

**STRANDS OF PRACTICE IN OR
(THE PRACTITIONER'S DILEMMA)**

by

C. CORBETT*
W. OVERMEER**
AND
L. VAN WASSENHOVE†

95/44/TM

* PhD Student, at INSEAD, Boulevard de Constance, Fontainebleau 77305 Cedex, France.

** Assistant Professor of Strategy and Management, at INSEAD, Boulevard de Constance, Fontainebleau 77305 Cedex, France and Stern School of Business, New York University, USA.

† Professor of Operations Management and Operations Research, at INSEAD, Boulevard de Constance, Fontainebleau 77305 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

Strands of Practice in OR

(The Practitioner's Dilemma)

Charles J. Corbett¹

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

Luk N. Van Wassenhove¹

¹ INSEAD, Fontainebleau, France

² Stern School of Business, New York University

May 5, 1995

Abstract

This paper is a first step towards a different way of looking at the practice of OR, with implications for practitioners and researchers alike. We make two key observations: (1) Looking at how OR practitioners perform individual projects will *not* lead to a full understanding of how OR practitioners perform individual projects. Instead, one has to take into account the series of preceding projects by that same practitioner, to infer and understand his 'strand of practice'. (2) A practitioner's strand of practice is developed as the result of a learning process. Therefore, practitioners operating in different environments can have widely different strands of practice. Consequently, operationalizable guidelines about how to perform OR projects will have to include an indication of the strand(s) of practice for which they are valid. Our observations are based on an exploratory series of interviews with practitioners and their clients.

1 Introduction

Among those concerned with effective practice of operational researchers, few would disagree that when such a practitioner conducts an OR project, he should be problem-oriented. That is, he should be able to enter a new situation with an open mind, attend to the unique characteristics of the situation at hand, and devise a customized solution to that situation by using whatever techniques are best suited. In this paper, we will argue that this prescription for a problem-oriented approach, on a case-by-case basis, is too limited a view of OR practice, and, if followed by novice professionals, leads to what we call the *Practitioner's Dilemma*. In its most extreme form, it would argue that an OR practitioner should be able to address any problem for any client using whatever technique is called for. When faced with the realities of actual practice, this has become a nearly impossible requirement. It may have been feasible in the early days of OR, but the vast expansion of the field and increased competitive pressures have changed that.

We will show how different European practitioners have found their own resolution of this dilemma by various degrees and directions of specialization, be it with respect to the tools they use, the class of problem they tackle, or in other ways. These specializations result in specific competences, accumulated over a series of projects, referring to ways of finding and framing problems, to planning projects, and to interacting with the client organization. In this way, different practitioners develop their own *strand of practice*, that allows them to respond adequately to the demands of *their* clients. (We will define the term 'strand' in Section 2.)

In our study, we saw one practitioner who focused on developing heuristics for combinatorial optimization problems in a wide variety of domains; we saw another who developed a powerful LP solver and focused on selling customized implementations of his system; we saw a third who acquired a lot of know-how about production processes for electronic components and about using simulation to study such problems; and we saw a practitioner who linked a series of projects into a larger product development effort. And there is no reason why other practitioners may not have developed their own strands of practice. These four practitioners clearly differ in the extent to which they are oriented primarily towards the case at hand, the tool to be used, the problem domain, and the development of a new product.

In this paper, we depart from traditional thinking and writing about OR practice in two fundamental ways. First, we argue that an OR practitioner's work cannot be understood by looking at the projects he performs as independent entities. Those projects are strongly interrelated, and the practitioner's approach and his success or failure in any one project can only be understood by looking at the series of projects that form that practitioner's prior experience. Second, OR practitioners develop different strands of practice; as a result of facing other circumstances, they develop very different types of competences. The consequence is that one cannot make normative statements about 'OR practice' in general, but that such statements must be qualified by adding to which strand(s) they apply. Much of what has been written about OR practice may be true and important for one particular strand, but less relevant or even misleading for another. In fact, the literature on OR practice can even provide a misleading and internally contradictory view, as we explain in Section 2. This is not to say that an 'epistemology of practice' is an impossible prospect, but that such an epistemology must acknowledge the existence of strands of practice and be sufficiently general to encompass whatever strands occur.

The ideas of this paper are based on a series of interviews with successful OR practitioners and their clients. In Section 3 we describe four projects, each of which belongs to a different strand of practice. We then shift from an essentially static view, focusing on individual practitioners and projects, to a dynamic view of series of projects carried out by consulting firms. In Section 4 we discuss how and why each of the practitioners involved has developed their own specific strands of practice, and examine the development of these strands as learning processes. Recognizing the existence of different strands of practice has far-reaching implications for the way OR projects are performed. In Section 5, we return to the project perspective, and analyse (in a speculative fashion) some of these implications by disentangling four project approaches underlying the strands of practice described. In Section 6 we examine the Lone Ranger issue in the light of our findings. Finally, in Section 7, we emphasize the limitations of the current work, and suggest areas for further research.

2 The prevailing view and the Practitioner's Dilemma

When novice professionals are introduced to the 'practice' of operational research during their extensive academic training, they are usually presented with two persistent images: projects are either centered around the application of an OR technique to a real-world problem, or they are seen as the step-by-step progression through a problem solving cycle by an operational researcher. A gathering of leading figures expressed a similar distinction between two prevailing opinions of systems science, in Advanced Research Institute (1979). From reading many of the examples of excellent applications of OR such as the Edelman finalist papers (published in *Interfaces*), a similar impression arises. The following quote from Royston (1993, p. 16) nicely illustrates these two views:

For instance a classification of our work in line with a 'general problem cycle' of the sort that is familiar to managers - for example:

- scanning for emerging issues
- formulating and diagnosing problems
- setting goals and objectives
- designing and developing options
- appraising options and making choices
- gaining acceptance for and implementing solutions
- monitoring and controlling implementation
- evaluating results

seems to give rise to a more marketable agenda for OR than one based on tools of the trade or narrowly technical problems.

