Vol. 13, No. 1, September 2012, pp. 17–27
ISSN 1532-0545 (online)

INFORMS Transactions on Education

http://dx.doi.org/10.1287/ited.1120.0093

A Seminar for Solving Client Problems in Project Teams

Ahti Salo
Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, 00076 Aalto, Finland, ahti.salo@aalto.fi

We describe the course "Seminar on Case Studies in Operations Research" in which teams of four or five students work on real problems posed by firms and governmental research institutions. In this course, which is taught at Aalto University School of Science, students use operations research to solve these problems and also develop teamwork, communication, and project management skills. Since 2001, the student teams have carried out more than 60 projects of which many have had significant business and societal impacts. We discuss experiences from this course and consider implementation details that have contributed to the attainment of learning objectives.

Key words: collaborative learning; group projects; teaching with projects; team teaching
History: Received: June 2011; accepted: July 2012.

1. Introduction

In typical operations research (OR) curricula, the majority of courses have a methodological focus on specific OR methods such as linear programming (Bazaraa et al. 1993), simulation (Law and Kelton 2002), and decision analysis (Clemen 2009). In these courses, the students solve textbook examples and submit individually completed classroom exercises. Yet these courses do not necessarily encourage students to think about which OR methods should be applied because this choice is implied by the course title. Moreover, textbook examples tend to be less complex and ambiguous than real problems in which several experts may have to collaborate in order to meet client expectations. A further difference is that the problem solving process must be systematically structured and managed to ensure that the relevant constraints of time, budget, and quality are not violated.

The above observations, together with the recognition that OR is an applied discipline for solving decision problems through the deployment of advanced analytic methods, suggest that there is a need for courses that help students develop skills for applying OR to real problems (Behara and Davis 2010, Cherney 2008, Revans 1998). Equipped with such skills, the students should be able to interact effectively with the client, to structure complex problems, to organize their work into a coherent plan, and to work together as a team. These skills—which are much desired of recent OR graduates—can be developed through work on real-world problems.

This paper describes the course "Seminar on Case Studies in Operations Research," which seeks to foster the above skills. In this course—which has been running since 2001 at what is presently Aalto University School of Science—students work in project teams of four or five students and tackle problems presented by firms and research institutions. Each year, about 20 students complete a total of four or five projects; these are the case studies that the course title refers to. Most course participants are M.Sc. students, but there are often doctoral students, too. In addition to using OR methods to solve client problems, the students develop project plans and interim reports, present their results in a final report, hone their presentation skills, and learn about the work of the other project teams. At the end of the course, the students provide feedback on their learning experiences and the overall organization of the course. This feedback has been overwhelmingly positive, and many projects have had significant impacts on business and society. The course design is generic enough to be adapted for use at other universities.

The rest of this paper is structured as follows. Section 2 describes the context of the course and its learning objectives. Section 3 describes the course design, §4 discusses implementation details, and §5 illustrates a selected project. Section 6 concludes.

2. Context and Learning Objectives

The author developed the course on "Seminar on Case Studies in Operations Research" during the academic year 2001–2002. Most students still come either from the MSc program in technical physics and mathematics or from the MSc program in industrial engineering and management. The two programs have some of the toughest entrance requirements in Finland. Almost invariably, the students will have studied systems and operations research as their major or minor subject. Thus, they will have taken courses on topics such as linear programming, decision analysis, simulation, optimization, and time-series forecasting. About 90% of the course participants are master's students who have completed at least three years of university studies and earned a total of more than 200 credits (including earlier B.Sc. studies) in the European Credit Transfer and Accumulation Systems[1] (ECTS) and about 30–60 credits in the specialization on systems and operations research. About 10% of participants are doctoral students. Overall, the course is taught to talented students with diverse backgrounds and a good knowledge of OR methods. Many students have already started to apply their professional skills through various forms of part-time employment.

The main learning objectives are the following:

Problem formulation: At Aalto University School of Science, the other courses on systems and operations research focus on specific OR methods and tools. The seminar course complements these by seeking to ensure that the students better understand how OR can be brought to bear on real problems. The students are therefore invited to work on real-world problems and to take responsibility for problem formulation as well (Behara and Davis 2010, p. 25).

Project-based teamwork: Because the students are accustomed to submitting individual assignments, most have not had much experience in working as a team. This motivates the second objective, i.e., learning about project-oriented teamwork for a real client. The clients are external in that they are not from the university. They are identified before the start of the course as described at the beginning of §4. Each client provides a project topic and is genuinely interested in the deliverables produced by a student team with a designated project manager.

Communication and presentation skills: The third objective is to improve communication and presentation skills, ranging from the exploration and clarification of client requirements to the presentation of project plans, interim reports, and final reports. These skills are particularly important during problem formulation and the delivery of final results.

Domain knowledge and self-confidence: By working on their projects, the students acquire knowledge about a new problem domain. This process of systematic knowledge acquisition enhances their self-confidence and improves their problem solving abilities so that they are better equipped to address problems in other contexts.

