatever seems wrong to you, and I will give you the new results." I could see that attempting to play the model didn't work; a change that would make one part better, often resulted in unexpected consequences elsewhere.

By the end of the decision conference, I knew this was the direction for my future work. Over the years that followed, my colleagues and I at Brunel University, later at the London School of Economics, developed decision conferences. Facilitating decision conferences, training new facilitators and meeting annually with colleagues in the International Decision Conferencing Forum provided research data for my colleagues, particularly Stuart Wooler,

Patrick Humphreys and Mary Ann Phillips, and we collaborated in developing both theory and practice. This paper brings together those many strands.

The decision conference

Decision conferences work best when they are about ‘hot’ issues—real concerns of an organization that require resolution. They can be strategic issues that require work over months or years, or operational ones, including immediate crises. Decision conferences don’t work very well for issues that are merely ‘interesting,’ or ‘nice to consider,’ but lacking any sense of urgency for their resolution.

Sometimes the problem has been studied for several months, providing information and data that can be brought to the table. Inevitably, the actual information needed to resolve the issues is substantially less than has been provided by the studies, and typically important information has been missed, so some decision conferences are organized at the very start of further exploration, with sensitivity analyses used to reveal areas that could benefit from additional data. This initial decision conference often helps to provide a new frame for the issues, helping to highlight what information is relevant and what is not. The outputs of such an initial decision conference guide the subsequent data gathering, making the process more efficient than the unguided search for information that may or may not be helpful.

More than 25 years and thousands of decision conferences later, conducted by many facilitators in over 15 countries, the elements common to all decision conferences are clear: attendance by key players, impartial facilitation, on-the-spot modeling with continuous display of the developing model, and an interactive and iterative group process.

The key players are chosen to represent all the main perspectives on the issues. While it is helpful if the decision makers are present, that isn't always possible, particularly for decision conferences in the public sector where it is elected representatives who hold the authority for final decisions. In these cases, the task of the decision conference is to make recommendations to the decision makers, so it is important that the perspectives of the decision makers are represented in the conference even if they can't attend.

At the start of the decision conference, the facilitator establishes the neutrality of information, encouraging participants to speak openly and freely, and asks the group to respect the privileged nature of the discussion. The Chatham House rule applies (RIIA, 1927):

When a meeting, or part thereof, is held under the Chatham house rule, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed.

In decision conferences the output is a group product, which was shaped by participants, so individual attributions are inappropriate.

Impartial facilitation, the second element of a decision conference, refers to the separation of the subject-matter content from the group process. The facilitator attends to content, but does not contribute to it, for to do so compromises the facilitator's impartiality, making it more difficult to function properly as a process guide. Most leaders of groups recognize the difficulty of attending and contributing to content and process. In a decision conference, the appointed leader remains in that role, but is relieved from guiding the process. Thus, all participants contribute actively to content, while the facilitator intervenes as appropriate in the group process to ensure that the group remains task oriented and achieves its objectives.

Special training in facilitation helps budding facilitators to achieve a reasonable level of impartiality. It is very tempting for the facilitator to take on a leadership role, especially if the

appointed leader is weak, or if the facilitator is a specialist in the topic under discussion. But over and over again, facilitators tempted to ‘help the group out’ by contributing their expertise at some point in the decision conference, report that at the end of the discussion they find it very difficult to return to a stance of impartiality, mainly because the group now sees the facilitator as trying to steer the group in directions that are not those desired by participants. With training, facilitators can learn to use their expertise as the source of questions to the group, not as providing answers.

On-the-spot modeling with continuous display of the developing model, the third element of a decision conference, ensures that every word and number input into the software is seen by participants, who are free to discuss, modify and edit the inputs. In this way, the model is built in small, digestible steps, with explanations given only when they are needed, ensuring that the transparency of the model is maintained throughout the whole process of creating it. A sense of ownership develops. If successive decision conferences are scheduled for complex problems, changes to the model in the days between decision conferences should wait until the start of the next decision conference, so participants can then approve them or not, continuing the model-building process from where it left off at the last decision conference.

Another contributor to the sense of ownership is the room in which the decision conference is held. Two basic principles apply, whatever the size or configuration of the room: (1) everyone should be capable of direct eye-to-eye contact, and (2) all displays, flip charts and white boards should be readable by everyone. This means that participants should not be arranged in straight lines behind end-to-end rectangular tables, and visual aids equipment must not obscure lines of sight. For groups of six to fifteen participants, chairs arranged in \cap -shape, work better than \square -shape. Groups larger than 15 are better arranged around round tables, cabaret style, for with only slight shifts of position, anyone talking can be seen by everyone else. A good discussion of

possible arrangements is found in Hickling (1990). Drinks, not just coffee and tea, and other refreshments should be continuously available at the back of the room so the facilitator can call breaks as milestones are reached, without being constrained by the timetables of the organization providing the room.

The final element, interactive and iterative group process, means that as the modeling process proceeds, requiring participants to be clear about each model element, thinking is clarified, particularly as information is exchanged in the group. As new ideas emerge they are captured in the model. Then, as the model combines the information given to it, results may stimulate new perspectives, which may require modification and revision of the model. Thus, there is a reflexive interplay between the participants and the model, as if the model is another participant, but a neutral one, merely reflecting back to the group the collective results of the information that has been fed in.

Dissatisfaction on the part of the participants with elements of the model, or its results, drives the dialectic in the group, resulting in further changes to the model. Unexpected results typically emerge as the group struggles to create a model that both informs and is informed by their experience and understanding of the data. Eventually, the process settles down as participants' understanding of the issues deepens and grows, and as the model becomes increasingly consistent and realistic. Finally, the model becomes 'requisite,' just good enough for the group to agree the way forward.

This point is not always reached in a decision conference. It may be necessary for a short period of reflection on the results to take place, for participants to work through the implications of their new understanding of the issues. A follow-through meeting, often with only the decision

maker and a few key players is sufficient to explore any remaining issues, and for the decision maker to take final decisions.

Stages in a typical decision conference

The decision conference process begins with an inquiry from someone in the sponsoring organization who recognizes a need: possibly a gap between desired and actual performance, or a recognition that changes in the environment call for new ways of operating, or that current strategies or policies are losing their relevance to new conditions. Whatever the source, it is helpful to establish if there is a motivation for change, for without it, decision conferences are not likely to lead to commitments that people will implement. If it can be established that a hot issue really exists, and that a decision conference could deal with the issue, then the facilitator engages the client in the next stage, preparation, as shown in Figure 1.