Both images have in common that they see the practice of OR as a strictly project-oriented activity, in which the practitioner enters each new situation with an open mind and develops an approach from scratch, pulling the appropriate tools out of his bag when necessary. This goes back to the classical tradition of OR as it came out of the WW II effort. It is also the prevailing view of what OR should look like, as witnessed by the many authors adhering to the 'back to the roots' movement.¹ Rosenhead (1989, p. 349), who also criticizes this view, describes it as "Each problem, it is asserted, should be tackled without preconceptions, or

¹This view is commonly juxtaposed with the derogative view of an OR practitioner as a boy with a hammer for whom everything looks like a nail.

pre-investment in technical apparatus which may bias the approach.” In his contribution to the debate on a better name for OR, Royston’s (1993) suggestion of ‘general analyst’ fits with this trend. These two images form what Argyris and Schön (1974) would call the ‘espoused theory’ of OR/MS practice. What we are arguing is that this espoused theory of practice differs, in potentially misleading ways, from a ‘theory in use’ that can be constructed on the basis of observations from actual practice.

In order to explore the key problem areas in OR practice, we interviewed seven practitioners about a project of their choice. Afterwards we also interviewed clients involved in the projects. Six of the practitioners came from three (highly) successful OR consulting firms. The firms employed in the range of 30 to 70 people and had existed as a separate organization for between 2 and 15 years. They have extensive contacts with universities, frequently to the extent of practitioners simultaneously holding a (part-time) faculty position. The ideas reported in this paper are largely inspired by the interviews but have been supplemented by our prior knowledge of the firms through various contacts over a considerable period of time.

Based on this series of interviews, we found that projects were characterized by a far higher degree of complexity and diversity than suggested by the two images above. Our practitioners readily agreed with the inaccuracy of the prevailing images and the complexities of carrying out OR projects. Yet they too were hard pressed to provide a better image of their own work. Although we had rich descriptions of a variety of cases, we found ourselves unable to move towards formulating more general guidelines for OR practice without being drawn into non-operationalizable platitudes such as “communication is important” or “data gathering can be difficult”. This, we now realize, was due to our initial focus on individual projects.

We tried to account for how practitioners were able to carry out projects with shifting problem definition, complex project management, and with widely different types of clients. Only as we reflected more about their respective firms, we came to see that their competence in carrying out OR projects could not be understood in terms of individual projects but as evolving over multiple projects within specific company environments, as a sometimes entrepreneurial and strategic response to the changing conditions each faced. In looking at a series of projects carried out by any one practitioner, we observed that many of them shared a number of common characteristics, constituting a distinct approach to OR practice. These

common characteristics are unique to each practitioner, and define his *strand of practice*. A practitioner may in fact have more than one strand of practice, as we will see later. By viewing all of these practitioners at once, we saw that they had each found their own response to the predicament faced by any OR practitioner – if he were to follow the prescriptions implied by the recommendations for effective practice – how can an OR practitioner be all things to all people? This predicament is what we call the *Practitioner's Dilemma*.

Our answer is that he or she cannot! Faced with demanding clients, fierce competition, and dependent on specific available resources (initially often universities), successful practitioners often developed a strand of practice; that is, they developed a distinct set of competences or know-how, tightly integrated with specific OR knowledge, that allowed them to build on previous projects whilst performing new ones, without having to start from scratch. Our four cases display four such strands of practice. Each of these strands can be viewed as a movement towards resolving the Practitioner's Dilemma. By focusing on a strand, practitioners accumulate experience within that particular strand, and in doing so reduce the uncertainty implied by the Practitioner's Dilemma. The strand or strands developed by practitioners can then be seen as a learning process in response to the changing circumstances they face.

The realization that some practitioners do not see problems as being independent of one another, but attempt to carry over knowledge or solutions to new projects, has been made before by eg Weingartner (1987) and Mitchell (1993). As Keys (1989, p. 113) states: "It is part of any attempt to improve future practice to reflect on past experience." Keys (1989, 1991) presents a systems model of an OR group, and discusses various ways in which the projects within a group interact. To illustrate that an OR group's projects are indeed often interdependent, Meyer zu Selhausen (1985; p. 199) found, in an empirical study of OR in Germany, that in 57% of the cases "more or less routine work was done in a given problem area and for given clients".

There have been earlier papers suggesting that OR/MS practitioners can play a variety of roles (eg Eilon 1974, Batson 1987, Fortuin & Korsten 1988, Bowen 1990), and others outlining intellectual paradigms within which OR/MS analysts act (eg Zeleny 1979, Pidd 1995). Our work differs from and adds to this in important respects. Our contribution lies in our hypothesis on how and why strands of practice develop. Also, this paper is research-based

(albeit exploratory), as opposed to having the more philosophical basis of many typologies. Meyer zu Selhausen (1985) forcefully argues the need for careful empirical research, the results of which “might be an *ideal starting point for a fundamental re-thinking* of the concept of OR as a whole and the OR process in particular” (p. 202; author’s italics).

3 Four projects

In this section, we describe four projects by four practitioners, representing four strands of practice. This section is at the level of projects and practitioners; in the next section, we shift to consulting firms and series of projects. The projects described below are, to a limited extent, typical of the type of project each of these practitioners perform most frequently, but the correspondence is far from perfect. In fact, each practitioner has also successfully performed different types of projects than the one described below. To aid in reading the cases, Table 1 provides an overview of the projects.

<i>Consultant and firm</i>	<i>Project</i>
A (ADOR)	decision support for tactical manpower planning
B (BOR)	implementation of a customized system for cutting stock problems
C (COR)	simulation studies to support design of a coloured LCD production line
D (ADOR)	development of an asset-liability management system

Table 1: Overview of projects and people interviewed

3.1 Consultant A: decision support for tactical manpower planning

Client A, then head of the 1200-strong maintenance department of a very large transportation firm, has an engineering background, and tries to apply analytical reasoning whenever he can. Consultant A, of ADOR, already had experience with developing and implementing planning systems of various types, but did not have any prior experience working in this particular type of environment. When he contacted the client, the latter suggested that he go and interview

the subdivision heads to see where opportunities for applying OR lay. The consultant returned with a list of possible projects; they agreed to work their way down, starting with the project promising the highest return. This was to develop a tactical system to help the client with manpower planning; it would allow him to determine how much staff he would need for a given set of activities in his division, for given work schedules, for given collective labour agreements, etc. The system would have to be developed from scratch.