[1] The bachelor's degree consists of three years of studies and 180 ECTSs, which is followed by the master's degree that takes two years and consists of 120 ECTS; see http://en.wikipedia.org/wiki/European_Credit_Transfer_and_Accumulation_System

3. Organization of the Course

3.1. Roles and Responsibilities

The roles and responsibilities of the stakeholders are as follows:

Clients are either firms or governmental research institutions faced with a real problem that is potentially amenable to the deployment of OR methods. The clients present the problem to the students, provide access to data, interact with the students throughout the project (usually three to five meetings of two hours each), and evaluate the results.

The two teachers are the *professor* and the *teaching assistant*. The professor runs the overall organization of the course. This includes the identification of clients and the initial selection and scoping of projects (i.e., assessing whether a problem, as proposed by the client before the course, is indeed suitable for the course); the assignment of students to project teams; the supervision of project teams, with an emphasis on scoping and choice of methodologies; and the delivery of feedback on the results. The teaching assistant offers practical assistance, for instance, by informing the students about client excursions and access to software tools.

The *students* work in *teams* so that each team addresses a project topic proposed by a client. Each team has plenty of freedom in problem scoping, and the instructors encourage the students to approach their research topic with an open mind. In particular, the students are not just "told" how to solve the problem. Rather, the project team is asked to scope its problem so that it expects to provide useful results to the client on time (in three months) and within the "budget," defined essentially by the credits the students expect to earn (five credits per student).

Each project team has a *project manager*, who has additional responsibilities for communicating with the client and the teachers, providing leadership to the team, taking the lead in the specification and assignment of project tasks, and ensuring that all team members contribute. The project manager is advised to talk to the professor, confidentially if need be, if the team encounters major difficulties of any kind.

Each project manager earns two extra credits for his or her work.

Each project team has an appointed *shadow team*, which is one of the other project teams. The students in the shadow team follow the activities of this other team throughout the course. For instance, the shadow team provides oral and written feedback on all the deliverables (project plan, interim report, final report) of the team that it shadows.

3.2. Initial Kick-Off Meeting

Held at the end of January, the initial three-hour kick-off meeting begins with an initial 30-minute presentation in which the professor explains the learning objectives and the practical course arrangements. Specifically, he gives a short motivational introduction to the course and outlines the roles and responsibilities of teachers, clients, and project teams. He explains the grading principles (pass/fail), describes the overall schedule (Table 1), and urges the students to study his teaching materials on project planning and management that cover good practices in teamwork as well. The professor then describes the procedure for building the teams and clarifies policies for the choice of reporting language (Finnish, English) and the signing of nondisclosure agreements (NDA). He stresses that problem scoping is essential and that each team is partly responsible for its scoping decisions.

Following this introductory presentation, each client gives a 20-minute presentation, which first covers facts such as the number of employees, organization, and main products and services. The client then describes the proposed project topic with an emphasis on the objectives and the relevance of OR methods. After each client presentation, the professor gives the students a handout on the proposed topic (see Appendix A for an example). This handout describes the topic, outlines the objectives of the project, and provides methodological suggestions. The handout has been developed by the client and approved by the professor in advance. This ensures that the client has devoted enough attention to the problem, and the handout helps the students build an informed opinion about the topic.

The client presentations are followed by a 10–15 minute pause during which the students indicate which topics they wish to work on. Specifically, the students are requested to rank their four most preferred project topics in order of preference. There can be pairs of students who wish to work together; this is permitted by allowing such pairs to submit a joint ranking of their preferred topics. Larger teams cannot submit joint rankings, however, because it could then be difficult to achieve a balanced assignment of students to project topics.

Next, the teachers analyze the students' rankings and form the project teams, assisted by an optimization model, in order to ensure that (i) all projects teams have four or five students, (ii) as many students as possible can work on their preferred topics,

Table 1 The Overall Schedule of the Course "Seminar on Case Studies in Operations Research"

Activity/role	Clients	Teachers	Project teams
November–January	Identification of prospective clients Screening of project topics	Identification of prospective clients Screening of project topics	Registration for the course
Kickoff meeting, late January	Presentation of the clients' activities and project topics	Presentation of course arrangements Assignment of students to project teams	Exchange of contact information Appointment of the project manager
February	In-depth discussions with the project teams Delivery of data	Meetings with teams to check the viability of project plans Suggestions for literature, methods, tools	Problem formulation Project planning Literature review Development of project plans
First excursion, late February or early March	Presentation of the first host client and its OR activities Feedback to project teams	Feedback to project teams	Presentation of project plans Feedback from the shadow teams
March	Guidance to project teams	Ad hoc meetings with projects	Work on projects Preparation of interim reports
Second excursion, late March or early April	Presentation of the second host client and its OR activities Feedback to teams	Feedback to teams	Presentation of interim reports Feedback from the shadow teams
April	Guidance to project teams	Ad hoc meetings with projects	Work on projects Writing of final reports
Third excursion, early May	Presentation of the third host client and its OR activities Feedback to project teams	Feedback to teams	Presentation of final reports Feedback from the shadow teams
May	Approval and clearance of projects	Final approval of deliverables Grading	Implementation of corrections to final reports Feedback on the course

and (iii) students are not forced to work on projects that they did not mention in their ranking. The first of these three constraints is strict. The second and third constraints help ensure that students need not work on projects they are not interested in. Although the optimization model provides decision support for team formation, allowances can be made based on other considerations (e.g., ensuring the breadth of skills in the team). The shadow teams are appointed soon after the project teams have been formed, based on the principle that students who could not work on their most preferred project topics should nonetheless be able to comment on them.

