

WHAT THE CASES DON'T TELL US

by

C. J. CORBETT*
W. J. A. M. OVERMEER**
and
L. N. VAN WASSENHOVE†

95/76/TM

- * PhD Candidate, Production and Operations Management, at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.
- ** Assistant Professor of Strategy and Management, at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France and at Stern School of Business, New York University, New York, USA
- † Professor of Operations Management and Operations Research, INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

A working paper in the INSEAD Working Paper Series is intended as a means whereby a faculty researcher's thoughts and findings may be communicated to interested readers. The paper should be considered preliminary in nature and may require revision.

Printed at INSEAD, Fontainebleau, France.

What the cases don't tell us

Charles J. Corbett¹

(Ph.D. candidate, Production and Operations Management)

Willem J.A.M. Overmeer^{1,2}

(Assistant Professor of Strategy and Management)

Luk N. Van Wassenhove¹

(Professor of Operations Management and Operations Research)

¹ INSEAD, Fontainebleau, France

² Stern School of Business, New York University, New York

August 18, 1995

Forthcoming in: *OR at Work: case studies on the application of OR in industry, service, agriculture and health care*, by Fortuin, L., P. van Beek and L.N. Van Wassenhove (eds.), London: Taylor and Francis, 1996.

1 Introduction

When operational researchers write cases about their practice, they do so with the intent to show others what they have done and how they have done it, and to provide guidance for others on how to apply existing or new techniques to old or new problems. Cases are an important medium for the transfer of knowledge from one practitioner to another, or from seasoned practitioners to novices and students. Sometimes a case will also illustrate to a (potential) client how an operational researcher might be of assistance to him or her.

Over the past few decades, the OR community has realized that focusing only on mathematical techniques is not sufficient for becoming a successful practitioner. In the OR literature, much has been written about the dangers of focusing too much on theory (see e.g. chapter 18 in this book or Corbett & Van Wassenhove 1993). In response, the professional OR societies have launched several laudable initiatives¹ to emphasize OR practice more heavily, eg by providing descriptions of how such techniques are used in practical situations. Yet, if a novice to the practice of OR were to read all 120-odd papers on the highly successful projects chosen as finalists for the Edelman award for excellence in OR/MS practice, it is unlikely that that novice would be able to practice OR. In fact, in most cases, even a competent and experienced practitioner would have difficulties performing the projects as described and prescribed. Why is this?

We will argue in this chapter that successful practitioners do a lot more than they write about, and much of what they do not write about is in fact critical for their practice. Some issues may remain undiscussed because the practitioner is simply not aware of them. Very often, though, their omission seems due more to a pervasive form of what looks like a collective self-censorship, leading to cases that represent OR practice as a linear progression of activities invariably leading to a successful outcome. As a result, the ropes of practice remain hidden behind a streamlined simplicity suggested by many cases, rendering the learning process for others, particularly the uninitiated such as students, novices, and clients, unnecessarily difficult.

¹ Among others, the creation of the journal *Interfaces*, devoted entirely to OR practice, the Edelman competition for excellence in OR/MS practice, and the recently announced EURO award for the best applied paper.

The editors of this volume asked us to begin to address this situation by writing about the practice of three of the practitioners who contributed to this volume. In addition, we studied four more consultants. To break out of the potential hazard of limited representation by consultants of their own practice, we also interviewed members of the client system with whom the consultants worked closely together. An overview of the projects and people involved is given in Table 1.²

Table 1 about here.

2 The limitations of prevailing views of the practice of OR

As our research progressed, we were struck by the two images of the practice of OR that seem to dominate the field. One is the perspective that OR practice is the mere application of OR theory, where the outcomes of (possibly black box) computations are provided to the client who then makes decisions. This view provides a very narrow definition of OR, divorced from most concerns of practice; the limits of this image are widely recognized within the OR community. The other image, which one might infer from reading OR case descriptions, is that OR practice involves a linear process of problem solving within a limited amount of time, proceeding by some well-defined sequentially linked stages such as problem definition, model building, data gathering, model solving, validation, and implementation (see, for instance, Fortuin, van Beek and Van Wassenhove 1992). This image, too, is limited. It does not capture the fact that practitioners are not always immediately successful in a first attempt to address a problem, that trying to solve a problem can lead to new information that throws a new light on the original problem, that problems sometimes have to be reframed as a result, that projects get stuck and sometimes unstuck, that minor and even major disasters occur during projects, and that practitioners respond to and often learn from all these events. These aspects of practice increase the complexity of projects in non-trivial ways, which are mostly unaccounted for in theory and in

The names of some of the individuals and companies involved have been camouflaged; also, in some of the cases, certain numbers have been changed.

accounts of practice, even though these are key ingredients of competent practice of the case writers.

The project came about when one of us, Charles Corbett, educated in a more or less traditional academic OR department and knowledgeable about OR techniques, joined a business school PhD programme and began to wonder exactly how such techniques were used in practice, and what practitioners actually do besides straightforward application. Together with Luk Van Wassenhove, a professor of OR with a concern for use of OR in practice, and Wim Overmeer, a professor in management focussing on strategies for organisational learning and inquiry, he decided to find out by talking to highly competent OR practitioners and map out how projects are conducted, not only based on the account of the practitioner but also that of the user-client.

The observation that there are many important, non-technical aspects that have an impact on the outcomes of OR projects is by no means new. Ever since the early days of OR, people have warned that communication problems between client and consultant occur frequently, and that all sorts of political issues within the client organization can play a role. However, these issues are rarely discussed in any depth in a case description or elsewhere, essentially conveying to readers either the somewhat fatalistic message that “things can and will go wrong, but there’s nothing you can do about that”, or relegating the problem to “organisational behaviour”. In this chapter, we intend to describe aspects of practice that cannot be ignored, and that will hopefully eventually lead to a new image of OR practice. Such an image is, in our view, essential, to mitigate the often slow, painful, and sometimes damaging learning process of beginning OR practitioners. Although just a first attempt, we hope it will constitute a step towards answering the question: “how can we educate someone to become a successful practitioner?” It might also begin to provide the more experienced practitioners with a way to better understand their own competences, and help them reflect on their own effectiveness.

We will begin by offering brief descriptions of the seven projects we studied and highlight critical aspects of the process followed by the practitioners in collaboration with their clients. In the second part of the chapter we will turn our attention to three key problem areas facing practitioners as well as concerned academics:

1. *Project management issues.* We will focus on the two related issues of problem framing and project planning. The view of OR as an application of techniques assumes that the client has set the problem and that the task of the OR practitioner is to merely select the best techniques given the way the problem is

formulated. The image of OR as a problem solving cycle assumes a sequential process of defining, analysing, and solving a problem, and then implementing the solution. However, we found that the problem need not be framed by the client, or that the framing of the problem shifts as the project takes place. Moreover, some consultants and clients start with just a hunch, or a very broad problem definition, while in other cases, they begin with a narrow problem definition. We will return to these issues in section 4.