Both parties preferred an informal, flexible mode of working; project planning was, accordingly, kept to a minimum. Consultant and client got along well, and the client trusted the consultant to come up with good results. 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”. The consultant’s view fits well with this: “We usually don’t build systems the way a computer scientist would advise, ie defining upfront what the system should do and adding increasingly more details.”

Unfortunately, by the time the system was completed, the client had moved to a different division. His successor did not share his interest in OR, and could not find any precise plans describing the ongoing project. Projects planned for the future were cancelled, and the consultant was relegated several levels downward in the hierarchy to complete the manpower planning project. However, staff at that level had different concerns than client A. For them, labour agreements, duration of lunch breaks, etc, were a given, and considered unchangeable. They wanted to use the system’s graphical output to quickly check how much work the division had lined up for the coming period, or to determine detailed work schedules. In the original system, detailed schedules were relevant only as a means for evaluating higher level decisions. Accordingly, good rather than optimal schedules had been sought after, and the consultant had to revise the system for this new use. The system is currently operational, but for a lower-level purpose than originally intended.

3.2 Consultant B: implementation of a customized system for cutting stock problems

Some staff members at client company B, an oil company in Europe, decided they needed customized software for their packaging foil cutting machines. Consultant B's firm, BOR, was already building a reputation for their experience with cutting stock problems. Company B staff visited BOR to discuss the possibilities, talked it over internally, and returned a year later. A more primitive system was in use within company B in the US, but, with some effort, the Americans were persuaded not to impose that system on Europe; permission to do a project with BOR was granted.

The project started with two pilot studies:

- A functional pilot study, to define the problem in detail, and to specify what the software should be capable of.
- A more organizational pilot study of hardware and software aspects such as how to build interfaces with existing systems, who should install which database, on which hardware, what type of output the system should generate, etc.

After the first pilot study, a presentation to general management was made; then a contract was drawn up for the main body of the project, distinguishing a number of stages. Although much of the software had to be developed from scratch, the system was built around BOR's flexible LP software. The 18-month project was tightly planned and smoothly executed.

Cooperation with planners, the future users of the system, 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. It was the consultant who had to coach the planners along, for example by being on site during the first days after the actual changeover, and by installing a modem link with the BOR office. The consultant summarized his coaching efforts as "I'll jump first, you just follow me." Several of the client's project team members changed during the project, but despite that they only deviated from the initial plan by two weeks. In consultant B'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.”

3.3 Consultant C: simulation studies to support the design of a coloured LCD production line

In 1988, client C became leader of a company C project team responsible for designing a coloured liquid crystal displays plant. This was the first such plant in Europe, so no experience with the production process was available. The project team consisted largely of engineers, each responsible for a part of the total production line, several hundred process steps in all. 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.

Client C had no previous experience with OR. A mechanical engineer, he had been involved in determining ram and pile tensions using simulation; as a result, he recognized that the static spreadsheet calculations performed gave insufficient insight into what were inherently dynamic processes. So, he asked COR, formerly an internal OR consulting group, to perform some simulation studies. Consultant C could contribute both his simulation expertise and his experience in modelling production processes from previous projects. The client first commissioned a small pilot study, and then a simulation of the entire line, which took some 50 days.

The consultant was then called in three more times for various small projects on subprocesses and, in 1992, when the plant design was nearly complete, for another simulation of the entire line. He remarked that the client and he had an efficient way of framing the problems: the client would call him to talk about some subprocess that needed detailed analysis, he would write a proposal describing the problem, the approach he had in mind, and the time needed. That way, the client could check whether the consultant's interpretation of the problem was correct.

The simulations led to a number of important insights for the project team. They brought out, among others, the drawbacks of too carefully balancing a line, the importance of maintenance staff location, 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 Consultant D: development of an asset-liability management system

The development history of the asset-liability management (ALM) system started around 1984, when a Dutch bank hired consultant D, from ADOR, to help develop a liability management system for the bank's own pension fund. When the system was completed, the consultant saw that, expanded to include a wider variety of assets, it would have great potential. Because of legal and economic changes, pension fund managers needed a new type of tool to evaluate investment strategies.

Consultant D then further developed the ALM system with 3 partners, including an insurance company where client D worked in the liability department of the pension fund management division. Client D set up a project team which started on the basis of the consultant's prior work. The intention was to combine the consultant's OR knowledge with the insurance company's knowledge of pension fund management, in order to develop a more powerful system for asset liability management. The consultant intended to sell the system to various pension funds, the client wanted to use it as part of the service they offered smaller pension funds. Most of the company's project team members knew nothing about OR.

Initially there was no project planning, it started as an unorganized joint venture. Consultant D was actually not operating as a consultant to the insurance company, and several company staff members invested a lot of their time in the project. During the first serious meeting, the team estimated the development would take about half a year. Eventually, "it took half a year or more just to understand each other", as a team member put it. A key success factor was the team's enthusiasm, particularly given the unofficial nature of the project; for most people, the work came in addition to their normal workload. The team member: "We found the subject so interesting that we all got carried away, we all became as enthusiastic as [consultant D]". They perceived the project as an intellectual challenge. This enthusiasm arose despite the initial communication gap; it took the team more than six months to reach a "common frame of reference".

Choices as to what to include and what to leave out, eg foreign currency issues, were made

continually throughout the project. Initially, the team became more and more ambitious, wanting to incorporate everything, but at some point in time, client D said, they “came back to Earth, both feet back on the ground”, and started throwing things out. The client decided in November 1991, well into the project, that some time pressure was needed. After taking stock of what still had to be done, a May 1992 deadline was set. From then on, everything was tightly planned, tasks were assigned week by week, people reported back and forth; they slipped from that plan by less than a month.