At a short meeting, typically less than two hours, usually held with the decision maker, perhaps supplemented by key members of his or her staff, the facilitator explores the nature of the issues to ensure that decision conferencing can help, and works with the team to establish the objectives of the decision conference. Next, the key players who will attend are identified: people whose perspectives can make a useful contribution to the resolution of the issues. This may include people from other parts of the organization who are not necessarily stakeholders, and possibly outside experts if specialist information is lacking within the organization. The outlines of a calling note, to be sent by the decision maker or lead person, are established. These include the purpose of the meeting; administrative details of when and where; a paragraph explaining that the meeting will be conducted as a decision conference, with an attachment providing an introduction to decision conferencing; preparation asked of participants, usually

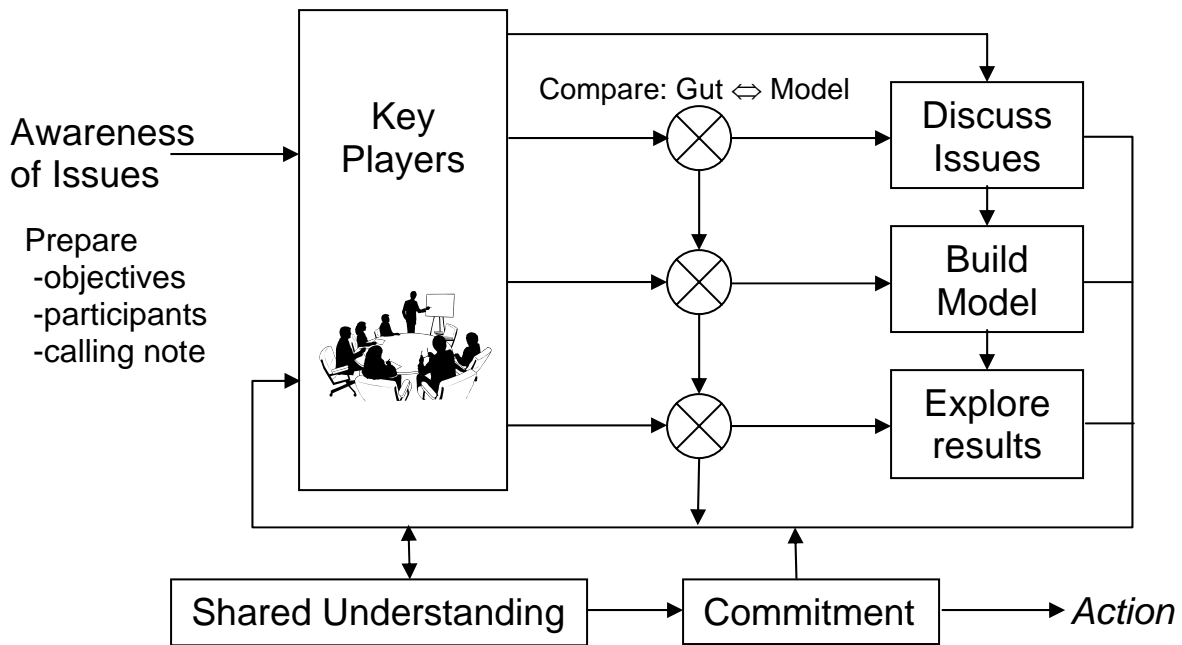


Figure 1: Schematic of the decision conference process.

little more than to think about possible options for the way forward; and a request to arrange diaries for uninterrupted attendance throughout the full two days.

The decision conference starts with a reiteration of the objectives, which participants are encouraged to discuss and possibly modify. Agreement to the objectives establishes the primary task of the meeting, grounds the group in a shared social reality, and provides the facilitator with a goal that legitimizes future interventions. The facilitator explains that information is to be shared freely, and treated as a neutral commodity. The three generic purposes of decision conferences (see next section) are explained, shared understanding, sense of common purpose and commitment to the way forward. To achieve these objectives, the facilitator explains, the group will create a model that captures the key elements we must address to resolve the issues. To do this, we will first discuss the issues, then build a model of the issues, and finally explore the results. Participants are encouraged to compare the outputs of the stages with their holistic judgment, their gut feelings, and to report any sense of unease, even if they don't know the

reason. Exploring the discrepancy may generate new insights, requiring changes to the model or to people's judgments, which help to build the group's shared understanding of the issues. Many sensitivity analyses will show the extent to which results are affected by imprecision in the input data and differences of opinion. As our understanding deepens, we will circle back to previous steps to make changes and revisions that reflect our new perspectives. As this process settles down, a sense of common purpose develops, and agreement about the way forward will most likely emerge. Then the work begins.

After the decision conference is over, the facilitator prepares a report of the meeting, and may hold a short follow-through meeting with the senior staff to resolve any remaining issues.

The purposes of a decision conference

In the early years, I thought that improving decisions was the goal of decision analysis, and this is how we positioned decision conferences in the UK. However, follow-through studies of decision conferences showed that recommendations from the decision conference were not always followed, yet participants valued the experience, particularly as compared to conventional meetings (Chun, 1992). Further questioning revealed three underlying reasons why decision conferences were valued: (1) they helped the group to generate a shared understanding of the issues, without requiring consensus about all issues, (2) they developed a sense of common purpose, while allowing individual differences in perspective, and (3) they gained commitment to the way forward, yet preserving individual paths. Senior managers often said their biggest problem was getting everyone to pull in the same direction, and that decision conferences helped to achieve agreement about the way forward, even if not everyone was agreed about the best decisions.

We found that the problem was not so much poor decisions as wrong decisions. For example, one group of sales managers from six Eastern European countries reported that the decision conference showed that the wrong strategy was being followed in each country. The strategies recommended by the end of the decision conference were not subsequently implemented because the realignment of the country managers' perspectives led them to grasp new opportunities when they returned to their countries. All the managers agreed that the decision conference had been worthwhile (Phillips, 1990). Indeed, profits and revenues that had been static for four years, more than doubled 18 months after the decision conference, and continued to rise steadily for many years afterwards. The managers attributed this improvement in performance to the new strategies.

Of the 80 organizations for whom I have engaged in consultancy work, only one came to me asking for help to improve decision-making in the organization. Perhaps this is not surprising, for as Francois de La Rochefoucauld observed, "Everyone complains of his memory, and nobody complains of his judgment." By the six standards of quality decisions (Matheson & Matheson, 1998), decision conferences lead to better decisions, but this is rarely mentioned by participants in debriefing. Of the six decision quality dimensions, only 'commitment to action' comes close, but even then it isn't so much 'action' that is mentioned by decision conference participants as 'commitment to the way forward.' This is a general orientation that can fall short of specific actions, more like the alignment of iron filings created by a magnet held below the paper on which they are scattered. Furthermore, a case can be made that because it is only individual managers who are held accountable for their decisions, not groups, allowing groups in organizations that are managerial accountability hierarchies to make decisions would undermine the authority of managers (Jaques, 1998).

Decision conferences are now positioned as helping managers to achieve committed alignment to the way forward. In managerial accountability hierarchies decision conferences make recommendations, not decisions. This distinguishes them from ordinary workshops, whose only purpose is to achieve a particular technical objective, such as the best design of a system, or an improved allocation of resources, or to choose the best of several alternatives. Decision conferences accomplish both technical and social objectives, so that is why they are positioned as a socio-technical approach to resolving issues of concern to an organization.

Evolution to *decision conferencing*

Although we often deal with large problems involving substantial resources within the two-day period, some problems, particularly those involving many separate groups of people with specialized knowledge from different parts of the organization, can't be handled adequately with just one decision conference. Then, a succession of interviews, workshops and decision conferences may be held, with the final decision conference bringing together all of the previous work.

We call this overall process of sustained working with a client *decision conferencing*, whose purposes are, again, both technical and social. The second case study in the next section provides an example, and issues concerning the design of an appropriate social process are discussed in the section below on process.

Case studies

To illustrate a decision conference and the decision conferencing process, two case studies are presented in this section. The first took place over just one day, an unusual application, and not recommended because it fails to take advantage of an overnight reflection that typically brings

fresh perspectives to the table at the beginning of the second day. The second case study required my involvement intensively over about two weeks, and periodically over three months. This second example is typical of sustained engagement with a client.