2. *The interaction between client and consultant.* With differing ways of framing problems and planning projects, the interaction between client and consultant takes center stage. Key actors in the system -- those who take the final decision, those who frame the problem, those who provide input, those who use the outcomes -- tend to change, and this can cause significant disruptions to the project. In such a context, the gathering of valid data is a non-trivial and often underestimated problem. Moreover, communication problems between client and consultant are likely to arise either during the framing, the implementation, or the re-framing. Finally, the meshing of the client's expertise about his own problems and that of the consultant about OR can be complicated: some (but not all) clients have considerable technical, mathematical or even specific OR expertise, while some consultants (but again not all) have gained substantial know-how about the client's domain through past projects. These issues are addressed in section 5.

3. *The learning effect.* This is frequently referred to as an important if not the main benefit of an OR project. However, it remains an elusive and even somewhat mystical effect, and frequently occurs as an unexpected by-product rather than as the primary goal of a project. What exactly does a client learn, and how? This issue will be addressed in section 6.

In the rest of the chapter, our main purpose is to illustrate that the images of OR as application of technique, and of OR as a linear problem solving cycle, are seriously incomplete overall prescriptions for OR practice. They provide a false sense of stability, a stability we have not found during our discussions with successful practitioners. The framing and re-framing of problems, the sometimes open-ended projects, changes in the client-system, the meshing of complementary and overlapping expertise, problems in data gathering and communication, all lead to what we believe is a more realistic image of practice that is messier, more unstable, and more ambiguous. Under those conditions of practice, practitioners display skills not captured in the traditional images of OR. Yet, we believe that those skills are central to competent and successful practice. Learning to be an effective OR practitioner hinges on acquiring those skills, which constitute the

ropes of OR. In this paper we point out some of the key areas in which skills need to be developed. In a follow-up paper (see Corbett, Overmeer & Van Wassenhove 1995) we have tried to develop ideas about how such skills are acquired.

The rest of this chapter is structured as follows. First, in section 3, we provide preliminary outlines of the seven cases on which this chapter is based. The outlines focus entirely on the non-technical aspects of each case, and in particular on the questions raised above. In sections 4, 5 and 6 we discuss each of those questions using the evidence provided by the cases. In the conclusions, we place the findings of this chapter in a broader context, and suggest how they relate to a number of concerns held by many within the OR community.

3 The cases

We now present seven outlines of cases that illustrate critical features of the process of OR. The first two cases -- "creating meaningful performance measures" and "optimising global sourcing" -- are two small projects by two different consultants. In both cases, the problem was relatively clearly defined by the client. The next two cases -- "simulating a complex production process" and "implementing a customised cutting stock system" -- are larger projects. In both cases, the client had a fairly clear sense of the problem. However, in the first case, the consultant was "on call", doing a series of smaller projects and two larger ones. In the second case, the client asked for two pilot studies before embarking on the project itself. In the fifth project, on "decision support for tactical manpower planning", the consultant was asked to identify promising projects. The last two cases -- "decision support for office allocation" and "jointly developing a decision support system for asset liability management" -- are different in that the client had a broad and open problem definition, and then engaged with the consultant in a process of narrowing and re-framing the problem definition.

3.1 Callebaut: creating meaningful performance measures

The client framing the problem. Frans Callebaut, director of operations at the Belgian chocolate producer

Callebaut Europe, had been implementing various improvement programmes throughout his plant. But he became aware that it was difficult or impossible to say whether these programmes were really working. He concluded that he needed a performance measurement system. Callebaut had a clear sense of what he wanted -- the formulation and implementation of performance measures that would be accepted by his employees as reasonable and practical throughout his operations division. But he did not know how to do this.

The consultant improvising on previous experience. Leo Fortuin, at the time with CQM, formerly an internal Philips staff consulting group, had been asked by a Philips director to write a report on the use of performance measures within Philips. When, some time later, Fortuin was interviewed by an engineering journal, about performance measures, Callebaut read that interview just as he was looking for someone to help him introduce performance measures in his division. He invited Fortuin for a first meeting to discuss the situation. Because Fortuin had no experience in chocolate production, he subsequently visited the plant, and let himself be guided around by employees to make sure he understood it well enough to be able to conduct the project confidently. From the many projects he had performed within Philips and CQM, Fortuin had accumulated a lot of experience related to production processes. This gave him an intuitive feel for the type of measures that would be needed in Callebaut's case, and how they could be implemented in practice.

Creating meaningful measures. Fortuin developed an action plan, which he discussed with Callebaut in a second meeting. They decided to let functional work groups of employees define their own performance measures, in brainstorming sessions, supervised by Callebaut and Fortuin. This way, Callebaut could be sure that the people who would have to implement and use the performance measures would support them and that the measures were practical.

The consultant as facilitator. The action plan was carried out in a one-week implementation session. During that week, Fortuin acted as facilitator in the group sessions. Callebaut was present all the time, so whenever a measure was agreed upon, he could authorize it directly when an employee would remark that such data were hard to get, Callebaut would instantly order that they find out why those data were hard to get, and how they could be collected in the future.

Extending the performance measurement system. After about three months, enough data had been collected to make the improvements resulting from Callebaut's efforts visible; this was as Fortuin and Callebaut

expected. On reviewing this chapter, Callebaut added that "the performance measurement system is working very well and the key performance indicators are now used during the Board of Management meetings as a base for further actions", and that "a similar performance measurement system has been implemented in the meantime as well at our subsidiary in the UK". He lamented that it is very hard to push the performance measurement system beyond the operations division.

3.2 Monsanto: optimising global sourcing

The client sensing an opportunity and preparing himself. François Cammaert was the head of the logistics department of Monsanto for Europe and Africa. Together with a financial analyst, he was thinking about the issue of global sourcing of agricultural products: for each customer, from which plant worldwide should the raw materials be sourced, where should the product be formulated and where should it finally be packaged? A trade-off has to be made between production costs, transport costs, and import duties and profit taxes, each of which depend on the routings chosen. Cammaert was convinced that Monsanto could do better than their current unstructured approach to this problem. He and his staff had already spent two years thinking about how to improve their product sourcing decisions. He had even formulated a simple mathematical model describing the problem. What he needed was technical expertise to assist in developing and implementing a working version of his crude initial model.

The client's problem formulation improved by a consultant entering a new domain. When Cammaert attended a seminar on linear programming organised by the Belgian OR society in which Beyers & Partners took part, he was struck by their emphasis on applications of LP, so he invited them to come to Monsanto. Beyers & Partners had no particular experience in global sourcing problems, but they were experts in applications of linear programming. During their first meeting they could instantly point out the shortcomings in Cammaert's own model and suggest a modular approach, which was exactly what was needed.

Little formal planning and quick results with a prototype for part of the problem. The Monsanto project was a small one, encompassing a period of less than two months, during which time it was not a full-time concern for either of the two Monsanto staff or for Anita Van Looveren, the consultant involved. No formal project planning was performed, although clear agreements were made on project content, time involved,

budget, and delivery date, for both of the two-month stages: 1) modelling and validation, and 2) application and documentation. One of the difficulties Cammaert encountered was in obtaining figures for import duties and tax rates. The fact that Cammaert had already devoted so much thought to the problem meant that he had a very deep understanding of the issues involved, and could quickly explain any surprises that occurred in building and using the model. After four or five weeks, the first prototype was completed. Even though only part of Monsanto's operations had been included, it already gave useful results. The remaining four weeks of the project consisted largely of improving what had already been done, such as improving output screens, and of documenting the model.