The assignment of students to teams usually takes fewer than 30 minutes and the project teams can therefore be announced at the end of the kickoff meeting. The team members then meet and each team chooses its project manager. At the kickoff meeting, each project team is given the contact information of its client (email, phone number) so that the project manager can contact the client in order to schedule the first client-team meeting. Moreover, every project manager is asked to confirm that all team members are fully committed: this helps exclude the possibility that a student would drop the course for not having been assigned to his or her most preferred topic. In order to foster good group dynamics, new students are not permitted to join the teams once the teams have been formed.

3.3. Project Planning

About a week after the kickoff meeting, each project team meets its client, with whom it discusses the project topic in more detail. The project team then carries out a review of relevant literature and develops a tentative project plan, which they present to the professor in a 30- to 45-minute meeting. In this meeting, the students describe how they perceive the topic and how they intend to carry out the project. This meeting helps ensure that the team has started working and is making progress. The professor often presents suggestions and may, for instance, ask the students to clarify and redefine their objectives. After the meeting, the team writes its 5–6 page project plan, structured under standardized headings (e.g., background, objectives, tasks, schedule, task assignment, risk management plan).

The meeting at which the project plans are presented—as well as the later meetings for the presentation of interim reports and final reports—takes place during excursions that are hosted by three different clients. During a typical excursion, the teams give their presentations first and the hosting client then gives an overview of its activities, often with a focus on its use of OR. This allows the students to learn about the client's activities.

The meeting for presenting project plans takes place in late February or early March. It is attended by the project teams, the teachers, and the client, who hosts the excursion. Two days before this meeting, the project managers send their project plans to the teachers and their shadow teams for comments. During the meeting, each team has about 20 minutes for presenting its project plan, which is followed by 10 minutes of discussion.

3.4. Interim Reports

In March, each project team works according to its project plans and, if need be, has meetings with its client and the teachers. Many of the client meetings are focused on the clarification of objectives and the scoping of the project topic. In more technical matters—such as which OR methods are likely to be most appropriate—the project team usually turns to the teachers for advice.

In late March or early April, all project teams present their interim reports during an excursion that is again hosted by one of the clients. Each interim report (i) summarizes the progress that the project team has made, (ii) describes possible changes to the initial project plan, and (iii) contains a revised risk management plan. The interim reports are typically two to three pages.

The presentations are given using the same format as the project plans. That is, each team gives a 20-minute presentation, followed by 10 minutes of discussion with the shadow team, the teachers, and the client (if present). In its presentation, the project team describes what it has accomplished, how the project plan may have been revised, and what the prospects for completing the project successfully are.

3.5. Final Presentations

In April, each project team devotes a significant share of its time to the writing of its final report. This report is usually about 30 pages and, following the policy at Aalto University, it will be made publicly available on the Internet. As a result, some reports contain illustrative results based on modified data, rather than confidential data that the team may use when delivering results to its clients.

In addition to the final report, each project team is required to write a two-page summary on "lessons learned," discussing what the team has learned and how well it thinks it has performed overall. The team is encouraged to be explicit about what it would have done if it were given the chance to start again. These summaries give the teachers insights into what difficulties the teams may have encountered and how they have sought to address these.

The final reports are presented in late April or early May during an excursion that is hosted by a client.

In its presentation, each team focuses on its achievements and the significance of the results. Feedback on this presentation is given by the shadow team, the teachers, and the hosting client, but the other project teams are also invited to comment. Based on the feedback, the team may be required to revise its final report.

At the end of this excursion, the professor explains the final practicalities such as getting a clearance from the clients, submitting the revised final reports, and completing the questionnaire for student feedback. He reminds the students that many of them will soon work in organizations in which they may encounter problems that are potentially suitable for the seminar course. In such a situation, they should consider possibilities for offering a project topic for the course. This mode of soliciting project topics has become increasingly important. It has the advantage that former students have had firsthand experiences with the course.

The students have given exceptionally good feedback ratings for the course. This feedback is collected systematically at the end of the course with a survey that has 18 questions on the relevance of project topics, the level of support provided by the client and the teachers, and the amount of work that the course required, among others. In 2011, for instance, all students reported that the course helped them develop their teamwork skills and that the course was useful. Also, the presentations by the other project teams were judged either very or somewhat interesting. Looking at all the feedback that has accumulated since 2001, it appears that this kind of feedback has been particularly positive in teams in which the student-client interaction has been intensive and the students have been very interested in the problem topic. This suggests that it is important to offer topics that are perceived as real problems by the clients and to assign them to those students that are highly motivated by these topics.