New business appraisal

A financial services company in the UK, part of a global organization, wanted to transform and evolve the world-wide Group's business by engaging in a new e-commerce activity. Uncertain about what business to pursue, they hired outside consultants to develop alternatives. The consultants narrowed down a list of 30 possibilities through a process of evaluation on several key criteria to a short list of three, here identified as SMB, Bank, and Benefits. Each showed promise on many criteria, but none was overall best in every respect. The managing director of the financial services company decided to call a one-day decision conference attended by his senior staff and the consultants, altogether 16 participants, to choose one of the three.

The meeting began with a presentation from the director of IT of the background to the project, followed by brief presentations from the consultants of the key features of the business alternatives. Each presentation provoked discussion as the group sought to clarify the issues. Participants then privately scored the three options on a scale of overall preference, with the most preferred option assigned a score of 100, the least preferred a zero, and the third option a score somewhere in between so that the differences between it and the least and most preferred options reflected their own differences in strength of preference. The majority, 10 participants including the four consultants, chose Benefits, while six chose SMB. Nobody chose Bank. Clearly, the group was not agreed about the best way forward.

The group then developed a multi-criteria decision analysis (MCDA) model in which the three options were appraised against two monetary cost criteria, five risk criteria and 11 benefit

criteria. Direct relative scaling was used for each criterion, including the cost criteria, though many of the scores were based on data provided by the consultants in a briefing pack provided to participants. The value tree took about two hours to develop, with the scoring another two hours. Swing weighting provided the scaling constants used to combine the scales. Lively discussions attended the scoring and weighting. On the few occasions when consensus could not be reached, the majority view was input to the model, with the disagreed figure the subject of later sensitivity analyses. The value tree in Figure 2 shows the structure of the model.

Many criteria were means objectives rather than fundamental objectives. It was necessary to include them because detailed financial models had not yet been built about any of the options. At this stage, the managing director felt that the group could make good relative judgments about the options, and if the MCDA showed one to be a clear winner, it would then be subject to closer financial scrutiny. Of course, most of the criteria in the value tree would be considered in creating the financial model; the MCDA model was simply considered another approach to obtaining an overall result. Some care was taken to ensure that the criteria were defined to be mutually preference independent.

Note that Risk was considered as one aspect of Cost; this is the way the group felt about risk, so each risk criterion was expressed as a 0-to-100 preference scale, with the option judged to be least risky on a criterion assigned a 100, and the most risky a zero. Participants were reminded that all criterion scales were relative scales, like Celsius or Fahrenheit temperature, and that zero represented least value, not no value.

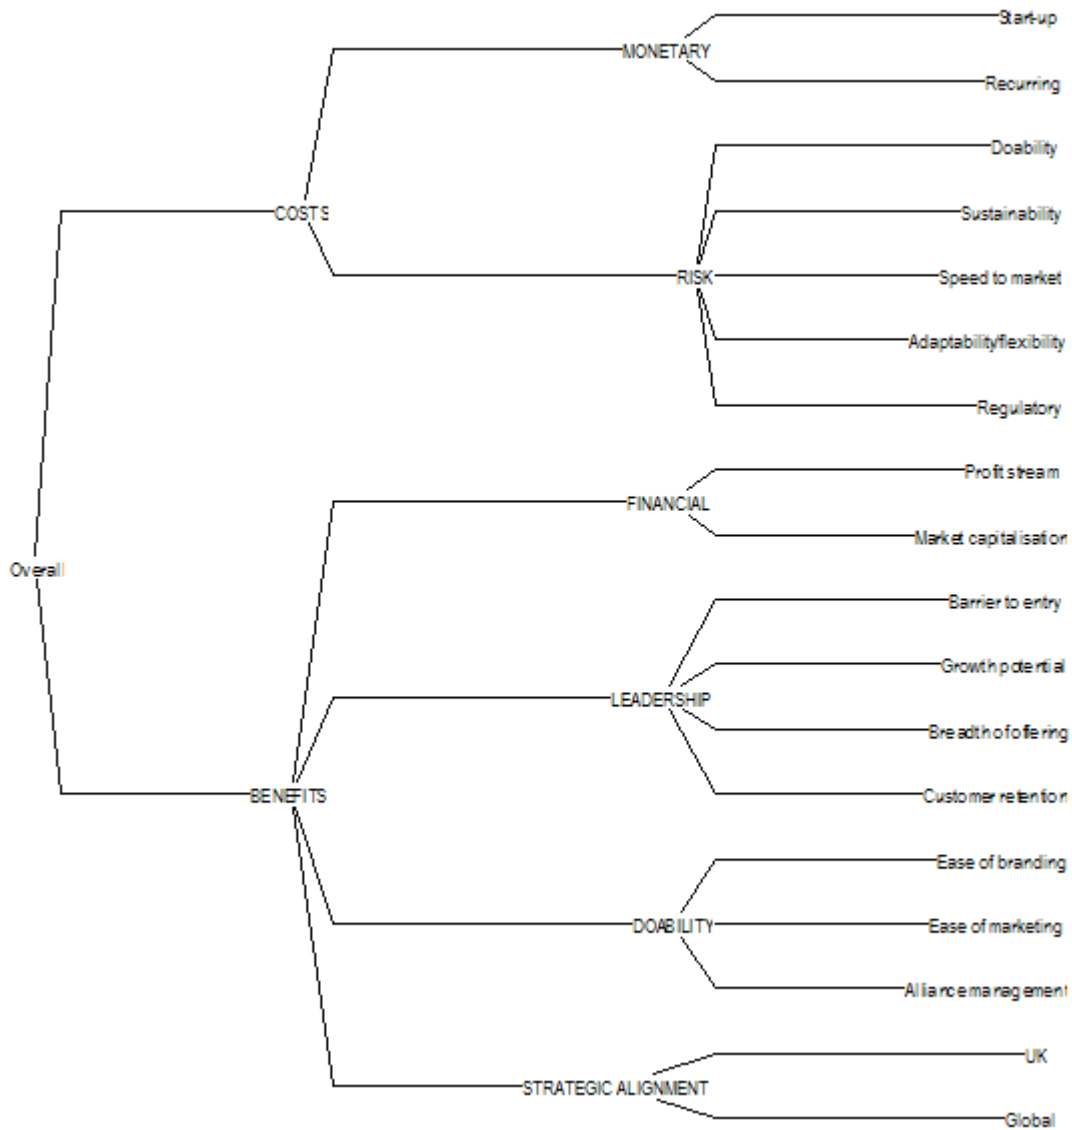


Figure 2: Value tree for the e-commerce new venture.

Overall results are shown in Figure 3 in two ways, as stacked bar graphs, and as overall benefits versus overall costs. (Note that the lower the Costs, the longer the upper portions of the stacked bar graphs because less cost was more preferred.) The left graph giving the overall results shows that while SMB is overall most preferred, that Bank and Benefits are very close to each other and considerably less attractive than SMB. Several of the 10 people whose holistic judgment had led them to choose Benefits challenged some of the scores and weights.



Figure 3: Weighted scores at the level of Costs and Benefits for the e-commerce new venture, and a plot of benefits versus costs for the three options.

A period of intense sensitivity analyses followed, mainly on the weights assigned to monetary costs and to risks. These were displayed graphically, in the usual way for MCDA, as overall benefits versus the weight on the criterion, resulting in three straight lines, one for each option. Further sensitivity analyses on the Financial, Leadership and Doability nodes confirmed the attractiveness of SMB. Overall, the dominance of SMB proved remarkably robust over large ranges of difference of opinion about weights, enabling the group to agree that SMB was the way forward. Additional comparisons showing those criteria that account for the advantages and disadvantages of each option showed why SMB was best, and so helped participants to understand why, contrary to the initial opinions of some, SMB emerged as the winner.