Results of the project, including a new project. The model has been used several times, to choose routings given the current situation, and also to evaluate possible changes in the network. It is a strategic model, and is run approximately once or twice a year, whenever the situation has changed sufficiently to merit a revision. It has already saved at least \$ 500.000 to \$ 1 million, and has given Cammaert and his staff various useful insights into the structure of the global sourcing problem. Although this was Cammaert's first encounter with Beyers & Partners, he came to trust them enough to engage them in a second, much larger project, even though Monsanto headquarters in the US advocated a different consulting firm, taking a certain personal risk in doing so. On reviewing this chapter some time after the interview, Cammaert wrote that "the evidence after another 2 years is however that the system is less used than in the past, since several of the constraints, which were at the base of the model, have been designed away since then. Maybe this is another hidden benefit of this modelling approach!"

3.3 Compeco: simulating a complex production process

A client with a hunch, and an atypical project for the consultant. In 1988, Jan Smits became leader of a project team within Compeco, responsible for designing a plant to produce coloured liquid crystal displays, with many consumer electronics and industrial applications.³ This was to be the first such plant in Europe, so

Some of the more technical details are described in chapter 12.

no experience with the production process involved was available. The project team consisted largely of engineers, each responsible for part of the total production line, which comprised several hundred process steps. In designing the line, the team had to trade off factors as throughput times, work-in-process, and capacity, but had little insight into how these factors related to one another.

While Smits had no previous experience with OR, he had, as a civil engineer earlier in his career, been involved in determining the tensions in ram and pile, using simulation; as a result, he recognized that the static spreadsheet calculations performed already in designing the plant gave insufficient insight into what were inherently dynamic processes. “Piling isn’t a static process either, it’s dynamic, you’re hitting something, and hitting is never static. It was a very different simulation programme, but that explains my enthusiasm to find out quickly how [the plant] would behave dynamically.”

A pilot simulation of a subprocess. Smits asked CQM, a consulting group which had close links with Compeco, to perform some simulation studies. Fleuren, the consultant, could contribute both his simulation expertise and his experience in modelling production processes from previous projects at CQM. Smits first commissioned a simulation of one of the subprocesses, by way of pilot study, to see whether he and Fleuren could get along well. Smits was sufficiently satisfied with the results to ask him, in 1989, to perform a simulation of the entire line, which took Fleuren about 50 days. To use the words of Smits, the aim of the simulations was to know more about key design parameters of the plant. “What should I pay attention to? If I turn this knob, how will the plant react?” Fleuren was then called in another three or four times for various small projects on subprocesses. Fleuren remarked that Smits and he had an efficient way of framing the problems: Smits would call Fleuren to talk about some subprocess that needed detailed analysis, after which Fleuren would write a proposal describing the problem, the approach he had in mind, and the time the analysis would take. That way, Smits could check whether Fleuren had interpreted the problem correctly.

Problems with data gathering, and finding an intermediary to help out. In 1992, when the design of the plant was nearing its completion, Fleuren was asked to do another simulation of the entire line. Particularly in the early stages of Fleuren’s involvement, the design of the as yet non-existent production line continually changed, but adapting the simulation model to each new design would have been prohibitively expensive, so a decision had to be made which configuration of the line to simulate. In the overall design project, certain decision points had been specified; the configuration as determined at one of these decision points was chosen

for the simulation. Fleuren pointed out that this was not a typical project: it is more common to be called in to study an existing production line rather than to be involved in the design stage. The problems they are then confronted with could often have been prevented with more careful design.

Gathering data was often a problem. To be able to perform useful simulations, Fleuren needed accurate estimates of machine breakdown frequencies, processing rates, etc. Some of the process engineers did not see the importance of such figures, and could not be bothered to come up with precise estimates. Fortunately, one project team member, Van Gool, helped make Fleuren's task easier. Being a member of the team (which Fleuren was not), Van Gool was closer to the process engineers, and knew how to get the data needed.

What the project team learnt from the simulations. The simulations led to a number of important insights for the project team. They brought out, among others, the drawbacks of a line which is too carefully balanced, the importance of the location of the maintenance staff, and the effects of different material handling systems. This resulted in various minor and major design changes in the production line, which represents a total investment of several hundreds of million dollars.

3.4 Oilco: implementing a customised cutting stock system

A long incubation. Two years before the project actually started, some staff members at Oilco in Belgium believed they needed customized software for their packaging foil cutting machines. BOR Consultants was already acquiring a reputation for their experience with cutting problems, a reputation which has grown since. The Oilco staff visited BOR, went back to talk it over, and returned a year later. There was a more primitive system in use at Oilco in the US, but the Americans were persuaded not to impose it on Europe too, so eventually permission to do a project with BOR was granted; the project was performed by a team, with David Bookman as consultant in charge.

Two pilot studies, and tight project planning. The project started with two pilot studies: 1) a functional pilot study, concerned with defining the problem in detail, and specifying what the software should be capable of, and 2) a more organisational study of hardware and software aspects, revolving around issues such as how to build the interfaces, who should install which database, on what hardware, what type of output the system

should generate, etc. A presentation to general management was made after the first preliminary study. For the main body of the project, a contract was drawn up, that distinguished a number of stages. Although much of the software had to be developed from scratch, BOR were able to plan and execute the 18-month project very smoothly and tightly, based on their previous experience in this domain. Despite some changes in the composition of the project team over its 18-month total duration, deviations from the plan were limited to a few weeks. Integrating newcomers into the team provided no serious problems.

Coaching the users. Cooperation with the planners, the future users of the cutting software, generally went smoothly. However, given the size of the plant and the cutting machines, it is understandable that planners sometimes felt uneasy about fully transferring control to the new software. In such cases, it was Bookman, the consultant, who had to coach them along, for example by being present during the first days after the actual changeover to the new system, and by installing a modem link with the BOR office. As Bookman summarised his coaching of the planners: "I'll jump first, you just follow me".

A new project. This project concerned the optimisation of an existing process; as a result of its successful completion, BOR were subsequently called on in a project involving the actual selection of machines.

3.5 Transco: decision support for tactical manpower planning

A committed client. William Gray had just been appointed Head of the 2500-strong Vehicle Maintenance Department at Transco, a very large transportation firm. He has an engineering background, and tries to apply analytical reasoning wherever he can. He did not have any particular problem in mind that urgently needed solving, but held a strong conviction that OR could help him run his division better. The in-house OR group at Transco had suggested to Peter Brown of ABC Consultants to get in touch with Gray. The latter invited Brown to talk to the heads of each of the subdivisions and draw up a list of possible projects.

The consultant defining projects. It was left to the consultants to come up with good proposals, which they were able to do, despite not being familiar with the problem domain. These were presented to Gray, and they started on the one promising the largest return. This first project was to develop a system to help Gray determine how much staff he would need for a given type of activities in his division, for given work

schedules, for given collective labour agreements, for given organisation structures, etc. The system would help with a range of medium and long-term manpower planning issues, such as negotiating labour agreements.