4 Four strands of practice

In this section, we change our unit of analysis from practitioners and projects to series of projects. Here we show how practitioners develop strands of practice, that can be interpreted as responses to the specific circumstances they face. The descriptions of the firms given here correspond to actual cases, but they have been modified where necessary for reasons of confidentiality or clarity of exposition.

4.1 ADOR: academics turned practitioners

Consultants A and D co-founded ADOR in the mid-80’s whilst pursuing PhD dissertation research, in order to get involved with more practical OR. Their first job was a small logistical application; a second project followed soon. With no real marketing efforts, orders kept coming. ADOR’s projects often involved developing and implementing tactical or operational systems, built around a limited set of heuristics for combinatorial problems. ADOR has grown to some 40 employees, including many junior consultants, recent graduates or students of OR. Projects still frequently originate through word-of-mouth advertising; once consultant A is ‘inside’ a company, one project often follows another. Apparently, nobody who hired consultant A has ever been dissatisfied, but others (as client A’s successor) are not to be persuaded of the value of OR.

ADOR had little or no real competition, and that country’s companies were readily prepared to spend on OR consulting projects until relatively recently. The projects are many but never

very large, rarely more than \$100,000. Much of the firm is highly informal, both internally and with clients. So far, consultant A has had little reason to deviate from his predominantly case-by-case approach to performing projects. His only complaint is that he gets to perform so few 'studies', projects involving strategic questions rather than operational planning systems.

ADOR's continued success shows that this approach works well for them. However, it is not without its dangers: for getting projects, they rely heavily on good personal relationships with clients. There are no clear specific competences with which they can market themselves, especially to a client who does not know about OR. Moreover, they will generally only be called in for projects for which no specialized firm is readily available; if there is a specialist (whether or not an OR consultant) for that problem, ADOR can hardly hope to compete. In extremis, their market niche may well become increasingly marginalized, although there are no clear signs of that happening yet. The next strands, with BOR and COR, illustrate two ways of getting round that problem of how an OR firm may be identified by a prospective client as offering the service she needs.

4.2 BOR: from selling software to application-based consulting

The origins of BOR are closely intertwined with the development of the LP solver LPS by consultant B and a colleague, both professors at the same university. For some time, they tried to sell LPS directly, but with little success, despite the system's potential. After having performed a project which involved developing software for a cutting problem, based on LPS, they started selling customized applications rather than raw OR software. This included taking responsibility for implementation of their systems within the client organization. Sorting out the interfaces with the client's existing hardware and software and training and motivating future users became crucial parts of BOR's activities, despite some of these issues often not being contractually required. One reason for this is the relative reluctance shown by firms in their country to invest in studies; they demand complete, operational solutions.

After performing several cutting projects, they gradually acquired a reputation as specialists in that subfield. They have to compete with fly-by-night amateurs selling 'cutting optimization software' for one tenth of their price. The quality of BOR's software is generally incomparably

better, but this is not obvious to prospective clients. They also have to offer superior consulting and implementation services. Accordingly, in their dealings with clients, BOR often leave a more conventional business-like impression than does ADOR. Their emphasis on software implementation has resulted in employing several people specialized in hardware and software-related aspects of computer science.

BOR have started looking for other applications of LPS. They are frequently involved in strategic studies on rationalization in the dairy industry, determining optimal product mixes for a given cooperative. Although seemingly very different from cutting problems, they both rest heavily on LP. One of BOR's main concerns is how to keep their relatively specialized (compared to ADOR) people, with skills in eg computer science, usefully employed and motivated.

4.3 COR: from internal to independent consultancy bureau

COR (also described in Fortuin & Zijlstra 1989) originated as the Centre for Operational Research within company C. Under its first director, who was well-connected throughout the company, they established a reputation as the wizards who could help with any problem. There are three specializations: operational research and material handling (which includes consultant C), statistics, and forecasting and planning support; we focus on the first. COR have close links with a technical university, with PhD students working on COR projects during their studies and joining the firm upon graduation. Several COR employees are also part-time faculty members.

Often, COR becomes involved when a client realizes she has a problem with an existing line; their task is then to pinpoint the problem and suggest solutions. COR has accumulated a wide experience with material handling systems, and have become expert users of simulation. They have their own specialized simulation-related software to enable them to perform projects more quickly.

In general, COR projects are studies. They analyse a problem and write a report, giving the client more insight but not actively implementing a solution, though the latter may well change in the future. Once COR has diagnosed a problem and found that software needs to

be developed and implemented, their involvement in a project is often over.

COR started as a cost center within company C, became a profit center, and recently have become independent. Though company C is still their largest customer, the proportion of external projects is increasing.

4.4 Consultant D: from consulting for clients to co-development with partners

Consultant D follows the ADOR strand of practice described in Section 4.1, with the notable exception of the development of the ALM system. The key difference came when he realized, after performing the initial bank project, that pension funds provided a promising market for OR applications. To work his way into this world, he set up joint ventures to develop a marketable product. His partners paid him just for the original bank system; from then on, it was a joint venture, rather than a client-consultant relationship. Much of his university work, and that of his students, was devoted to developing the theoretical background.

At the same time, consultant D continued to work on ‘conventional’ ADOR projects, to keep money coming in. It was only after a long period and an enormous investment of his and his juniors’ time that the first ALM system was sold. Such a large investment with no direct payoff is a flagrant deviation from ADOR’s dominant case-by-case approach, based on getting each hour worked paid by a client. Now, the system is repeatedly customized and sold, the original investment paying off handsomely. Clients include pension funds with no particular knowledge of OR and no previous experience with ADOR. The competition with consulting firms traditionally more closely associated with pension funds is intense, but ALM’s technical superiority and consultant D’s gradually acquired expertise has enabled ADOR to corner a significant share of the market. Sales are often made on the basis of references from previous clients; some pension funds even buy the system without having seen it.