4. Implementation Details

The above course design is supported by the following implementation details:

Solicitation of project topics: When the course was first established, the professor solicited all topics proactively through his contacts. Presently, more than 60% of topics are proposed by prospective clients in three different ways. First, companies and research institutions may approach the professor with OR-related problems, which are not intended for the course, but which can be shaped into project topics. Second, some clients—such as the Technical Research Centre of the Finnish Defence Forces—have offered topics almost every year. These clients—whose share is about 30% of all clients—understand what they can expect from the project teams so that the professor needs less time to screen the topics proposed by these repeat clients. Third, some topics are proposed proactively by course alumni who work in organizations that have problems that are potentially amenable to OR methods.

Not all project topics pass the screening phase. To qualify, the topic must fulfil three criteria: specifically, it must be (i) *real*, meaning that the problem is relevant to the client and that there is a committed contact person who is prepared to collaborate with the team; (ii) *feasible*, meaning that the professor believes it is likely that a student team can tackle the problem constructively in four months without investing more than five credits of work per student; and (iii) *instructive*, meaning that the topic is not too straightforward and offers opportunities for learning about OR. For example, topics have been turned down because they would have necessitated laborious software implementation efforts or because they would have required too much expertise about an unfamiliar problem domain. During the negotiations, the clients are also informed about the course practicalities (e.g., liability issues, intellectual property rights, nondisclosure agreements) that they must accept if the project topic is to start.

A client may wish to propose a project topic in the expectation that it can soon give data for the students to work on. However, because there may be unexpected delays in the acquisition and delivery of such data, a project topic should be approved for the course only if the requisite data are surely available. A project team will be very discouraged if it cannot proceed because of problems of data availability that are beyond its control.

Forming student teams: All students want to work on the project they are most interested in. But because some topics tend to be more popular than others, the building of teams calls for a systematic procedure that first allows the students to express their preferences for the proposed topics and that then maximizes the match between student interests and topics, subject to the constraint that each team will have four or five students. The professor has the final say so that other factors such as the match between topics and the students' educational background can be accounted for. The students are not given information about the preferences expressed by the other students: instead, they are given only arguments about why the final project assignment should be accepted (e.g., 80% of students will work on their first or second most preferred topic; no one has to work on his or her least preferred topic). The reason for *not* conveying information about preferences is that this minimizes the possibility of complaints (i.e., "why did I not get to work on my most preferred topic, but my friend did?") at a time when it is critical to build cohesive teams and foster project momentum.

Team size: The "optimal" team size seems four or five students. In smaller teams of three, the amount of work required of each team member may become excessive. Also, if one of the three students were to drop the course for whatever reason, the project would have to be finished by two students and the true team dynamics would be lost. Conversely, in teams of more than five students, it would be more difficult to ensure that all team members bear a comparable workload.

Course schedule: Because the course runs from late January until early May, negotiations with prospective clients can be carried out from November through January. It would be harder to start the course in September, because negotiations with clients would have to be carried out during the summer (which would be difficult because of holidays) or in May and June (in which case the client interest could wane because of the long delay between negotiations and the start of the project).

The number of students who enroll for the course in mid-January depends on what other electives are offered during the academic year; but as discussed above, the project topics must be negotiated one or two months before the start of the course. The students are therefore encouraged to submit their tentative registration in December. Those who register early are more likely to be assigned to their most preferred topics, and those who come to the initial meeting without advance registration may be assigned to topics that they prefer less. Moreover, the clients are given an advance warning that the project for their topic will be initiated only if there are enough interested students. Every client therefore has an incentive to give a good presentation at the kickoff meeting, because otherwise there is a risk that there will not be enough enthusiastic students to work on its project topic, although all topics generally receive some interest. However, if the number of participating students is less than expected, topics can be discarded or postponed (the share of such topics is less than 5%). In such an event, the professor discusses with the client possibilities for other forms of collaboration, for instance by starting a master's thesis project.

Overall, it is useful to impose a strict schedule so that the students know when they must present their deliverables. This creates a sense of urgency and puts social pressure on the students. No team wishes to report that "we have accomplished hardly anything" when the other teams are making good progress.

Client interaction: The interaction between the project team and its client is guided by recommendations about the expected number of meetings between them. Specifically, the client and its respective team are advised to have three to five meetings that last up to two hours each. This helps ensure that there is sufficient and continuous interaction between the team and the client. The teachers also solicit feedback from the client on the project. This is an important mechanism of quality control that also helps verify that the project deliverables do not contain confidential information.

The deliverables are posted on the Internet and can therefore be accessed by other stakeholders. For instance, the evaluation of helicopters for the Finnish Border Guard—which was one of the projects in 2009—caught the attention of health care administrators who then approached the professor with the proposal that a comparable analysis of helicopter emergency medical services should be conducted. This demonstrates that the dissemination of results in the Internet may generate good project topics.