It is worth observing that the four consultants initially preferred the Benefits option. Their criteria were less complete than those shown in Figure 2. This illustrates a major difficulty with the ‘doctor-patient’ model of consultancy: the consultant’s reality does not necessarily correspond to the client’s. It is easy for ‘doctor-patient’ consultants to become overly prescriptive, failing to understand fully the client’s history, culture and preferences. Decision conferences work with the client’s beliefs and values, both understanding and developing them, a

constructive approach (Gergen, 1993) that allows the client to keep in touch with the problem, its formulation and results.

Prioritization of projects

The next case study concerns a major pharmaceutical company in the United States that used the decision conferencing process to prioritize projects in its late development portfolio. Decisions in many R&D organizations are made on an individual project basis, which inevitably results in an overall use of the available resource that is not collectively optimal, a condition akin to the well-known Commons Dilemma (Hardin, 1968). A better approach is to model the trade-offs between projects, which is effectively done by applying MCDA. It is then possible to find the best combinations of options for a given level of resource, creating a genuine portfolio that will realize more overall benefit as compared to the sum of individual 'silo' decisions (Phillips & Bana e Costa, 2006). An overview of the approach is given in chapter 13 of Goodwin and Wright (2003).

This type of assignment benefits from a three-tier approach. A Strategy Committee consisting of the R&D director and a few senior staff first meet with the facilitator(s) to establish the scope of the project: the purpose and context for the prioritization, the projects to be included and the criteria against which the project activities will be appraised. A Working Group, whose chair is one of the senior people on the Strategy Committee, is established to organize the entire process. Among other things, they identify who will be involved in each of the team meetings—not just the team working on each compound, but also the many other key players who can contribute to the prioritization workshops; internal staff representing regulatory affairs, marketing, finance, clinical, etc, and possibly an outside expert. They also identify those staff who will attend all meetings, usually the chair of the Working Group, one or two people from finance, and the

analyst who is familiar with the software and builds the model. A key role for these people is to ensure the realism and consistency of inputs from one team to the next. I often work ahead of time with the finance specialists to ensure that just the right level of financial modeling is carried out to support the teams.

The process began with a Kick-Off Meeting, attended by the R&D director, some of his senior staff, all members of the Working Group, the finance specialists, and two senior people from each team, the current Project Leader and Program Manager. The R&D director started the meeting by explaining why this *decision conferencing* approach is being taken, encouraging participants to afresh and creatively, and he explained the current overall strategy of the R&D Division. This was followed by an explanation from the facilitator of the modeling approach. The benefit criteria suggested by the Strategy Committee were explained to the group, followed by discussion, modification and clarification of the criteria. At this stage, the experience of the facilitator can be helpful, for most organizations are unaware of the attributes for good criteria (Keeney & Raiffa, 1976), nor are they aware of the principles of value-focused thinking (Keeney, 1992). This expertise of the facilitator has to be used carefully, questioning, not telling, to ensure that the criteria are fully owned by the group, and not imposed by the facilitator.

The facilitator warned the group that getting the criteria right the first year was unlikely. Criteria reflect the underlying core values and strategic intent of the organization, and even if prior thought has been given to them, using them in an MCDA always prompts changes to them. Getting them roughly right the first year is all that can be expected, and is sufficient for the three purposes of decision conferencing. The meeting ended with a question and answer session on how the subsequent team meetings will be conducted. The Kick-Off Meeting lasted about two hours.

It will be helpful at this point to show the structure of the model developed for this case; see Figure 4. Each of the 12 towers represents a different compound, here referred to as a project. The building blocks correspond to strategies associated with the compound. The term ‘strategy’ is used in the sense of ‘what and why,’ that is, what the allocated resource is used for, and why. At this stage, the ‘how and by when’ is not considered, though for every strategy there must, of course, be a ‘how.’ Strategies that are currently underway are indicated by a ‘P,’ whereas new strategies proposed to join the portfolio are shown by a ‘+.’ Many strategies are to develop drugs for specific indications (Ind), which require testing in humans. Some strategies are to conduct studies, to create combinations of compounds, to develop new formulations, even to open new markets. A complete statement of the strategy also includes an explanation of why it is attractive, such as ‘to develop a new oral formulation that would be more convenient for patients because it would be taken once a day instead of the current three times a day.’

Each project team met with the facilitator in a one-day or half-day workshop to develop the

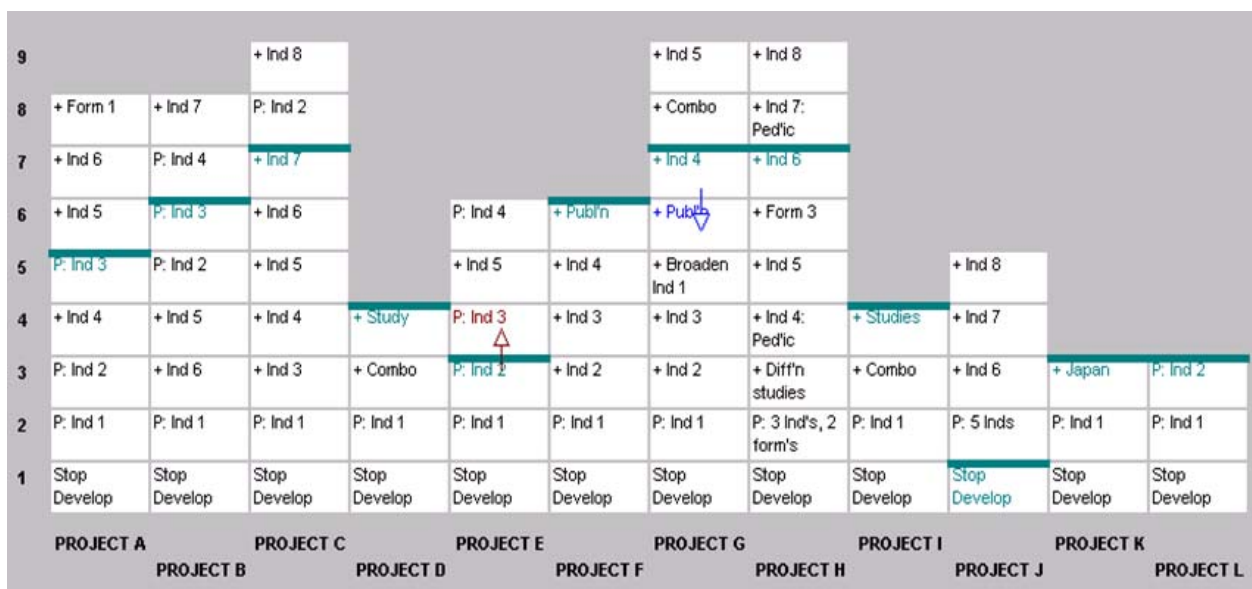


Figure 4: Structure of the portfolio model, with each tower a different compound, and the white boxes showing current Plan (P) and new (+) options. All strategies up to the bold lines define the affordable frontier portfolio.

model. The group defined the project and explained the current overall strategy for the compound, if there was one, along with a description of the currently funded strategies. The facilitator next engaged the group in a SWOT (Strengths, Weaknesses, Opportunities and Threats) analysis (Ascher & Nare, 1990), listing on flipcharts the internal strengths and weaknesses of the compound, the team and the organizational infrastructure, and the external opportunities and threats posed by the extent and seriousness of unmet medical need, untapped markets, developing technology, regulatory restrictions and market competition. The group then used the SWOT analysis to suggest new strategies that would build on the strengths, fix the weaknesses, grasp the opportunities and stave off the threats. Voting on the resulting list left a shorter list of new strategic options to be included in the prioritization model.