The notion of ‘project’ becomes ambiguous in this strand; it can refer to customising the ALM system for an individual customer, but also to the entire co-development of the system with and selling it to a series of customers. (It was in fact this ambiguity that opened our eyes to the concept of strands of practice.) Perhaps the key bottleneck here lies in the entrepreneurial

abilities needed to pull it off. Consultant D firmly believes ADOR could follow a similar approach with several other systems, developing them into more widely applicable products and then selling customized versions to multiple clients.

4.5 Strands of practice as a learning process

Let us re-emphasize that an OR firm's strategy is not necessarily composed of a unique strand, nor does any practitioner adhere exclusively to one particular strand. For example, COR has various divisions, each with their own specialization, while BOR also offer services based on their financial planning software; we have focused on a single division. Moreover, within ADOR, consultant D generally followed a similar approach to projects as consultant A; only with the ALM did he deviate from that approach. By definition, however, any practitioner can develop only a very limited number of strands.

The question of interest here is: how and why did those strands arise? What we suggest above is that these strands each represent a practitioner's way of getting around the Practitioner's Dilemma. As practitioners and consulting firms gain more experience with performing OR projects, they learn which approaches work for them, under which circumstances, with which type of project and client, etc. In several cases, the firms are obliged to modify the strand of practice underlying their strategy as they learn more about market pressures. Eg, BOR set out to sell a powerful LP solver, but found they had to shift towards selling powerful applications (with consulting) built around that solver; consultant D saw that by working on a purely case-by-case basis, ADOR was missing many valuable opportunities to sell results of past projects to other clients.

In other words, each strand of practice can be seen as a practitioner's learning process. It determines that practitioner's competitive positioning in the market for services. It determines whether the practitioner finds himself competing with other OR consultants (or even academics and students eager to perform practical projects at very low cost), with general management consultants, with software houses, or with yet other providers.

This learning process may, but need not be, a deliberate, conscious one. Consultant D may have seen that the system developed with the bank had potential to be expanded and sold

elsewhere, but an important reason for his seeing that was that pension fund legislation just happened to be changing, focusing attention on such issues. BOR may have consciously realized that selling an LP solver as such was not the way to go, but to what extent their subsequent focus on cutting problems was the result of a deliberate decision rather than the coincidence of their being presented with one or two cutting projects is not clear. This development of a strand of practice resembles what Mintzberg (1989) refers to as ‘emerging strategy’, a series of actions from which an implicit long-term strategy can be inferred, but which are not the result of such an explicitly formulated strategy.

OR groups can also deliberately cultivate a variety of strands, allowing them to compose project teams covering the range of skills needed. Looking back on his experience with a 25-man office in the US Air Force, Miser (1995b) realized they had various strands, some methodological (statistics, programming, etc.), others related to aspects of air force operations (air defense, strategic air power, etc.). In truly large-scale OR or systems analysis projects, appropriately combining various strands is key (Miser 1995a).

The strands each possess their own internal inconsistencies; no strand is perfect. For instance, consultant A’s approach works well with some clients (such as client A), but poorly with others (such as his successor). BOR, heavily depending on powerful in-house systems, are faced with the challenge of developing these systems, with very limited resources, while performing other projects to keep money coming in. By focusing on one or more domains, COR potentially restricts the size of its market, and may have to reach out to other domains, again contrary to their strand of practice. And finally, consultant D’s approach relies on finding ‘customers’ prepared to and capable of co-developing a system rather than buying a turnkey system from a consultant; such clients may well be few and far between. Although each of the strands represents a way of circumventing the Practitioner’s Dilemma we started out with, they also carry within them the seeds of new practitioners’ dilemmas. Our data do not allow pursuing this point here, but we intend to address it in future work.

5 What strands of practice mean for projects

We have seen four strands of practice, and how these have emerged over time as a response to four practitioners' circumstances. These strands clearly differ fundamentally in several respects, which is reflected in the practitioners' project approaches and corresponding problems. The key point here is that practitioners have each developed their own strand of practice, different from one another but each a viable response to their specific circumstances. This could go a long way in explaining why Hildebrandt (1980), in his useful survey of implementation research, found that no general guidelines emerged from the normative literature, and that most empirical work failed because it ignores precisely that message – (p. 8) “that implementation and its attendant difficulties are to a considerable extent situation dependent”.

In this section, we attempt to better understand the strands described above by disentangling their key characteristics. To do so, we will describe a caricature of four hypothetical project approaches, by heavily simplifying and exaggerating these key characteristics. The level of discussion returns from strands of practice, as in the previous section, to individual projects. We first describe the four project approaches, then we analyse the implications for the problems OR practitioners face in performing individual projects. An important note is that this section is essentially speculative in nature, as opposed to the previous sections. The project approaches and the subsequent discussion may have been inspired by the data outlined earlier, but they are by no means proven by those data. The project approaches constitute a step towards a framework which may help in the fundamental re-thinking of the OR process called for by Meyer zu Selhausen (1985).

5.1 Hypothetical project approaches

It is useful to disentangle the strands for a moment and focus on four 'idealized' project approaches that should in fact be viewed as extremes of a multidimensional continuum. Of course, these extreme forms will not occur as such, but are elements which can recur to a more or lesser extent in any practitioner's project approach. The practitioners we interviewed combine elements of each of these approaches in often highly sophisticated ways; they are far more competent and professional than the hypothetical caricatures that follow.

- *Case-by-case orientation*, the type most closely related to what is often referred to as ‘classical’ or ‘old-style’ OR: the practitioner moves from case to case, hopping from one domain to another, and using a wide range of tools to arrive at a solution. He is solely oriented towards the project at hand.
- *Tool orientation*, in which the OR practitioner develops a tool, and then focuses exclusively on selling and implementing that tool. An example could be an LP solver, which can be used in widely different domains.²
- *Domain orientation*, where the practitioner performs projects exclusively within one domain, using whatever tool is appropriate for each individual project. An example could be production lines for electronic components, where, depending on the precise question, simulation, queueing network analysis, or LP could be the appropriate tool.
- *Product development orientation*, where the practitioner intertwines performing projects and developing an OR product. In each project, the product is expanded or improved upon according to the current client’s requirements and in close collaboration with that client; over time, a complete and marketable product is developed.