In the rare case of a major failure on the client's side, an entire project team can be given a new project topic in the middle of the course. There has been but one such case when the client contact had to take a long leave because of unexpected health problems. In this case, the team could not complete its project on the same schedule as the others; but its students were nevertheless able to complete the course and to earn credits, albeit after some delay.

Presentations skills and interaction with shadow teams: The three excursions during which the teams present their project plans, interim reports, and final reports are useful because each team must respond to questions posed by the shadow team, the teachers, and the hosting client. These discussions help the teams become clearer about what their problem really is and why their approach is suitable. Working in the shadow teams trains students to assimilate results produced by others because each shadow team is explicitly required to provide feedback on all the deliverables (project plans, interim reports, final reports) produced by the team that it shadows. Among other things, this contributes to an effective interaction among the teams, whose performance may be enhanced by the suggestions of their respective shadow teams. Moreover, the shadow teams reduce the teachers' workload by making valid observations so that the teachers can focus on points that have not been made yet.

Compensation policy and liability issues: Although many projects produce results of commercial value, the projects are carried out *pro bono* without payments to the students or to the University. This policy follows from the need to treat all students on equal terms (i.e., no student can complain that some other student was assigned to a project with a higher compensation). Nor are there contractual agreements between Aalto University, the clients, or the students. This is partly because the development of such contracts would complicate the preparation of project topics and because it would call for excessive administrative support. Contractually, Aalto University is not liable for the results, and the clients can use the results "as is," with no formal quality guarantees. The intellectual property rights (IPR) to the software tools and other deliverables belong to the students. Thus, if the clients want access to the IPR, they must negotiate an agreement with the students.

Nondisclosure agreements: If a client requires that the students sign nondisclosure agreements, or if it expects to use the software tools produced by the project team, the client must state such requirements explicitly when presenting its project topic at the kick-off meeting. The students are consequently aware of such requests and, in particular, any student who does not accept them can exclude these topics from his or her list of preferred topics. In particular, the teachers cannot require that a student should sign an NDA against his or her will. Instead, the professor is responsible for ensuring that a student can, if he or she so desires, work on a project topic that does not require an NDA.

Grading policy: The course is graded as pass/fail, in the understanding that the students will pass provided that they work hard enough. This is for the following reasons: (i) the project topics are varied and comparisons among them are therefore difficult, and (ii) the teachers cannot ascertain how significantly the team members have contributed. If the project deliverables (including deliverables for shadowing, too) are not acceptable at the time of final excursion, the team may be required to continue working until the professor approves the deliverables. If the project manager reports that there is a nonperforming student on the team, then, based on face-to-face discussions between this student, the other team members, and the teachers, the professor may require that such a student produces an additional independent report in relation to the project topic. This student will receive credits only after his or her additional report has been submitted and approved.

Continued activities: If the client is interested in continuing the work that the project team has started, it usually seeks to hire students from the team and, in numerous cases, a student has subsequently been employed by the client. The client may also acquire IPRs to the project deliverables from the students. Some projects have lead to the establishment of research contracts between the client and the university. For example, the course project *Multicriteria decision analysis for the optimal selection of road pavement projects* in 2004 was followed by several research projects that, for instance, were recognized in the INFORMS Decision Analysis Society's Practice Award Competition in 2007 and contributed to the development of journal publications in *Decision Analysis* and *European Journal of Operational Research* (Liesiö et al. 2007, Mild and Salo 2009).

Publication activities: After the course, some project teams have written a refereed paper based on the final report. Usually, this is possible only if the team has a doctoral student who needs publications for his or her dissertation. For instance, Renjish Kumar et al. (2005) is a streamlined version of the final report on the evaluation of the end-user usability of office applications. Although publications are desirable, possibilities for publishing more extensively are limited because (i) there are only a few doctoral who take the course, (ii) the MSc students do not have experience in writing papers for scientific journals, and (iii) the teachers cannot offer much support for writing these publications. This notwithstanding, there are close links with the projects and the research that is carried out at the department. For example, the 2006 project *Prioritization of new product features in software development* motivated the numerical example that was published in Liesiö et al. (2008). Furthermore, the results of methodological research at the department have been adopted by the project teams. For instance, methods of Ratio-based Efficiency Analysis (Salo and Punkka 2011) were applied in the project *Analyzing the efficiency of Finnish health care units* in 2011. This project has paved way for related M.Sc. theses on the efficiency analysis of health care services.

Costs: The costs to the university are small because the students can complete their projects by using available infrastructure (e.g., access to computers and software tools). The university may buy relevant books and other scientific references if the professor thinks these would be needed for the successful completion of the project. If the students wish to use advanced software tools that are not available on the university's computers, the client has at times provided the students with such tools. However, as a rule, Aalto University does not buy specialized software for an individual project because the project topics can usually be tackled satisfactorily with standard software. When excursions are made to clients that are far away from Aalto University, the travel costs are borne either by the university or the clients. There are no charges for participating in the course. This is in keeping Aalto University's broader policy of no tuition fees.