Working in an informal way. Project planning had not been done in a very formal and precise way. This is precisely how Gray wants it. He remembers that “that project was somewhat untransparent, but if we have a good relationship with a consultant we don’t make very formal fixed contracts.” This fits well with Brown’s view: “We usually don’t build systems the way a computer scientist would advise, i.e. defining upfront what the system should do and adding increasingly more details.”

Changes in the client system, and re-framing of the project. Unfortunately, by the time the system was completed, Gray had moved to a different division within Transco. His successor was not interested in the system, and Brown was relegated several levels downward in the hierarchy. However, for the people at that level, the type of problem Gray had been concerned with did not exist; for them, labour agreements etc. were a given, and considered unchangeable. They do use the system, but for purposes of a far more operational nature than intended, such as determining how much work the division will have during the coming period, or determining detailed work schedules. In the original system, the detailed work schedules were not relevant in themselves, only as a means for evaluating the impact of higher-level decisions. The system has never been used for its original purpose. Gray, however, is still very enthusiastic about OR, and since having moved to a different department, he has initiated several projects there with Brown and ABC.

3.6 RGD: decision support for office space allocation

A client senses an opportunity. RGD is a Dutch government agency in charge of housing for all government agencies. The director of the research department of RGD had been struggling for some time with the problem of allocating buildings to organisations. When he happened to see a personnel ad from ORTEC he realized, from the way ORTEC described itself, that they were the kind of consultants he needed. As a result, ORTEC and he set up a project, called ‘BOSS’, to develop a system for the building allocation problem; for ORTEC, this was an entirely new field. While that project was underway, the Ministry of Health and Environment (VROM) announced a move to a new building. Such a large-scale move often entails employees

having to change office several times. Kees Wassenaar, a member of the research department and an expert in the realm of housing for government agencies, wanted to help minimize the number of such internal office changes that would be needed during the move. He knew of BOSS, so he contacted Guus Boender of ORTEC. Together they initiated a new project, which was “strongly aimed at reducing the number of times people had to change offices during an internal move within one building”, as Wassenaar put it.

Client and consultant favour informal cooperation. The intention was to complete the project in time to be able to assist in planning the upcoming move. But, rather than draw up a contract and a detailed work programme for the entire project, Wassenaar and ORTEC decided roughly which direction to take and worked on the basis of periodically renewed fixed-term fixed-sum agreements. There was no clear goal to be met in each period, both parties proceeded on the basis of mutual trust. Every time Wassenaar and a colleague discussed the problem with ORTEC, the consultant would write a brief report to test their understanding of the problem with Wassenaar. After three months a proposal resulted, with a suggestion for an approach and an estimate of costs.

Underestimating complexity leads to re-framing the project. The original goal turned out to be much more complex than either side had foreseen, and they were not ready by the time VROM moved into its new building. In response, they decided to simplify the problem. The focus of the project evolved from allocating agencies to buildings into allocating office space within a building to departments of agencies. Wassenaar: “We decided to concentrate on simply matching organisations and buildings, before trying to do something about moving organisations around. We couldn’t even find a good matching within one building, so how can you even think about reducing the number of moves in going from one matching to another?” Allocation of office space within an agency used to depend on which department had the most power; Wassenaar hoped that with a system like this the allocation procedure could be made much more objective.

At the same time, Wassenaar recognized the need to explicitly identify a new target group, to keep the project moving, because he felt that “there has to be a deadline, then people try harder.” With the change of focus, the group of potential future users of the system also changed from site managers wanting to reduce the number of office changes to housing consultants within RGD, who check whether certain buildings are suitable for certain organisations. For instance, if an organisation has five buildings to choose from, the consultant has to determine how well the organisation can be accommodated in each of the buildings.

The original approach appears based on a false assumption. One of the major complications in the early stages of the project was that the pilot study was based on what turned out to be a false assumption. The approach proposed was to generate a large number of admissible allocations meeting all relevant criteria and then choose the best of these allocations. However, this presumes that such admissible allocations exist! It turned out that no allocation met all criteria (hence, perhaps, the traditional power struggles), so the problem had to be reframed as finding a collection of allocations that would meet as many criteria as possible, and would let the final decision-makers choose from that collection, based on which criteria they perceived as more important.

An unintended further shift in problem focus. A natural approach in tackling the problem of allocating office space to entire departments is to determine an allocation of individual people to individual offices. As Wassenaar said, “the concept of allocating people to offices became very dominant, so the software’s output was in terms of people and offices, but that isn’t really relevant. What matters is the differentiation of office space”. Although both parties seemed to understand the problem, the project unintentionally shifted from its original more global goal to the more detailed level. It was only when Wassenaar decided to stop working on the algorithms for a moment and focus on the output screens that this shift in problem framing became apparent.

The client contributes to the algorithmic development. Boender believes that the client’s contribution to the project was critical for its success so far: “the people with whom we dealt had an exceptionally deep knowledge of the issues they were facing and grasped our approach and our contribution quite quickly”. Wassenaar confirmed this when he said: “At some point in time you want to know how the algorithm works, you think they’ve understood the problem but when you see the algorithm you see that everything works in one particular way and you wonder, why not some other way?” Hence, it took ORTEC’s experience in developing solution methods combined with Wassenaar’s insight into the problem to construct an efficient solution method.

3.7 AMEV: jointly developing a decision support system for asset liability management

The consultant senses an opportunity to expand on previous work. The history of the development of the asset liability management (ALM) system started around 1984, when the Rabobank hired Guus Boender of ORTEC to help them develop a system for liability management for the Rabobank's own pension fund. When the project was completed, Boender saw that if the system were expanded to include a wider variety of assets, it would have great potential. Because of changes in the legal and economic environment, pension fund managers needed a new type of tool to evaluate investment strategies.

The consultant sets up joint ventures with clients to further develop his system. Together with three partners, one of which was AMEV, a major insurance and pension fund management company, Boender further developed the ALM system.⁴ The connection with AMEV was made through a former colleague of Boender's who then worked in the investment department at AMEV. This former colleague persuaded AMEV's Voûte, then in the liability department, to cooperate with ORTEC. Voûte saw the possibilities of the consultant's system but also identified serious shortcomings for the applications he had in mind. To further develop the system in AMEV, Voûte set up a project team, consisting of a business economist (himself), a macroeconomist, a lawyer (Coen van Dedem), some actuarians, and an econometrician; most of the members of the project team at AMEV knew nothing about OR. The project started on the basis of what Boender had developed for the Rabobank. The intention was to combine Boender's knowledge of OR with AMEV's knowledge of pension fund management, in order to develop a more powerful system for asset liability management. ORTEC intended to sell the system to various pension funds in the future, AMEV was looking for a tool to help them assess the impact of different investment strategies and pension schemes, as part of the consulting service they offered to small pension funds.

From very loose to very tight project planning. Initially there was no project planning at all, the project started as a very loosely organized joint venture. ORTEC was not operating as a consultant to AMEV, and several AMEV staff invested a large amount of their own time in it. During the first serious meeting, the team estimated they would need about half a year for the project. Eventually, "it took half a year or more just to understand each other", as team member Van Dedem puts it.