Next, the team scored all the current and new strategies, now referred to as options, against a set of cost and benefit criteria suggested by the Strategy Group and accepted at the Kick-Off Meeting. These included the following, which were defined for the teams in more detail: total cost, net present value, extent to which the option will meet unmet medical need, business impact, future potential and probability of technical success.

All input probabilities assessed by the teams were converted to subtractive penalty scores by a logarithmic mapping, resulting in a proper scoring rule (Bernardo & Smith, 1994). This somewhat unconventional procedure for turning an uncertainty into a risk criterion captures the intuitions of people in the pharmaceutical industry that improving the probability of success from, say, 0.10 to 0.20 is more valued than from 0.80 to 0.90. In addition, treating risk as a criterion makes it possible to do sensitivity analyses on the weight of this criterion, allowing risk to be more or less influential on the final result, another desired feature expressed by

participants. Of course, an entirely additive model enhances transparency of the results, and as an added bonus makes possible the use of software that only accommodates additive modeling.

Financial staff assisted in assessing total costs and net present values. Linear value functions transformed the NPVs into preference values. Preference values for unmet medical need, business impact and future potential were directly assessed by the group using ratio scaling techniques with balance-beam (Watson & Buede, 1987) consistency checks.

The group turned next to assessing swing weights on the five benefit criteria scales, thereby equating the units of value across the criteria (Clemen & Reilly, 2001; Goodwin & Wright, 2003). This important step circumvented judgments of ‘absolute importance’ for the criteria by posing questions that elicited value trade-offs.

With scoring and weighting completed, the computer program calculated a single risk-adjusted benefit and single forward cost for each strategy. Dividing the risk-adjusted benefit by the cost resulted in the priority index for each strategy, the recommended basis for resource allocation in cost-benefit analysis (HM Treasury, 2003) and corporate finance (Brealey, Myers, & Marcus, 1995). As I have noted elsewhere (Phillips & Bana e Costa, 2006), this simple criterion is rarely used by organizations, yet it is easy to demonstrate that any other basis for prioritization will not identify the best value-for-money portfolio.

The computer reordered the team’s strategies for each project in value-for-money order; that is the order shown in Figure 4, with value-for-money after level 1 declining up the tower. A plot of the cumulative risk-adjusted benefit versus cumulative total cost is shown in Figure 5. Note that two new strategies, Ind 4 and Ind 5 fall in highest-priority position, ahead of indications 1, 2 and 3, the current plan. This is a typical result when a team prioritizes for the first time on the

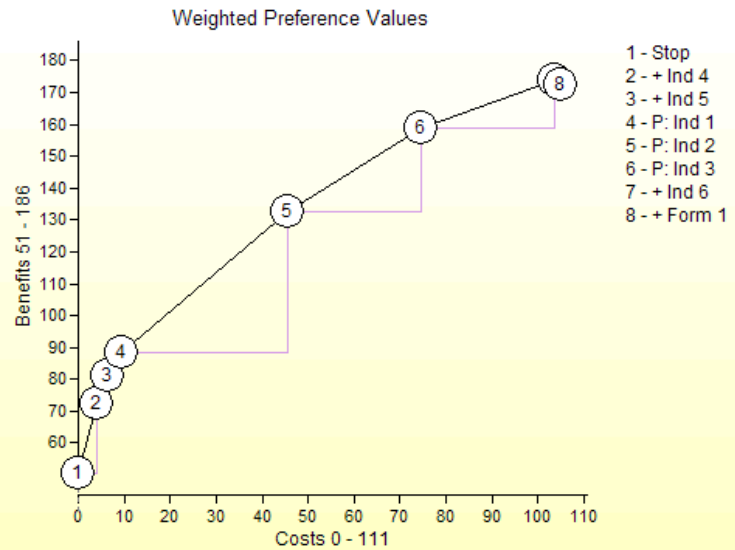


Figure 5: Cumulative risk-adjusted benefits versus total costs for Project A.

basis of value-for-money; current strategies are lower priority than things they aren't now doing, a finding which usually leads the team to pause and rethink their current strategies.

When the teams were satisfied with their work, it was submitted to a review panel of vice-presidents who checked the scores for realism and consistency. Any discrepancies were referred back to the team leader for revision. Final versions were collected into a briefing pack that was sent to all participants in the final Merge Meeting.

The main technical task for the Merge Meeting was to bring together all the work of the individual teams to form a single Order of Priority, or efficient frontier. To do this, participants engaged in a process of assessing swing weights, whose value trade-offs established scale constants between the projects and across the criteria, a double-weighting system that over-rode the across-criteria weights assessed by each team. The result is an efficient frontier, shown in Figure 6 with each point representing another strategy.

The shaded area shows the locus of all possible portfolios, formed by all combinations of the strategies, in this case, over 1 billion possibilities. Point P shows the current portfolio, while B

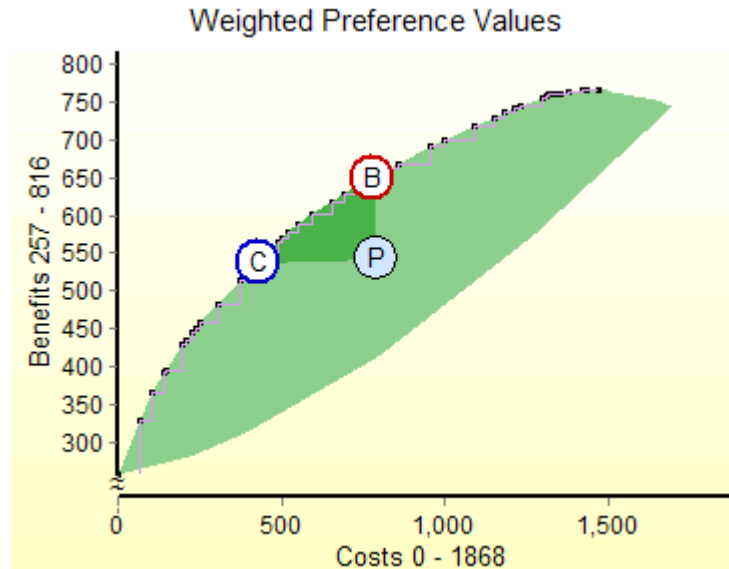


Figure 6: The final Order of Priority of all the strategies. P represents the overall risk-adjusted benefit and cost of the current portfolio. B is a better portfolio at the same cost, and C is a less costly portfolio at the same benefit.

and C are better and less costly portfolios, respectively. The darker shaded area defines the locus of all portfolios that are better than the current one. Participants were keen to see the composition of point B, which is shown in Figure 4 as all strategies up to the bold line for each Project. Note that closing down Project J was recommended. This led the team to rethink their approach of wrapping up five indications into one strategy. It would have been better to disaggregate this strategy and let the prioritization process select the best indications to take forward, as most of the other project teams had done. Some modifications to their approach were later agreed.

Another finding that surprised the group was the apparent imbalance in the portfolio: five current strategies fell outside the Better portfolio, and 27 new projects were included within it, a result that is not uncommon in larger organizations whose decisions are taken on a project-by-project basis rather than at the portfolio level, which requires trade-offs to be judged between projects. Trade-off analysis on the completed MCDA model brought this point home. Indication 4 for Project E was considered by many participants to be essential, that it couldn't be dropped. A

trade-off analysis brought Indication 4 into the Better portfolio, but then showed what strategies in the portfolio would have to be forgone in order to pay for Indication 4. Because that strategy was expensive, many other projects would have to be dropped, and this, too, was considered unacceptable by the group. Clearly, this result was such a substantial deviation from the current position that it could not be adopted; too many commitments would have to be broken to stop current low-priority strategies to enable resources to be transferred to new high-priority strategies.