Assessing impacts: Most projects, or close to 90%, have been successful in the sense that the clients have been satisfied with the results, which they have found useful, too. This feedback has been solicited by the professor through direct discussions with the clients who have commented on the deliverables and the performance of their teams. The benefit of these discussions is that they are interactive, which has made it possible to assess the students' performance and the project impacts in more depth than what would possible with a standardized questionnaire. Also, client feedback is needed quickly because the students should not be kept waiting for their credits. A drawback of soliciting feedback through a questionnaire is that some clients might reply after an unacceptable delay.

There is considerable variability in how the clients have used results. For instance, in the projects related to *Benchmarking and best practices in sales forecasting*, Nokia asked the students to benchmark its forecasting models with other kinds of models. This project indicated that the models in use were relatively good, even if improvements in accuracy could be achieved by developing tailored models for different market segments and by deriving more accurate estimates. This project did not create immediate business impacts because the changes to the models were relatively minor. Notwithstanding, the project was still valuable because Nokia now knew how good its models were, so it did not need to start potentially expensive consultancy projects on the topic, for example.

One concern in assessing impacts is that these may be difficult to quantify in monetary terms during or immediately after the course. For instance, the project *Impact of medical service on the combat performance of infantry*—which was proposed by the Technical Research Centre of the Finnish Defence Forces—developed a simulation model for predicting how the combat performance of infantry troops depends on the siting of medical stations (i.e., proximity to frontier, number of medics in platoon, mode of transport). After this project, the Defence Forces organized field exercises through which the numerical parameters were estimated by observing the performance of recruits in different terrains. Using these parameters in the model then provided support for the planning of medical services, yet it would be difficult to ascribe monetary value to these more informed siting decisions.

From an administrative perspective, requiring all projects to produce quantifiable business impacts could complicate the solicitation of project topics and limit the scope of topics that can be addressed. Such a stringent requirement would favor pressing short-term problems with immediate and demonstrable business impacts. Clients faced with such problems would likely insist on contractual agreements that are not part of the course design and that Aalto University does not wish to establish because it does not want to be liable for the students' work.

From the perspective of possibilities for publications and long-term impacts, exploratory topics that are either new or that can be fruitfully revisited by applying new methodologies seem particularly promising, and projects on such topics may generate significant impacts over time. For instance, the majority of Finnish road districts have now for many years used multi-criteria decision analysis to optimize their maintenance programs for bridge repairs as a result of the developments that started with the case study project *Multi-criteria decision analysis for the optimal selection of road pavement projects* in 2004.

Learning outcomes: In their summaries "lessons learned," the students have often highlighted learning experiences arising from teamwork and collaboration with an external client. These summaries suggest that such skills can be strengthened by forming teams in which the students are given the opportunity to work with students with whom they have not worked together. The teams also value learning about systematic project planning and management, even if the project teams are not large. Importantly, many, if not most, projects have encountered unexpected challenges that have necessitated actions such as modifying the problem scope, pretreating the data, or using modeling tools other than initially planned. As a result of the course, the students are therefore likely to embrace a healthy "expecting the unexpected" attitude in the early years of their careers.

Even other observations on learning can be made. First, projects do not usually have a unique "right" solution because the topics can be addressed through different approaches; this, in itself, is in contrast to many textbook examples. Second, because the projects must be finished under a strict schedule, it is crucial to eliminate technical risks early on. For instance, it may be advisable to proceed quickly with the parallel exploration of alternative strategies and then to focus on the most promising one instead of evaluating different strategies in sequence. Third, unconventional topics—whose "value" to the client may hard to quantify—can offer opportunities for creativity. In the 2008 project *Sculpture variations*, for example, the students developed advanced optimization models and visualization tools for a professional sculptor who sought to design sculptures consisting of horizontal and vertical beams that had to be connected in specific ways. With these tools, the sculptor was able to explore sculpture designs virtually before the actual assembly and welding. Fourth, for effective communication, teams may need to be encouraged to focus succinctly on the problem and the validity and significance of results (see Grossman et al. 2008). Otherwise, the teams may overwhelm their audience with technical detail about all the work they have done. Finally, most teams tend to underestimate how much writing the final report requires. Thus, the development of the final report should be started earlier than what is foreseen in most draft project plans.

5. Examples of Completed Projects

As shown in Table 2, projects have addressed a broad range of topics, spanning a large diversity of clients, problem domains, and methodologies.

To illustrate a particular project, we describe the evaluation model that the students developed for the design of national helicopter emergency services. This was one of the six projects in 2009.

Finland is a sparsely populated country, with 5.3 million people and 41 inhabitants per square mile. Helicopters are therefore essential in the delivery of emergency medical services. Starting in the early 1990s, these services were offered by independent associations that were responsible for the emergency helicopters in their respective geographical areas. These local associations were not subjected to nationwide regulation. Prior to 2009, no nationwide studies had been conducted to determine the optimal sites of helicopter bases or the selection of helicopter types.