Choices as to what to include and what to leave out, such as whether to include foreign currency issues etc.,

Chapter 10 contains a more detailed discussion of the ALM system that resulted.

were continually being made throughout the project. Initially, the team became more and more ambitious, and wanted to incorporate more and more, but at a given point in time, Voûte said, they “came back to Earth, with both feet on the ground”, and started eliminating options. For example, the team initially tried to capture career paths of individuals within a firm in a very precise way in the system, but then realised that this would be far too complicated. ORTEC suggested a much simpler, approximate method, which was then used. The framing of the problem did not change in any fundamental way during this project, but it gradually became more specific.

It was Voûte who decided, in November 1991, well into the project, that some more time pressure was needed. After taking stock of the work that still had to be done, a deadline was set for May 1992. In fact, Voûte announced that a presentation about the project would be given, and invited a large number of people from outside the team, to set a firm deadline. From that moment on, everything was planned very tightly, tasks were assigned week by week, people reported back and forth; they slipped from that plan by less than a month.

The client reorganises as a result of the project. For AMEV, the project had major consequences. Traditionally, the firm had an investment department and a pension liability department, so each pension fund had to deal with people from two departments within AMEV. The ALM system formally unites the two perspectives, which had previously been operating separately. Correspondingly, AMEV reorganized their pension fund management service, so that they now have one special department dedicated to supporting pension funds, of which Voûte is director; most members of the project team now also work in that new department. What ORTEC now has is a “system as basis for customisation”. With the larger pension funds, there are enough actuaries and econometricians with whom ORTEC can perform studies, the smaller funds are serviced by ORTEC’s partners.

Enthusiasm carries the project along. Voûte, without prior experience with OR, found that it was initially not easy to understand the OR consultants. For those involved, most of the work had to be done in addition to their normal workload. Due to a communication gap that existed initially, it took the team more than six months to reach a “common frame of reference”. A key success factor was the enthusiasm of the project team, particularly given the unofficial nature of the project. Van Dedem, one of the members, remarked: “We all found the subject so interesting that we all got carried away, we all became as enthusiastic as Guus

[Boender].” They perceived the project as an intellectual challenge.

Having sketched the seven cases, it is now possible to look into several critical issues such as project management, client-consultant system issues, and the learning effects related to each of the cases.

4 Project management issues

With respect to project management, two issues seem particularly important: that of deciding what the problem is, and that of planning the project. The two are obviously related, but we discuss them separately here.

Who frames the problem? While the prevailing images of OR practice assume that a problem has been framed early on, and that it has been fixed, we found much more variety in this area. In the Transco case, we saw that the consultant was invited to look for problems and define them himself, whereas in the Callebaut and Monsanto cases the respective clients followed the opposite approach by framing the problem before calling in a consultant. Callebaut had done some preparatory work before inviting Fortuin, and believes that “when you engage consultants, it’s very important that you’ve written down a rough draft of what’s expected. I believe that’s essential, that you make the consultant’s task very clear from the beginning. You have a sort of proposal, about three pages, then you can see how he reacts to it.” Similarly, it was Monsanto’s Cammaert who wrote down a brief description of what he wanted from the consultants: “I think that’s what we should do, we should write down all we know and say “look, this is how far we’ve got in our thought process.”” In Cammaert’s opinion, such a degree of preparation on the client’s part is necessary before embarking on an OR project.

When is the problem framed? The problem as framed by Cammaert was very precisely defined: “To be able to apply OR you first need a basis. Not all problems are sufficiently well developed to bring them within the reach of OR.” In the Compcoco case too, Fleuren and Smits also made sure that they agreed on what to study. Alternatively, the framing of the problem is not necessarily completed at the beginning of a project. The ALM project, for instance, started without an unambiguous definition of what the aim of the project was. The framing of the problem did not change in any fundamental way during this project, but it gradually became

more specific. Van Dedem: “Some parts of the system became more detailed, of others we decided not to include them in a detailed way but in a rougher way instead.”

Shifts in problem framing? Lastly, it became clear that a problem, once framed, is not immune to shifts, which can be caused by a variety of reasons. The focus of the Transco project shifted to a far more operational level when the project was delegated to a lower level in the hierarchy, as a result of Gray’s departure to a different department. In one of the cases, although agreement had been reached on what the project would involve, one of the clients changed his mind halfway through: suddenly he wanted the project to embrace the total production control system rather than the one particular issue originally singled out. This would of course have required a complete revision of the project, so the consultant could not and did not agree; the project remained as initially intended. Sometimes a shift in focus is desirable, though: in Bookman’s experience, clients often come with some operational manifestation of a problem, whereas the actual cause is often at a deeper, more strategic level.

The RGD project witnessed three shifts. First when it became clear that the project would not be completed in time to assist with the Department of Environment’s move, Wassenaar and Boender decided to set themselves a less ambitious target. Second, when the development of algorithms was in full swing, the focus shifted unnoticed from allocating office space to departments to allocating individual people to individual offices. It was only after the decision to stop working on the algorithms for a moment and focus on the output screens that this shift in problem framing became apparent. Third, when admissible allocations turned out to be non-existent, the objective had to be reframed as finding a collection of allocations meeting as many criteria as possible, and letting the final decision-maker choose from that collection, based on which criteria he perceived as more important.

4.2 Project planning

Tight or loose planning? Whether a project should be tightly and precisely planned does not seem to depend on the size of the project. Each of the simulations performed by Fleuren constituted a relatively small project, but for each of them he would write a proposal describing the problem, the approach he had in mind, and the time the analysis would take. The Monsanto project was also a small one, but here little formal project

planning was performed. “Let’s just begin, we’ll see where we end up,” as Cammaert said. He sees no advantage in formally planning such a small project, as so few risks are involved; in fact, he does not really believe it is possible, “you would be introducing a handicap”. The only formal planning that was done was intended more to pacify the purchasing department, who controlled all expenditures.

Early or later? We find the same variation in large projects, as in the Oilco and ALM cases. The Oilco project was planned in detail, with two pilot studies, and executed accordingly. According to Bookman, distinguishing well-defined stages in a project contract has three reasons: first, to monitor progress; second, to define checkpoints at which tests are to be performed or documents handed over, and to link payments with these; and third, the client also needs to write certain interfaces, so they need to know when certain parts of the software will be available to them for testing. In the ALM project, there initially was no project planning at all. Voûte decided, well into the project, that more time pressure was needed, and from that moment on, everything was planned very tightly.

Reducing the impact of changes in the client organisation. Although precise planning was perceived as undesirable by several consultants and clients, it does seem to be a way of reducing the impact of changes in the client organisation. Compare the Oilco project where, despite some changes in the project team over its 18-month total duration, deviations from the plan were limited to a few weeks, to the Transco case, where Gray’s departure led to a drastic change in focus of the project and to the abandonment of the other proposals which Brown and Gray had informally agreed on. Gray himself, however, is in favour of a flexible approach. About ABC, he says “they’re flexible, informal; they don’t keep on nagging about what’s formally in the contract; that’s good, because as client you don’t always know exactly what you want, other things sometimes turn up unexpectedly.”