Other organizations, when confronted with such discrepancies, recognize that the model provides a strategic direction from them to move toward over a period of time. It may take several years for the higher value-for-money strategies to be adopted as the lower-values ones are dropped or allowed to perish. This evolutionary approach to change can be very successful, as has been demonstrated by Allergan in the movement over several years of their current portfolio toward the efficient frontier (Phillips & Bana e Costa, 2006).

The reader may wonder why these two case studies show MCDA models and no decision trees or influence diagrams. The reason is that my clients over the past 25 years have brought issues that are more about the conflict of objectives than they were about risk and uncertainty. Perhaps this echoes, realistically or not, the finding of MacCrimmon and Wehrung (1986) in their study of over 500 American and Canadian managers:

“Rather than taking the chances of potential loss, magnitude of loss, and exposure as fixed, the managers tended to adjust the risky situations to make them more attractive. Choices tend to be made only when further adjustment cannot be made.”

Echoing this finding, Beach (1990)observed that

“...probabilities mean little to decision makers and have surprisingly little impact on their decisions. Probability is of little concern because decision makers assume that their efforts to implement their

decisions will be aimed, in large part, at *making* things happen. Controlling the future is what decision making really is about.”

Often I find it necessary to explain to clients what we decision analysts assume: that expected utility is the only reliable guide to making coherent decisions, so if the future is uncertain, then some representation of uncertainty is necessary. Too often I find clients making decisions on the basis of judging the overall benefits of projects, ordering the projects from best to worst, then allocating resources down the list until the budget is consumed, an approach that is guaranteed to fail the value-for-money test. So, although the vast majority of my clients’ models are based on MCDA, most of them also provide means for risk-adjusting the benefits.

The foundations

Is all the above simply a collection of ad hoc methods, or is there some theoretical foundation on which decision conferencing is built? My belief is the latter, and this section presents the theories that inform decision conferencing practice.

Requisite decision models

Years ago, when clients told me how valuable the decision conference was even though they implemented something different, I recognized the need for a new class of models. The model whose recommendations weren’t followed didn’t prescribe optimal behavior or describe actual subsequent behavior, nor could the model be considered as satisficing, in the sense of identifying an option that meets aspiration levels (Simon, 1955). Immersing myself in the literature on modeling, I came to the conclusion that the models developed in decision conferences are ‘good enough’ models, which I called ‘requisite’ (Phillips, 1982, 1984). They differ from operations

research models in their definition, generation, process of construction, criterion for ‘good enough,’ and what they represent.

First, I define a requisite model as one that is sufficient in form and content to resolve the issues at hand. The key issue may not be a decision. More typically, it is the failure to agree, a conflict of objectives, which moves the modeling in the direction of MCDA. Second, the generation of a requisite model is through the iterative and consultative interaction among key players and specialists, facilitated by an impartial decision analyst, with all the main perspectives on the issues represented by participants in the decision conference. Third, the process of creating a requisite model uses participants’ sense of unease about model results. The facilitator can encourage participants to express their discomfort or unease, even if they don’t know why they feel it. Exploration of the discrepancy between holistic judgment and the model’s results links people’s emotional and deliberative systems (Kahneman, 2002), helping them to access their experience and knowledge, make it explicit and work on it with the help of the group.

Fourth, at what point is a model requisite? After considerable work on this issue, with the help of Elliott Jaques, I realized that it is when no new intuitions arise in the group. At that point the sense of unease has largely disappeared and the dialectic in the group has ceased. It is then usually possible to summarize the shared understanding of the group, and gain commitment to the next steps. But this doesn’t always happen. In working for one organization, three two-day decision conferences over two years on top-level strategy threw up recommendations that the Chief Executive Officer found unpalatable. Only after the third meeting did he finally understand why; it was a fundamental error in the initial framing of the problem which he felt would become invalid within just a few years. He then said he now felt comfortable with a new way forward for the organization, which within a few years proved to be absolutely right.

Finally, what does a requisite model represent? Unlike models in the physical and natural sciences, this model represents the shared understanding of the participants in the decision conference. Interestingly, the model participates in the reality it creates, a feature that gave me some trouble until I found that it is justified from a philosophical perspective because the source of the model, decision theory, is different from the source of its content.

A requisite model is at best conditionally prescriptive; it suggests what could be done given the frame, assumptions, data and judgments. Its purpose is to help decision makers to construct a new reality, one that more effectively achieves their objectives. Given the temptation for decision analysts to construct beautiful, all-encompassing models, requisite modeling imposes a discipline that clients appreciate. All these issues are elaborated elsewhere (Phillips, 1984).

Process consultancy

The dual role of academic and practicing decision analyst imposed on me a stance to my clients that I soon found was not at all helpful: I knew what was best for them. Decision theory can easily be seen as a prescriptive science: if you believe this, this and that, then you are obliged to accept the course of action associated with the highest expected utility. If you don't, then the model is incomplete, so we must go back and fix it. This is the 'doctor-patient' model of consultancy, once again. In the concluding personal note in his book on process consultation, Schein (1999) eschews the selling of products, programs, diagnoses and recommendations in favor of a helping relationship with the client. He says:

“...help will not happen until the right kind of relationship has been built with the various levels of clients we may have to deal with, and ... the building of such a relationship takes time and requires a certain kind of attitude from the helper.” (p. 248)

While the 'doctor-patient' model is appropriate in some instances, I have found greater success with the process consultancy approach. Schein has said it all better than I can, so I

recommend in particular his ten principles of process consultation, well described in the first three chapter of his book—which is essential reading for my students.

Group processes

Before turning to the topic of facilitating decision conferences, a few words about my views on group processes are required. I have been fortunate to study group processes at the Tavistock Institute of Human Relations whose research on group processes extends back to WW II. At that time, many returning service men and women were in need of psychological help with the traumas they had experienced. As there were insufficient numbers of trained clinicians, the Tavistock adopted group therapy instead of one-on-one help. It soon became evident that quite apart from the healing effect, group processes themselves were being revealed. This led to a shift to studying group processes in experiential courses, where temporary learning communities form and reform, conducted in the ‘here-and-now’ with people drawn from all walks of organizational life, helped by consultants who are specialists in group processes. Much of this work has now been summarized in an anthology of readings, *The Social Engagement of Social Science* (Trist & Murray, 1990).

From this work, and my experience in decision conferences, I take the view of groups as similar to individuals in that the group has a personality, a character and an emotional life. It does not, of course, have a memory. But as individuals in the group take on roles, they find that they shape the group’s personality, which then has an impact on the group. In other words, there is a reflexive interplay, not a cause-and-effect one, between the individual and the group. A common experience of individuals in a group is a disjunction between individual feelings and the group life, which can lead to anxiety.

Mature, productive groups acknowledge the anxiety, tolerate it and hold it, enabling them to get on with the primary task of the group. But often the anxiety remains hidden, and then, just below the surface, the group shifts its emphasis away from working on the task and instead attempts to deal with the anxiety, a covert 'basic assumption' that diverts the group from the overt task as they work on dealing with the anxiety. So far, research has identified five common diversions: fighting with each other or the facilitator, pairing of two people who the group expects will protect the group's security, becoming dependent on an individual in the group or on the facilitator to 'save' the group, developing an intense sense of belongingness, and acting as if the individual is the only reality (Bion, 1961; Lawrence, 1996). This wholly inadequate summary belies the power of these basic assumptions, which take hold of a group as it attempts to deal with the anxiety created by the complexity, uncertainty and conflicting objectives of the primary task.