In 2010, the administrative situation changed when the Helicopter Emergency Medical Services (HEMS) were brought under the administrative control of the Ministry of Health and Social Services and the five Finnish university hospitals. In connection with this reform, both the sites of helicopter bases and the choice of different helicopter types were to be subjected to a thorough review under the direction of the HEMS Administrative Unit, which was the client of the project.

Table 2 Examples of Project Topics

Year	Selected project topic (out of 4–6 projects per year)	Client
2002	A calibrated model for the pricing of credit default swaps	Sampo Pankki Co
2003	Models for forecasting mobile phone replacement sales	Nokia
2004	Multi-criteria decision analysis for the optimal selection of road pavement projects	Inframan Co
2005	Assessing the impacts of first-aid medical delivery processes on combat performance	Finnish Defence Forces Technical Centre
2006	Prioritization of new product features in software development	Nokia Networks
2007	Optimizing the yield of a bioreactor	Medix Biomedica Co
2008	Spatial regression models for adjusting locally dependent treatments in field trials	Kemira GrowHow Co
2009	Optimal timing of harvesting of forest stands	UPM Co
2010	An evaluation model for designing national helicopter emergency services (HEMS)	HEMS Administrative Unit
2011	Analyzing the efficiency of Finnish health care units	National Institute for Health and Welfare
2012	Estimation of credit rating transition probabilities	OP-Pohjola Co

The project team was given two objectives in the redesign of HEMS operations: (i) to evaluate the suitability of alternative helicopter types for medical emergency services and (ii) to give recommendations for the optimal siting of helicopter bases. To this end, the project team developed a large-scale optimization model consisting of three parts: a search algorithm, a simulation model, and a performance evaluation model.

In the search algorithm produced by the project team, a major input was the number of helicopter bases, which varied from five to seven. The algorithm then formed decision alternatives in the search space consisting of 52 prospective sites for helicopter bases and nine helicopter types (these types differed with regard to transportation capacity and maximum speed, for example). A one-year simulation was executed to determine how an alternative, defined as a combination of (i) the type of helicopters, (ii) the number of helicopters, and (iii) the assignment of these helicopters to different sites, would perform with regard to two minimization criteria: (i) the mean time of reaching patients and (ii) the share of patient calls that would not be handled because of simultaneous patient cases or performance limitations. Because the search space was very large, Pareto-optimal solutions were generated using tailored genetic algorithms.

The model for simulating patient calls used past data and Monte Carlo techniques to generate a one-year sequence of patient cases calling for an emergency helicopter. These cases were generated on a geographical lattice in which Finland was divided into 20,000 hexagonal cells with different population characteristics. Expected patient case frequencies for these cells were inferred from statistical population data, and estimates about the range and speed of helicopter types were obtained from the manufacturers' performance charts. These estimates accounted for weather conditions (temperature, elevation, air pressure) and technical parameters (mass, fuel consumption, maximum cruise speed). All the above information was combined to determine if a given patient call could be handled, considering the location of the patient, the nearest available helicopter, the nearest hospital, and the range and speed of the helicopter. Emergency services would be delivered only to those patients who could be reached by an available helicopter, and the other calls would be left unhandled.

The project team sought to communicate its results so that they could be readily understood by aviation specialists, medical professionals, and decision makers. Figure 1, for example, shows (1) the earlier situation in 2010, (2) an alternative decision in which five helicopters are based at university hospitals and

Figure 1 The Pareto Frontier for Six Helicopter Bases at Different Sites

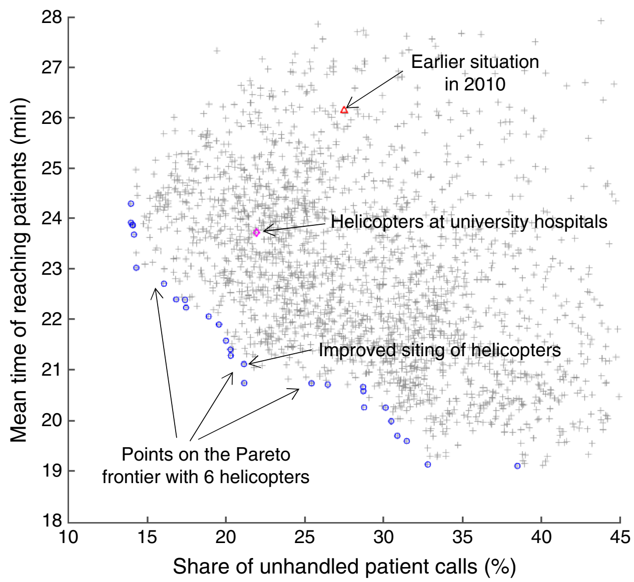

one in Lapland, and (3) numerous other alternatives defined by different combinations of 6 helicopters sites. The Pareto solutions, indicated by small blue circles, are those that cannot be improved simultaneously with regard to both criteria. An examination of this frontier suggested, for instance, that relocating helicopter bases can be expected to reduce the mean time of reaching patients from 26 minutes to 21 minutes, and the share of unhandled patient calls can be expected to drop from 27% to 21%. Moreover, comparisons among Pareto-optimal sets for different numbers of helicopters were made to determine how much the service level could be improved by having more helicopters.