5 Client-consultant system issues

5.1 Changes in the client-consultant system

Effects of changes. Especially when dealing with larger organisations and with longer projects, consultants

have to be prepared for changes in the people they deal with. Several consultants remarked that such changes have become a fact of life in long-term projects. In Bookman's experience, it is normal "that all the people with whom you made the initial agreements are no longer there when it's finished. That can be very disturbing, that's why a good pilot study is so enormously important." Van Looveren expressed similar experiences. Fleuren, too, had to convince each newcomer of the value of his simulations every time the composition of the design team changed. Wassenaar, who conducted the office space allocation project with one colleague, admits that "if either of us had been given a different function, the project would probably have fizzled out." We have seen the effects of Gray's departure to a different department within Transco. An open question, though, is whether a more formal approach would have prevented Gray's successor from relegating the project downwards in the hierarchy. People changing place during projects is a common occurrence, Brown says, and to make oneself less sensitive to such changes one might have to accept the "necessary evil of defining much more clearly what you're doing".

Coaching newcomers. Clearly, consultants need to spend a lot of effort on coaching newcomers into ongoing projects, to ensure that the newcomer feels equally involved as his predecessor. But even after a project is completed and, for instance, software is delivered to the client, its continued use is far from guaranteed. When the system had proven its due, Cammaert made sure it was properly documented, so that if he were to change to a different function, his successor could also benefit from it.

5.2 Gathering data

Users participate. That collecting data can be a tiresome task is no news. Particularly if the people who possess the data are not actively involved in the project, they may see no reason to waste effort getting the precise data asked for. During the group sessions with Fortuin, that problem was pre-empted by Callebaut's presence throughout the sessions. Whenever an employee would remark that the data needed to implement a new performance measure were hard to get, Callebaut would instantly order that they find out why those data were hard to get, and how they could be collected in the future.

Using an intermediary. The way data were collected for Fleuren's simulations and to formulate Cammaert's model are interesting, so let us look at each in turn. To be able to perform useful simulations, Fleuren needed

accurate estimates of machine breakdown frequencies, processing rates, etc. Some of the process engineers did not see the importance of such figures, and therefore could not be bothered to come up with precise estimates. Fortunately, one project team member, Van Gool, dedicated much effort to making Fleuren's life easier. Being a member of the team (which Fleuren was not), Van Gool was closer to the process engineers, and knew where to get the data needed. Smits: "Fleuren didn't have the connections to do that." Van Gool was so enthusiastic about Fleuren's work that he became "emotionally involved" in the simulations, and, where necessary, used that emotional involvement to influence the engineers into cooperating. By showing engineers how different the results of the simulation would be depending on whether they gave a 95% or a 99% machine reliability estimate, they began to see the importance of coming up with good estimates rather than seat-of-the-pants figures. Fleuren also addresses this issue in chapter 12.

Showing gaps in the data. One of the difficulties Cammaert encountered was in obtaining figures for import duties and tax rates. For whatever reason, fiscal specialists did not provide him with precise figures, which he needed in order to complete the model; instead, they would say things like "between 10 and 50%". However, by using the model and performing sensitivity analyses on the unknown parameters, Cammaert could establish which parameters really mattered and when each parameter became important. If it turned out that his product sourcing decisions should depend on whether a particular tax rate was below or above 35%, he could go back to the fiscal department and ask them which was the case. Apparently, the fiscal staff felt more comfortable with this type of question, and could now give Cammaert more useful answers. Cammaert did remark that obtaining and interpreting the cost data and financial data was facilitated by the fact that his collaborator was a financial analyst.

5.3 Communication

Differentiated responses to OR, even within the same client organisation. Different people, even within the same organisation, can react very differently to OR. As the composition of the design team changed, Fleuren had to convince each newcomer of his contributions. During each of the subprojects, however, the people with whom Fleuren directly dealt did not change. Interesting is the difference in reaction to Fleuren's simulation results. Process engineers on the project team often took Fleuren's results as the truth, eliminating the need for further discussion among themselves about "how the plant would behave if ..." On the other hand, when

the results were presented to people from outside the team, who had not been involved in the design of the plant, reactions were sometimes much more sceptical; people found it hard to accept that simulations could provide very useful insights if the input data were not exactly correct. Smits acknowledged that how people react to OR is very variable; although he is convinced of the benefits, he knows of many other people who, in his place, would not have seen the need to perform simulation studies. Callebaut adds that, in any project, it is important that the consultant and the client get on well together; establishing this is the other main purpose of a first meeting, besides discussing the project itself.

A project champion in the client organisation. Recall the difference between Gray's positive attitude towards OR and the dismissive attitude displayed by his successor. Gray recognizes that many managers do not have a quantitative mindset; what surprises him is how many managers do have an analytical background but have "deserted", too often as a result of "intellectual laziness". Brown recognizes that Gray and other clients often enjoy talking with consultants, and that this can be an important reason for them to initiate projects. Gray realises, however, that his enthusiasm is very personal. Indeed, his successor was not interested in OR, and would not let Brown change his mind.

The client invests in understanding the consultant. A sceptical attitude towards OR may pose a challenge, but not an unsurmountable one, which is clearly illustrated by Voûte. His opinion is unambiguous: "I have an aversion to anything mathematical. Just say it in plain English, I always say." Having no prior experience with OR, he found that it required an investment to understand the OR consultants; he believes this should not be the case, that OR consultants should improve their communication style. "I think the big problem is that the OR consultant doesn't make it sufficiently clear what he's doing. He shouldn't use formulae and programmes, because any client who isn't an OR person himself is wary of anything that's like a black box." An important consideration here is that Voûte and his staff have to understand the system and its philosophy well enough to be able to explain it to their own clients. Over the years, Boender has become sufficiently expert in the realm of pension fund management that he can compete with non-OR consultants traditionally more closely associated with pension funds.

Complexity may frighten. Capturing the interest of future users proved difficult in the RGD project. When selecting a new target group, Wassenaar did check whether the problem, in its new incarnation, was considered relevant; this was widely confirmed. But to get people to use the system will still take some time

and effort, because, as Wassenaar explains, “they haven’t really been involved, it isn’t their project. It’s something which has been invented and of which they see it could be useful, but the complexity frightens them.” Some of the regional staff, the current target group for the system, were closely involved in its development, and they are enthusiastic. Others, however, were not involved in the 20 sessions leading up to the development of the system. It is turning out to be difficult to persuade them of its potential. ORTEC and Wassenaar intend to deal with this issue by organizing information sessions and by closely supervising a number of office allocation projects with people actually using the system.

Starting with the user interface or with the "inside" of a system? According to Boender, this problem was caused because they had started with the “inside” of the system, the algorithmic part, leaving the design of the user interface until later. Usually, ORTEC takes the opposite approach. When they are dealing directly with planners who are future users of a system, they start by building the input and output routines, to give the users an impression of what the system should be able to do, and only when the users’ interest has been awakened, they develop the mathematical inside of the system. Boender: “I blame myself that we aimed too high, in the early part of the project we tried too hard to squeeze everything in, while one of ORTEC’s strong points is to always start with something small and then expand it. Now both the RGD and we believe that one of the reasons why the system isn’t taking off as fast as we’d like is that it’s too complex.” When asked as to why this was allowed to happen, Boender hesitantly responded “I think it’s because the initiative was theirs”.