Facilitation skills

What can the facilitator do to keep the group task focused when such powerful forces under the surface are diverting the group? The key is to focus on the group, not on individuals, and understand the group life (Phillips & Phillips, 1993). This can be done in four ways. First, by observing verbal and non-verbal behavior, particularly the latter, for it contains clues to what is happening in the group beneath the surface. Second, by observing roles and role relationships. Participants often speak from their roles, and relate to each other not so much as individuals, but in their roles. The scientist argues with the marketer, bringing to the table perspectives derived from their roles. Third, by making inferences based on overt and symbolic content of group discussion. "What is this group really talking about when it constantly refers to the inadequacy of its absent leader?" Perhaps it is their own inability to work on strategy, the topic assigned them

by the leader, because they don't know what strategy is. Fourth, and most important, by monitoring one's own feelings, the only route to becoming a self-aware, impartial observer. Attempting to become the scientific, detached facilitator is doomed to fail, for we are empathic creatures, and turning off our feelings while facilitating a group is bound to impose them on the group. This will be sensed by the group, and the facilitator's impartiality will be compromised along with his or her effectiveness in this role.

Only when facilitators feel they have a reasonable grasp of the group life, is it possible to intervene to help the group. Content interventions must be carefully judged, for contributing to content can interfere with effective work if the group feels the facilitator is attempting to impose his or her views. Contributing content can leave the group feeling de-skilled, or encourage it to become dependent on the facilitator, who may have been tempted into contributing content if the group is working on the basis assumption of dependency. Content delivered by the facilitator can hinder the group's ownership of results, and impede implementation, as so often happens in 'doctor-patient' consultancy. Interpretations of group or individual behavior are very inadvisable; they anger participants and can exacerbate basic assumption behavior.

When it is appropriate, several forms of intervention are available to the facilitator. *Pacing the task* helps the group to achieve their objectives within the two days of the decision conferencing. *Directing* the group down paths that the facilitator's experience suggest will be more productive, though any resistance from the group must be faced, with the facilitator ready to work with the group to choose a different path. *Handing back in changed form* is one of the most powerful interventions. It consists of taking in information and drawing a logical conclusion that is presented to the group, who now see the situation in a different light. In one sense, that is what the decision model does; it reassembles the pieces and shows the results,

allowing new properties to emerge. *Reflecting back* to the group, perhaps in different words, “You seem to be saying such-and-such,” is less powerful, but particularly effective in the early stages of a decision conference. Finally, *questioning* and *summarizing* are both interventions that help a group to conceptualize their work in useful ways.

While numerous books have been written on the topic of facilitation, many are little more than a collection of techniques. For mature adults, these techniques can often seem childish, and distracting of the real work. I have found it more helpful to deepen my understanding of group processes, a never-ending task, and to find ways of helping that feel comfortable for me. By spending time developing self awareness, using one’s feelings as data while working with a group, and reflecting later on the group experience, consultants can find their own ways of becoming effective facilitators. One size does not fit all. *The Skilled Facilitator* is one of the better guides to good facilitation (Schwarz, 2002).

Do decision conferences work?

Evidence for the effectiveness and limitations of decision conferences exists in two sources: case studies and systematic research. As examples of the many published case studies, decision conferencing has been used in many public sector cases: to help the Bank of England relocate its Registrar’s Department outside London to reduce operational costs (Butterworth, 1989), to help the UK’s National Radiological Protection Board to develop guidance on relocating the public in the event of a release of radioactivity (Aumônier & French, 1992), to various issues faced by the US Department of Defense (Bresnick, Buede, Pisani, Smith, & Wood, 1997; Buede & Bresnick, 1992), to manage cuts in the budget of the Bedfordshire Police Force in the late 1990s (Holbourn, 1998), and to assist water resource planning in South Africa (Stewart, 2003). Carlos Bana e Costa is a prolific user of decision conferences in the public sector, where he has tackled

problems such as the resolution of conflicts (Costa, Silva, & Vansnick, 2000), bid evaluation (Bana e Costa, Correa, De Corte, & Vansnick, 2002) and evaluation of flood control measures (Bana e Costa, Da Silva, & Correia, 2004).

Private sector applications include the use of decision conferences to evaluate alternative ways of replacing or upgrading ageing production facilities (Phillips, 1986), to help participants develop more effective strategies for dealing with a strike-prone factory (Wooler & Barclay, 1988), to develop better advertising, promotion and distribution strategies (Phillips, 1989), to allocate resources across different sales outlets (Phillips, 1990) and for evaluating and prioritizing projects and creating portfolios in the pharmaceutical industry (Charlish & Phillips, 1995).

Sustained working with clients using decision conferences characterized the development of long-term environmental planning in Hungary (Vari & Rorhbaugh, 1996), and to assist the Social Services Department of Dudley Metropolitan Borough Council to produce a range of budgets that officers felt they could deliver (Morgan, 1993). While these studies convey the flavor and scope of applications of decision conferences, they are a small sample of the thousands of real-world applications. Apparently, decision conference facilitators are too engaged in their work to write it up for publication.

On the research side, McCartt and Rohrbaugh (1995) studied 26 decision conferences held mainly in the public sector. Those rated more beneficial were smaller, hosted by organizations more open to change, and agreed more decisions. In a study of 22 decision conferences in the US and UK, Chun (1992) found that decision conferences were consistently rated higher than ordinary meetings on all of the 12 criteria measuring the effectiveness of decision processes, overall attitudes toward the system, and decision qualities. He also showed systematic

differences between the three facilitators in the study on those 12 criteria, though the effects are not strong, and are compounded with other variables. The perceived effectiveness of decision conferences was greater at senior executive level, than at lower levels in the organizations. Effectiveness was greater for smaller groups (4 to 8 participants) than for medium-sized (9 to 11) or large groups (15 to 18), but an interaction with the facilitator was also observed.

An important insight into why decision conferences work so well is seen in the work of Patricia Regan-Cirincione (1994). She found that small, interacting, facilitated groups performed significantly better than even the most capable members in the groups, contrary to the findings of much research in social psychology which indicated that groups rarely out-perform their best members. She traced the improved performance to the integration of three factors: group facilitation, decision modeling and information technology. The improved performance of the group appears to be the result of participating in the discussion and receiving feedback from others and from the projected model, and from the help of the facilitator in providing structure to the discussion without contributing to the content. It is reassuring to see systematic research providing support for processes developed through observation and experience on decision conferences.

The Big Picture

Since its birth in 1979, the decision conference has become a well-proven approach to engaging a client in a helping relationship. It combines social processes with technical modeling, creating a socio-technical approach to problem solving that benefits from skilled and impartial facilitation of the group, on-the-spot modeling and application of information technology. The social process provides an implicit agenda: state and agree the meeting's objectives, discuss the issues, build a

model, explore the results, and agree the way forward, with considerable iteration of those stages. Modeling provides a tool to aid thinking, not to give an optimal solution or ‘the right answer.’ The model provides a language, which identifies the key elements, and a grammar, which shows how the elements combine. Creating the model provokes thoughtful discourse from participants as they formulate the inputs and interpret the results. The model also polices coherence, revealing inconsistencies in data or judgments, which are then dealt with in the social process as participants explore discrepancies. Overall, the decision conference provides a forum for participants to engage constructively with each other, structuring and focusing the conversation without constraining it. It improves communication across disparate parts of the organization, stimulates creative thinking and improves team-working. Overall, it generates smarter, defensible decisions. Just what the Westinghouse team found, though now we understand why this process-consultancy, helping approach works where more prescriptive approaches have been less successful: *decision conferencing* ensures that the client continues to own the problem and the solution.