In the Introduction we mentioned ‘problem orientation’. As this need not mean the same to everyone, we deliberately do not use the term elsewhere in this paper. We should re-emphasize that other strands than the four we describe can be found, and from them other project approaches inferred. The above orientations will not occur in real life in the pure form described above, but these orientations recur to a certain extent in the strands. of Section 4. In the first strand, consultant A does not restrict himself to any particular domain, but does tend to formulate problems in mathematical programming terms and has a predilection for certain heuristic solution methods. He essentially follows a case-by-case approach, but with an element of tool orientation. In the second strand, we saw BOR developing a powerful and flexible LP solver. On this basis, they developed a system for solving cutting problems, and customized and sold this to various clients. They also found other applications for their LP software. Here, tool orientation is clearly present, but combined with a focus on some limited classes of problems in several domains. The third strand is that of consultant C, whose

²Rosenhead (1989, p. 6) writes of this: “Some consultants have carved out a niche as purveyors of particular techniques. For external consultants this offers a marketable commodity to be exploited by repeat sales. For internal consultants it offers a reliable downhill path to a role of minor technical auxiliary. In either case the original OR orientation has been turned on its head.”

generally performs studies of production lines (for electronic components), and who often uses simulation to gain insight into these processes. This is a predominantly domain-oriented approach, combined with a degree of tool orientation.

The fourth strand, that of consultant D, is less easily characterized, but is oriented primarily towards new product development. It starts with the choice of a particular domain (in this case, asset-liability management for pension funds), and no preconceived notion of which tools should be used. Step by step, a tool is developed to tackle projects within that domain, until a system is available which can be sold (after customization) to various clients. There is a shift from no tool orientation to a high tool orientation. Also, initially the domain is fairly broadly defined, but by performing the joint development projects, it gradually becomes clearer what are relevant problems for clients; the domain is narrowed accordingly, until it is defined as 'the class of problems for which the tool can be useful'.

So what might these project approaches mean for the way OR practitioners perform individual projects? Let us speculate about how they could interact with the key issues we identified elsewhere (Corbett, Overmeer & Van Wassenhove 1995): problem framing, project planning and project uncertainty, and interacting with the client organization.

5.2 Problem framing

Before a client will engage a case-by-case-oriented OR practitioner, she must have framed her problem as an OR problem, or at least be convinced that OR can be of use to her. Frequently, the client does not have a particular problem in mind, just a 'mess' or a feeling of unease. In these cases, the practitioner has to be able to enter a previously unfamiliar domain and find the root cause of the problem, before deciding on an approach to tackle it. A case-oriented practitioner will not hesitate to tell a client that the cause of her problem in fact lies somewhere else than the symptoms had her believe, and he should be able to help her solve that problem. However, the practitioner being unfamiliar with the problem domain can easily lead to often unnoticed misunderstanding about what the client really wants. We have seen an example of a client and consultant working together closely on developing a software system, but only when the client suggested to define the system's output format did they discover a subtle but

important difference between client's and consultant's problem interpretation.

For different reasons, this risk also confronts a tool-oriented practitioner: once the client has decided her problem can be solved with that practitioner's tool, it will be difficult, for the client, for the practitioner, or both, to accept that that tool may lead to a marginal improvement, but that some other change may be far more effective. For instance, once the client has decided her cutting operations are not efficient enough and contracted a consultant specialized in cutting software, the latter may quickly realize that the real bottleneck lies not in the cutting software but in the production line design. By specialising around a particular tool and a corresponding collection of narrow classes of problems, the practitioner relies on the client to frame her problem as one that can be solved with precisely that tool. The consultant will then be visible to the client as having the necessary expertise and software; the danger is, however, that clients frequently do not recognize the true causes of their problems. And indeed, consultant B remarked that clients often call his firm with some operational manifestation of a problem, the underlying cause of which is frequently at a deeper, more strategic level.

A domain-oriented practitioner relies far less on the client for correct problem framing. The wide experience and deep domain knowledge is precisely his primary asset which allows him to quickly and accurately diagnose the cause of any problem within that domain. A client wondering why throughput in her assembly line is so low, or with any other production logistics problem, can rely on COR to rapidly pinpoint the source of her problems. Indeed, once a client has diagnosed her problem more precisely (whether correctly or not), eg as a problem with the efficiency of her cutting machines rather than just one of low throughput, she may well bypass the domain-oriented practitioner and seek out a (tool-oriented) specialist. She may also do so once the domain-oriented consultant has diagnosed the problem for her. This is not to these practitioners' advantage: clients have used COR reports to serve as problem description for software houses to work from, and a COR consultant observed that there is often a lot more money involved in developing and implementing a solution than in identifying the problem.

5.3 Project planning and uncertainty

A tool-oriented practitioner's projects essentially consist of (lightly) adapting his system to client requirements and then implementing it. Doing this repeatedly allows the practitioner to plan the project tightly, and to accurately estimate cost and duration in advance. Because the problem is already relatively narrowly defined (otherwise the client would not have contacted a tool-oriented consultant; see above), there is little uncertainty in the project; the client knows what she wants and gets it.

The domain specialist faces more uncertainty, although his experience within that domain allows him to quickly recognize the type of problem and how to study it. For the case-oriented practitioner, uncertainty is greater still. Not being familiar with the domain, and possibly having to use an unfamiliar approach, he can hardly be expected to come up with a precise proposal for the entire project on day one. A multi-stage approach, in which the project is gradually defined more precisely, is one possibility, but in any case both client and practitioner will initially have to live with a relatively large extent of indeterminism. The client's expectations are often not clear in advance, neither to the practitioner nor even to herself, so the actual outcomes are more likely to be a surprise (positively or negatively) to her.