Understanding the threat to the users. Dealing with future users of a planning system is always a sensitive affair, particularly when that system will replace some of the current planners. One of the reasons for Gray’s satisfaction with ABC is the way they approach planners. Although several of the systems ABC has introduced at Transco had the potential to be threatening for planners, by making a number of them redundant, the consultants managed to identify closely with the planners, and gain their full cooperation. Brown is very conscious of the challenges getting closer to planners sometimes poses. For instance, he has noticed a tendency by OR consultants to underestimate the complexity of some standard OR concepts for planners. A concept as sensitivity analysis, for example, is harder to convey to planners than is often realized.

Coaching future users. Other consultants are also familiar with these issues. In projects involving implementation of software to be used on a routine basis by planners, the process of training and coaching them is often considered, explicitly or implicitly, the responsibility of the OR consultants. Obtaining a

planner's cooperation requires that management give him enough time to work with the consultants. Van Looveren recommends making that very clear from the start, to prevent problems later on. In the Oilco case, this went smoothly, but whenever a planner had doubts about the project, it was Bookman, the consultant, who had to coach him along, for example by being present during the first days after the actual changeover to the new system, and by installing a modem link with the BOR office. As Bookman summarised his coaching of the planners: "I'll jump first, you just follow me".

5.4 Expertise of client and consultant

Some clients are knowledgeable about OR. Clients of OR projects can have any degree of knowledge about OR. We have seen Cammaert, who was sufficiently knowledgeable about OR techniques to have formulated a linear program himself. We have seen clients as Smits, Gray, and Wassenaar, each of whom had a clear idea of the possibilities of OR, even if they were not familiar with the actual techniques used. Recall that it was Wassenaar who, by asking questions about the precise working of the algorithms, was able to suggest an approach more tailored to the problem at hand. And we have seen clients who knew little or nothing about OR, such as Voûte and van Dedem. Clearly, the client knowing about OR is by no means a necessary condition for a project to be successful.

OR projects can be perceived as risky. Conversely, the consultant may or may not be an expert in the precise domain concerned. BOR clearly are experts in the realm of cutting problems, which can have the important benefit of reducing the risk as perceived by the client. The support for the Oilco project generally came from people who were not high enough to actually authorize it, the people above agreed to the project but largely delegated it to the people below. At a higher level within the firm, a large-scale project as this one entails a certain amount of risk. Bookman: "A client deciding to initiate a project like that is touching the organisation of the firm, the way people plan, the way people do business, that man is taking a personal risk. It's easier to sell desk chairs for the same amount, there's no risk involved in that."

The consultant has to learn about the client's unique problems. Other consultants may be experienced in the general domain involved without being an expert on the particular problem under consideration. Fleuren was a frequent user of simulation in the context of production processes, without being a specialist in production

of liquid crystal displays. Fleuren: “If you don’t include the material handling system, you can never advise your clients. You surely have to have some feeling yourself for which elements matter in a plant like that.” Fortuin, also with a wide experience of production processes and having studied performance indicators, had not actually implemented performance indicators in a plant before. And Boender, at the start of the joint venture with AMEV, had some knowledge of pension fund management, but was not yet an expert in the field. Finally, there were cases in which the consultants had no prior knowledge or experience with the particular problem concerned. Before the projects with the RGD, Boender had not worked on office space allocation-type problems, nor had Brown any prior experience with medium-term manpower planning for maintenance personnel.

6 What did the client learn?

A very important but often underrated benefit of many OR studies is the increased understanding the client gains of his own business. Although the economic value of this learning is often impossible to measure, insight is frequently mentioned as the most important result of OR projects. Unfortunately, this learning effect is very poorly understood. It is not clear what types of insights clients gain, nor when and how these occur or can be made to occur. To illustrate the importance and the diversity of the phenomenon, we take a brief look here at what some of the clients learnt during their projects.

6.1 Monsanto

Gathering systematic data allowed comparisons across products. Cammaert obtained various types of useful insights from the project. To begin with, gathering the data already proved a valuable exercise. “From gathering this cost data, it transpired that some product formulations cost about as much as what they sold for. We’d never seen things from that perspective before, because we were now forced to collect these data in a systematic way, so that they could be compared with each other, otherwise you cannot use them in a model.” Furthermore, it became clear that there were important differences in the relative importance of each of the types of costs involved. The ratio between transport costs, production costs, and duties and taxes appeared to

be in the order of 1 to 10 to 100, which strongly suggests that duties and taxes should be the dominant factor in sourcing decisions.

Making decision rules explicit and testing them. An additional benefit of the modelling exercise was that it made explicit many of the rules Cammaert and his staff had been using in making sourcing decisions, and therefore allowed them to be critically checked. “The results were rather impressive, particularly because the model showed us things we’d forgotten, we hadn’t found those with our logical reasoning. That’s not surprising, when you think about it, we work with so many daily constraints; we think that something isn’t allowed, though we’ve forgotten why, but unless you explicitly include that in the model it won’t take that rule into account. That way we found certain routes of which we thought “that can’t be possible”, but analysis showed that we’d simply overlooked a few constraints. Then you wonder, “if we’d forgotten them, were they really essential?” So basically that model has helped us to better understand our business.” Cammaert summarised by saying “you have these rules from the past, that never get checked. The model allows you to check these rules.”

6.2 Compco

Making explicit trade-offs. Fleuren explained the goal of the simulations as getting insight into the behaviour of the production system, in terms of capacity, work-in-process, and throughput times. The results of the simulations were often presented in the form of graphs of the trade-offs between these key variables. Smits knew that there would be capacity losses due to congestion effects, but the team had no idea exactly how high these losses would be. After the first simulation, they knew that capacity losses would be about 11%. That first simulation also resulted in a major insight for the project team: the perfectly balanced line they had initially designed would have a terribly poor performance, as every process step would turn out to be a bottleneck, not just the most expensive steps. Fleuren: “The interesting thing is that people in a design environment are inclined to balance the line perfectly, but what happens then: that line’s behaviour is truly hopeless. [...] That’s a way of thinking that grows over time. As consultant I’ve noticed that they have gradually learnt to think along those lines.”

Finding design flaws. Various other insights occurred as a result of Fleuren’s work. For instance, the data

suggested there would be a relatively large number of “short” machine interruptions and only a few longer ones. Therefore, it is important that the maintenance staff have a short response time. As a result of the simulations, the maintenance staff and a large collection of spare parts are now stationed on the shop floor rather than in a separate workroom. Smits: “with hindsight, it sounds obvious.” Also, in the initial design, one particular material handling system had been specified. Fleuren however felt that the variability inherent in that system could be significant, and when he compared his simulation results under a “fast” and a “slow” scenario, the difference in capacity could amount to 30%. This was a total surprise for the project team. Smits: “The impact of the material handling system was shocking to me.”

6.3 RGD

Becoming aware of the problem's unsolvability. Perhaps the most important lesson in this project was that the office space allocation problem was simply unsolvable, when taking all the criteria that appeared relevant into account. Rather than find the best of a large collection of admissible allocations, the system would have to concentrate on finding the allocations violating the fewest criteria and letting the actual decision-maker choose between them.