References

- Ascher, K., & Nare, B. (1990). Strategic planning in the public sector. In D. E. Hussey (Ed.), *International Review of Strategic Management, Vol. 1*. Chichester: John Wiley & Sons.
- Aumônier, S., & French, S. (1992). Decision conference on emergency reference levels for relocation. *Radiological Protection Bulletin, 133*.
- Bana e Costa, C. A., Correa, E. C., De Corte, J.-M., & Vansnick, J.-C. (2002). Facilitating bid evaluation in public call for tenders: a socio-technical approach. *Omega, 30*, 227-242.

- Bana e Costa, C. A., Da Silva, P. A., & Correia, F. N. (2004). Multicriteria evaluation of flood control measures: The case of Ribeira do Livramento. *Water Resources Management, 18*, 263-283.
- Beach, L. R. (1990). *Image Theory: Decision Making in Personal and Organizational Contexts*. Chichester, UK: Wiley.
- Bernardo, J. M., & Smith, A. F. M. (1994). *Bayesian Theory*. Chichester: John Wiley & Sons.
- Bion, W. R. (1961). *Experiences in Groups*. London: Tavistock Publications.
- Brealey, R. A., Myers, S. C., & Marcus, A. J. (1995). *Fundamentals of Corporate Finance*. New York: McGraw Hill.
- Bresnick, T. A., Buede, D. M., Pisani, A. A., Smith, L. L., & Wood, B. B. (1997). Airborne and space-borne reconnaissance force mixes: A decision analysis approach. *Military Operations Research, 3*(4), 65-78.
- Buede, D. M., & Bresnick, T. A. (1992). Applications of decision analysis to the military systems acquisition process. *Interfaces, 22*(6), 110-125.
- Butterworth, N. J. (1989). Giving up 'The Smoke': A Major Institution Investigates Alternatives to being Sited in the City. *Journal of the Operational Research Society, 40*(8), 711-717.
- Charlish, P., & Phillips, L. D. (1995). Prioritizing projects and creating portfolios. *Executive Briefing, 16*, 33-36.
- Chun, K.-J. (1992). *Analysis of Decision Conferencing: A UK/USA Comparison*. London School of Economics & Political Science, London.
- Clemen, R. T., & Reilly, T. (2001). *Making Hard Decisions with Decision Tools*: Duxbury; Thompson Learning.

- Costa, C. A. B. e., Silva, F. N. d., & Vansnick, J.-C. (2000). Conflict dissolution in the public sector: a case-study. *European Journal of Operational Research*, in press.
- Gergen, K. J. (1993). *Toward Transformation in Social Knowledge, 2nd Edn.* London: Sage.
- Goodwin, P., & Wright, G. (2003). *Decision Analysis for Management Judgment, 3rd edition.* Chichester: John Wiley.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162, 1243-1248.
- Hickling, A. (1990). Decision spaces: A scenario about designing appropriate rooms for 'activity-based' decision management. In C. Eden & J. Radford (Eds.), *Tackling Strategic Problems: The Role of Group Decision Support* (pp. 167-177). London: Sage.
- HM Treasury. (2003). *The Green Book: Appraisal and Evaluation in Central Government.* London: The Stationery Office.
- Holbourn, M. (1998). Decision conferencing--a tool for budget allocation. *Focus on Police Research and Development*, May 1998(10), 22-23.
- Jaques, E. (1998). *Requisite Organisation: A Total System for Effective Managerial Organization and Managerial Leadership for the 21st Century.* Arlington, VA: Cason Hall.
- Kahneman, D. (2002). *Maps of bounded rationality: A perspective on intuitive judgment and choice.* The Nobel Foundation. Retrieved December 8, 2002, 2002, from the World Wide Web:
- Keeney, R. L. (1992). *Value-Focused Thinking: A Path to Creative Decisionmaking.* Cambridge, MA: Harvard University Press.
- Keeney, R. L., & Raiffa, H. (1976). *Decisions With Multiple Objectives: Preferences and Value Tradeoffs.* New York: John Wiley.
- Lawrence, G. (1996). The fifth basic assumption. *Free Associations*, 6(1 (No 37)).

- MacCrimmon, K., & Wehrung, D. (1986). *Taking Risks: The Management of Uncertainty*.
London: Collier Macmillan.
- Matheson, D., & Matheson, J. (1998). *The Smart Organization: Creating Value through Strategic R&D*. Boston, MA: Harvard Business School Press.
- 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.
- Morgan, T. (1993). Phased decision conferencing. *OR Insight*, 6, 3-12.
- Phillips, L. D. (1982). Requisite decision modelling: A case study. *The Journal of the Operational Research Society*, 33, 303-311.
- Phillips, L. D. (1984). A theory of requisite decision models. *Acta Psychologica*, 56, 29-48.
- Phillips, L. D. (1986). Decision analysis and its applications in industry. In G. Mitra (Ed.), *Computer Assisted Decision Making*. Amsterdam: Elsevier Science Publishers.
- 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. (1990). Requisite decision modelling for technological projects. In C. Vlek & G. Cvetkovich (Eds.), *Social Decision Methodology for Technological Projects* (pp. 95-110). Dordrecht, Netherlands: Kluwer Academic Publishers.
- Phillips, L. D., & Bana e Costa, C. (2006). Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Annals of Operations Research*, *accepted*.

- Phillips, L. D., & Phillips, M. C. (1993). Facilitated work groups: Theory and practice. *Journal of the Operational Research Society*, 44(6), 533-549.
- Regan-Cirincione, P. (1994). Improving the accuracy of group judgment: A process intervention combining group facilitation, social judgment analysis, and information technology. *Organizational Behavior and Human Decision Processes*, 58, 246-270.
- RHIA. (1927). Royal Institute of International Affairs. Retrieved, from the World Wide Web:
- Schein, E. H. (1999). *Process Consultation Revisited: Building the Helping Relationship*. Reading, MA: Addison-Wesley.
- Schwarz, R. (2002). *The Skilled Facilitator: A Comprehensive Resource for Consultants, Facilitators, Managers, Trainers and Coaches*. San Francisco: Jossey-Bass.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.
- Stewart, T. J. (2003). Thirsting for Consensus: Multicriteria decision analysis helps clarify water resources planning in South Africa. *OR/MS Today*, 30(2), 30-34.
- Trist, E., & Murray, H. (Eds.). (1990). *The Social Engagement of Social Science: A Tavistock Anthology* (Vol. I: The Socio-Psychological Perspective). London: Free Association Books.
- Vari, A., & Rorhbaugh, J. (1996). Decision conferencing GDSS in environmental policy making: developing a long-term environmental plan in Hungary. *Risk Decision and Policy*, 1(1), 71-89.
- Watson, S. R., & Buede, D. M. (1987). *Decision Synthesis*. Chichester: John Wiley & Sons.

Wooler, S., & Barclay, S. (1988). Strategy for reducing dependence on a strike-prone production facility. In P. Humphreys & A. Vari & J. Vecsenyi & O. Larichev (Eds.), *Strategic Decision Support*: North Holland.