Another source of uncertainty to which the case-oriented practitioner is especially vulnerable is changes within the client. Particularly in relatively long projects with large organizations, the person initiating it from the client's side will often move to a different post during the project. If her successor is not as favourably disposed towards OR, the project may well be discontinued, unless there is a clear project plan or unless it is obvious to her that the consultant is the right person for the job; ensuring the latter will be easier for a domain- or a tool-oriented practitioner than for a case-oriented one.

5.4 Interacting with the client organization

Communicating effectively with clients has always been cited as a sore point by the OR community. The consultant essentially needs to 'learn the client's language', which can differ

strongly from one client to the next. Because both domain- and tool-oriented practitioners frequently deal with similar clients, they have the opportunity to learn this language; for a case-oriented practitioner, this can require a lot of skill and flexibility on his part.

Another frequent stumbling block is data gathering. Here, tool-orientation puts the practitioner at an advantage: having repeatedly implemented highly similar systems, he knows exactly which data are needed and can specify these in advance, to prevent awkward surprises later on when it turns out that the required data are not systematically collected. He has also seen which people in an organization typically hold those data, and has experience in getting their cooperation.

Practitioner orientations presuppose varying degrees of familiarity with OR on the client's part. A client engaging a tool-oriented consultant need know nothing about OR, and need not even believe in OR; often, she will not be aware that the approach is in fact OR-based. Before engaging a domain-oriented or particularly a case-oriented practitioner, the client must be convinced that an OR approach is appropriate, whether based on familiarity with OR or on something else. Otherwise, why not select some self-professed specialist offering cut-price rates or a far more widely known general management consulting firm? It is clear that consultant D's approach to co-developing products in joint ventures with 'clients' places special requirements on the latter's competence and willingness to cooperate.

This list of orientations is clearly not complete. Eg, the various systems-based approaches, as formulated and practiced largely in the UK and discussed in among others Keys (1991), could be seen as one or several separate orientations. The key underlying competence is not so much knowledge of a particular domain or a problem-solving tool, but a problem-structuring methodology, which places different requirements on consultant and client. Systems analysis, as described in the *Handbooks* by Miser & Quade (1985, 1988) and Miser (1995a), is different again: working in or leading interdisciplinary teams for large-scale projects of strategic importance presupposes other specific skills, and constitutes yet another orientation.

It is important to re-emphasize that these project approaches are an abstraction of reality, and that they do not correspond to the strands of practice outlined in Section 4. But, every strand does lean more heavily towards some project approaches than towards others. Also, the implications outlined here are speculative in nature, given the exploratory nature and

limited scale of the research underlying this paper.

6 Lone Rangers and the Practitioner's Dilemma

The existence of different project approaches has profound implications for the OR profession and the way it is practiced. Consider the case of the Lone Rangers or Independent Consultants, OR practitioners operating as individuals within an organization, and whose plight has recently attracted much attention. Paul Gray (1993) for instance investigated the needs of Lone Rangers while he was TIMS president. Ian Mitchell (1993b) reports on the Independent Consultant Network set up by the British OR Society.

Each of the previous section's project approaches has strong and weak points. The pure form of case-by-case orientation, though, seems difficult to maintain. When clients expect the consultant to have the experience of a domain-oriented practitioner and the powerful toolkit of a tool-oriented practitioner, a case-oriented practitioner faces a formidable task. By their very nature, Lone Rangers frequently face this daunting challenge. In other words, Lone Rangers are an exemplification of the Practitioner's Dilemma. Writing from his own experience as a Lone Ranger, Mitchell (1993a) believes that the "OR practitioner seems to be something of a Troubleshooter". As opposed to following a 9 to 5 mentality, he finds that "A focus on project-orientation has become the only way to structure my time". From this perspective, it is no surprise that many Lone Rangers cannot live up to their clients' expectations, and suffer a high rate of burn-out or decide to leave the OR profession.

By facing up to the existence of strands of practice and the near impossibility of managing without one, we believe that a Lone Ranger can form more realistic expectations about his practice. He will generally have to develop his own strand, though not necessarily one of those we discussed. Gray (1993) observed that "they see themselves as either experts in problem-solving generally or, more typically, as experts in solving problems in specific industry or functional areas". Ken Bowen's work (see eg Bowen 1992) could be seen as built on problem-formulation skills. A Lone Ranger may develop a strand based on his unique and intimate knowledge of his organization, with all its idiosyncracies in terms of project management approaches, communication styles, etc.

7 Conclusions and further research

To conclude, let us reiterate the two key points made in this paper.

- First, to properly understand an OR practitioner's work, it is not sufficient to focus on individual projects; on the contrary, one needs to study his entire 'strand of practice', the specific set of competences he has developed over a series of projects and how he puts them to use in future projects.
- Second, OR practitioners each develop their own strand of practice, as a (possibly strategic) response to the particular circumstances they face. Normative statements about OR practice will generally not hold for all strands, but only for some; to be useful, any such normative statement must include a delimitation of its domain of validity.

Much of this paper revolves around the tension between the espoused theory (Argyris & Schön 1974) as often found in OR literature and the theory in use we observed. On the basis of the latter we construct a view of OR practice that is perhaps less heroic but more feasible than that implied by the prevailing view which, when extrapolated, leads to the Practitioner's Dilemma. This gap between espoused theory and theory in use clearly has implications for OR education; to the extent that the theory of practice does not correspond to observed practice, practitioners involved in education are implicitly induced to apply the maxim "don't do as I say, do as I do". By helping practitioners describe their own strands better, they can achieve a clearer understanding of their practice, and so respond better to the contradictions inherent in it. This should also make it easier to capture and transfer their know-how to novices in the field.

This paper has argued that writing about (successful) projects, a tradition in the OR literature, serves many useful purposes but is certainly not sufficient to help novices understand OR practice. It may actually be counterproductive in that it may create a naive and unrealistic view of the profession. We should follow successful practitioners during a series of successive projects and write about strands of practice, how to develop them and to nurture them.