6.4 AMEV

Counterintuitive findings. Voûte and his team gained several fundamental insights into the business of managing pension funds. For instance, it was previously believed that investing in real estate was a reliable way for a pension fund to hedge against inflation. However, studies with the ALM system showed that, depending on the type of people insured by the fund, real estate could be a poor hedge, and that shares could be much more reliable. An insight that grew during the cooperation with Boender concerned the relative importance of the revenues from contributions and the revenues from investments: this turned out to be in the order of a 20% to 80% ratio, suggesting that decisions concerning the investment policies of a pension fund could sometimes have a significantly larger impact than those concerning contribution schemes. More generally, Voûte feels much more confident and comfortable discussing such issues with his clients as a result of the project: “You can communicate much more easily. And you can show things. You know much better

what you're talking about."

7 Conclusions

An OR project as a "kayak trip". Perhaps the most important conclusion to be drawn here is that there is indeed much more to the practice of OR than meets the eye. There are many issues, of which we have only discussed a few here, which are critical to the success of any OR project, but that are hardly, if at all, mentioned in case descriptions. Performing a successful project involves much more than simply applying techniques, and projects are rarely the linear progression that case descriptions frequently suggest they are. To use Bookman's words: "it's like a kayak trip. You can keep your head perfectly above water, but from time to time you need to take a bend. And when the water gets very calm, the waterfall's usually just round the corner."

Practitioners have difficulties articulating their practice. We've seen that the practitioners with whom we spoke are all capable of performing highly successful OR projects, but they all found it difficult or impossible to articulate exactly what it was that made them successful. This state of affairs is obviously a major obstacle on the road to providing better education for students of OR, a better preparation for a career as an OR practitioner. What we need to do, as a profession, is capture the knowledge of experienced practitioners by documenting their work much more accurately than is commonly done, focusing especially on those issues they themselves do not write about. Incidentally, this situation is not unique to OR, but appears to be a general phenomenon in many professions. In this context, Argyris & Schön (1974) refers to the discrepancy between 'espoused theory' (what do practitioners say they do) and 'theory in use' (the theory that can be inferred by observing what they really do). Much of what practitioners do is in fact based on tacit knowledge, making it difficult for them to express. To prepare students for practice, we need to teach them about the 'theories in use' of successful practitioners, but what we actually teach them are the 'espoused theories'.

Let us quickly look at how all this relates to three issues of concern to many within the OR community: the 'crisis debate' (the subject of Chapter 18), the increasing marginalisation of OR departments in many business schools, and the fact that so many OR graduates abandon their field so soon after graduating.

The crisis in OR. Authors contributing to the crisis debate during the last five decades have done the profession a service by calling attention to the need to emphasize OR practice more heavily, resulting in several laudable initiatives. On the other hand, the resulting polarisation between theory and practice has tended to overshadow the fact that what is really needed is a theory *of* practice. The literature is replete with statements that OR practitioners need to possess (more) social skills to reduce the frequent communication, political, and other problems that occur in OR projects. But, if we are to teach students these skills, some formalized body of knowledge is needed about these skills, an 'epistemology of practice'.

Marginalisation of OR departments. There is a trend among business schools to shift the emphasis in their teaching away from quantitative skills, resulting in downsizing and closing of many OR departments. Generally, MBA students are far more likely to end up as the client of an OR project than as the consultant. From the 'projects discussed here, we see that the client knowing about the techniques of OR is not a necessary condition for a project to be successful (though it may sometimes help). At least as important is that the client have a good understanding of how to manage the process by which the project is executed, i.e. precisely those issues not normally described in OR cases. So, richer descriptions of cases are a necessary (but not sufficient) condition for training future practitioners, as well as for educating future clients!

Apprenticeship in OR practice. Lastly, we should emphasize that we definitely do not claim that OR practice is something that can be learnt from a book. Good descriptions and a good theory of practice may help a lot, but experience will always remain an essential learning mechanism. One approach often employed in management consulting firms is that of apprenticeship: by pairing a junior consultant with a senior consultant, the former can quickly learn a lot from the experience acquired over the years by the latter, and a lot of typical beginner's mistakes can be avoided. People who have to start practicing OR as an individual without such an apprenticeship, such as the Lone Rangers, can easily, by making a few minor mistakes in their first projects, turn the tide in their organisation against them, leading them and their organisation to abandon OR much too quickly.

The need for serious research on the practice of OR. To conclude, we strongly believe that there is an urgent need for further, serious research on the practice of OR. One reason why such research is rare probably lies in the fact that the subject matter and the appropriate methodologies for such research are generally not those for which people within the OR community are well-equipped to deal with, while scholars in other disciplines

with the necessary skills generally have no particular incentive to apply those skills to studying OR. Although trying to bridge two gaps at once, that between theory and practice of OR and that between OR and other disciplines, is a challenging prospect, in our experience it has proved highly rewarding and stimulating, and we hope that many others will have similar experiences.

Acknowledgements

We would like to warmly thank the OR consultants and their clients involved for their willingness to discuss these projects with us and their openness in doing so. We hope that this chapter reflects something of the enthusiasm with which they conduct and talk about their work.

References

Corbett, C.J. & L.N. Van Wassenhove. 1993. "The Natural Drift (What Happened to Operational Research?)", *Operations Research*, Vol. 41, No. 4, July-August, pp. 625-640.

Corbett, C.J., W.J.A.M. Overmeer & L.N. Van Wassenhove. 1995. "Strands of Practice in OR (The Practitioner's Dilemma)", *European Journal of Operational Research*, forthcoming.

Fortuin, L., P. van Beek & L. Van Wassenhove. 1992. "Operational Research Can Do More For Managers Than They Think!", *OR Insight*. Vol. 5, No. 1, January-March, pp. 3-8.

Argyris, C. & D. A. Schön. 1974. *Theory in Practice*. Jossey-Bass, San Francisco.

Table 1. Overview of projects and people involved

client organisation	people interviewed	consultant interviewed
Compeco (multinational electronics company)	J. Smits (project leader) J. Van Gool (project team member)	H. Fleuren (CQM)
<i>Project: simulation studies to aid in the design of a new plant for coloured liquid crystal displays.</i>		
Callebaut Europe (European industrial chocolate manufacturer)	F. Callebaut (director of operations)	L. Fortuin (CQM)
<i>Project: developing and implementing performance measures.</i>		
Monsanto Services International N.V. (multinational chemical company)	F. Cammaert (logistics manager Europe-Africa)	A. Van Looveren (Beyers & Partners)
<i>Project: developing a system for supporting agricultural product sourcing decisions.</i>		
Oilco (multinational oil company)		D. Bookman (BOR Consultants)
<i>Project: developing and implementing optimisation software for cutting packaging foil.</i>		
RGD (Rijksgebouwendienst) (government housing agency)	C.L.G. Wassenaar (research and development)	C.G.E. Boender (ORTEC)
<i>Project: development of a system for decision support in allocation of housing to government agencies.</i>		
Transco (transport company)	W.P. Gray (Head, vehicle maintenance department)	P. Brown (ABC Consultants)
<i>Project: development of a decision support system for medium and long-term manpower planning decisions.</i>		
AMEV Levensverzekering N.V. (insurance company)	D.A. Voûte (adjunct director) C.W. van Dedem (project team member)	C.G.E. Boender (ORTEC)
<i>Project: development of a system for asset liability management</i>		