This research was conceived as an exploratory project, and this paper should be viewed with that in mind. We believe we have clearly demonstrated that different strands exist; we do

certainly not claim to have described all possible strands. Moreover, our explanation of how and why these strands have evolved may seem plausible and even compelling, but is as yet no more than hypothetical. And clearly, the implications in Sections 5 and 6 for performing OR projects are speculative in nature. However, we hope to have demonstrated that these issues are of sufficient importance for OR practice to merit further, deeper research. We are currently working on a deeper analysis of our data, to further clarify the strands we found and their implications for practice.

A key question is how OR practitioners transfer knowledge from one project to another, over time, and across individuals. The entrepreneurial strand of practice, mixing performing projects for clients with a step by step co-development of a system with those clients, is an intriguing and surprising one; its ramifications need to be understood more clearly. Also, further research should uncover other strands of practice, and how new and viable strands are developed. Finally, much of what we discuss here is closely related to various findings in business strategy. A useful exercise would therefore be to explore in more detail what the links are between strands of practice in OR and the literature on content and process of strategy formulation.

Acknowledgements

We are very grateful to Ken Bowen, Michael Lapré, Hugh Miser and two anonymous referees for their helpful suggestions. We also wish to express our warm gratitude to the practitioners and clients involved, for their openness and enthusiastic cooperation. This project was funded by INSEAD R&D budgets 2010.2902 and 2118R.

References

Advanced Research Institute on 'Education in Systems Science' (1979), held at Gras-Ellenbach, Germany, February 21-25, 1978, Report and Proceedings; Taylor & Francis Ltd, London.

- Argyris, C. & D.A. Schön** (1974), *Theory in Practice*, Jossey-Bass, San Francisco.
- Batson, R.G.** (1987), "The modern role of MS/OR professionals in interdisciplinary teams", *Interfaces*, 17(3), pp. 85-93.
- Bowen, K.** (1990), "An Eighth Face of Research", *OMEGA*, 18(2), pp. 215-216.
- Bowen, K.** (1992), "Understanding and Structuring New Problems", presented at *A Conference on New Management Technologies and Transition to a Market Economy*, on the "Sally Albatross", Helsinki-St. Petersburg, July 1-3.
- Corbett, C.J., W. Overmeer & L.N. Van Wassenhove** (1995), "Learning the Ropes of OR", manuscript in preparation, forthcoming in Fortuin, van Beek & Van Wassenhove, *OR at Work*, Taylor & Francis.
- Corbett, C.J. & L.N. Van Wassenhove** (1993), "The Natural Drift: What Happened to Operations Research?", *Operations Research*, 41(4), pp. 625-639.
- Eilon, S.** (1974), "Seven faces of research", *Omega*, 2(1), pp. 1-9.
- Fortuin, L. & A.T.M. Korsten** (1988), "The OR worker in practice", *European Journal of Operational Research*, 37, pp. 325-327.
- Fortuin, L. & M. Zijlstra**, (1989), "Operational Research in practice: Experiences of an OR group in industry", *European Journal of Operational Research*, 41, pp. 108-121.
- Gray, P.** (1993), "Voice of the Lone Ranger", *OR/MS Today*, February, p. 6.
- Hildebrandt, S.** (1980), "Implementation—the bottleneck of operations research: the state of the art", *European Journal of Operational Research*, 6, pp. 4-12.
- Keys, P.** (1989) "MS/OR projects and their interaction", *OMEGA*, 17(2), pp. 113-122.
- Keys, P.** (1991), *Operational Research and Systems (The Systemic Nature of Operational Research)*, Plenum Press, New York.
- Meyer zu Selhausen, H.** (1985), "Organizational conditions of OR processes in German business firms—Results of an empirical investigation", *European Journal of Operational Research*, 22(2), pp. 191-202.
- Mintzberg, H.** (1989), *Mintzberg on Management: Inside Our Strange World of Organizations*, The Free Press, New York.
- Miser, H.J. & E.S. Quade** (eds.) (1985), *Handbook of Systems Analysis (Overview of Uses, Procedures, Applications, and Practice)*, John Wiley & Sons, Chichester, England.
- Miser, H.J. & E.S. Quade** (eds.) (1988), *Handbook of Systems Analysis (Craft Issues and Procedural Choices)*, John Wiley & Sons, Chichester, England.

- Miser, H.J.** (ed.) (1995a), *Handbook of Systems Analysis (Cases)*, John Wiley and Sons, Ltd, Chichester, England (forthcoming).
- Miser, H.J.** (1995b), personal communication.
- Mitchell, G.** (1993), *The Practice of Operational Research*, John Wiley & Sons, Chichester, England.
- Mitchell, I.** (1993a), "Of Lone Rangers and Terminators: Some thoughts on the solo practitioner", *OR Newsletter*, March, pp. 19-20.
- Mitchell, I.** (1993b), "New network set up for solo practitioners", *OR Newsletter*, October, p. 5.
- Pidd, M.** (1995), "Pictures from an exhibition: Images of OR/MS", *European Journal of Operational Research*, **81**, pp. 479-488.
- Rosenhead, J.** (ed.) (1989), *Rational Analysis for a Problematic World (Problem Structuring Methods for Complexity, Uncertainty and Conflict)*, John Wiley & Sons, Chichester, England.
- Royston, G.** (1993), "The wonder of OR? Diagnosing the Woolworth syndrome using the Dillons dipstick", *OR Newsletter*, April, pp. 16-19.
- Schön, D.** (1983), *The Reflective Practitioner (How Professionals Think in Action)*, Basic Books, Inc.
- Weingartner, H.M.** (1987), "The Changing Character of Management Science", *OMEGA*, **15**, pp. 257-262.
- Zeleny, M.** (1979), "The Last Mohicans of OR: Or, It Might Be in the "Genes" ", *Interfaces*, **9(5)**, pp. 135-141.