The project results—which were referred to in the media and even in national television broadcasts—suggested that most helicopter bases should be located in densely populated areas where the number of patient calls is high and where contemporaneous calls were the main cause of unhandled patient calls. Practically all helicopter types were deemed suitable for HEMS operations in all weather conditions, with a few exceptions for the most northernmost part of the country. Maximum cruise speed was the most important technical attribute because fast helicopters reach patients more quickly. The cost of helicopters was less significant because the overall costs were driven mostly by the salary costs of medical experts.

After the completion of the project, the decision to locate the helicopter bases in the vicinity of university hospitals has been taken and the selection of helicopter operators has been made through a competitive bidding process. The HEMS Administrative Unit has acquired the rights to the software tools developed by the project team and has been using them in its planning activities.

6. Conclusion

This seminar course has been designed to help students improve their skills in applying OR methods to real problems by working in small teams on problems posed by clients. The course—which has received very good feedback from the students—complements more theoretical and methodological courses in which the students have learned the methods and tools that are needed for successful completion of real projects. Overall, this course design is likely to work best in the later stages of a master's program. At this stage, it is also easier to get client commitment because the clients can confidently expect useful results and because they may also be interested in providing employment offers to the students.

Finally, organizing this kind of a course can be an instructive experience for teachers. First, because the teachers can proactively identify interesting project topics, creating and coaching highly motivated project teams gives opportunities for learning about new topics and for creating impact by leveraging fresh talent on exciting problems. Second, the close interaction with students reveals what they have learned in other courses and, specifically, where the students' weaknesses and strengths lie. Third, because many clients are former students, the course helps illuminate what former students think of their university education. And, on a personal note, it is rewarding to see how students move on in their lives and careers, building on the skills they acquired during their formative years at the university.

Acknowledgments

The author is thankful for the high quality of refereeing provided by the editors of this special issue, the associate editor, and the two referees. He also thanks Alec Morton and Stefan Rice for their insightful comments and suggestions.

Appendix A. An Illustrative Description of a Project Topic

Modeling Long-Term Electricity Prices

Danske Markets

Background and Motivation. Electricity is a commodity with special characteristics. There is no convenient and economic way of storing electricity because consumption must match production at any time. Furthermore, the dynamics of electricity price differs from other commodities because of high seasonal effects and the high volatility of short-term prices.

Estimating and modeling prices of long-term electricity contracts is important for many players in the electricity market. Most importantly, producers are committed to long-term (up to tens of years) investments. Currently, there is an ongoing shift in production technology toward new technologies (such as natural gas and renewable sources). Market prices, however, exist typically only for relatively short future periods (up to 5–6 years). Although producers can hedge part of future production by selling either physical or financial contracts, there is a significant time period that cannot be hedged in the electricity markets. In addition, electricity markets still exhibit relatively low liquidity, which creates challenges for efficient hedging.

Although the correlation between short and long period prices in commodities such as oil is quite high, this is often not the case in electricity. In fact, spot and long-term dynamics of electricity prices can be very different because of different factors. This means that long-term commitments cannot be efficiently hedged using short-term contracts—which is the approach that is often used with other commodities. For modeling, this means that using spot price, or short-term price, in long-term price modeling may fail to capture the true dynamics and consequently yield misleading results.

Long-term electricity price modeling faces other challenges, such as availability of relevant market data. Furthermore, it may be difficult to detect and include a seasonal component in long-term pricing modeling. For example, although a seasonal component is present in spot dynamics, this component may not be very visible at times of high stress from demand or supply side. In addition, only a limited number of monthly and quarterly forwards are quoted in the market.

Research Problem. The project has several aims, all of which center on long-term electricity price dynamics:

- Build an understanding of short-term price versus long-term price dynamics of electricity.
- Present a literature survey of electricity price models and discuss approaches to the modeling of long-term price dynamics.
- Develop a viable model for long-term electricity prices. Attention is to given to issues such as data availability and user friendliness. Time series co-integrated vector autoregressive model is recommended as a possible approach.
- If possible, introduce a seasonal component to the model.
- Discuss whether the risk premium over spot price is stable or not.

References

Bazaraa, M. S., H. D. Sherali, C. M. Shetty. 1993. *Nonlinear Programming: Theory and Algorithms*, 3rd ed. John Wiley & Sons, New York.

Behera, R. S., M. M. Davis. 2010. Active learning projects in services operations management. *INFORMS Trans. Ed.* **11**(1) 20–28.

Cherney, I. D. 2008. The effects of active learning on student's memory for course content. *Active Learn. Higher Ed.* **9**(2) 152–171.

Clemen, R. T. 2009. *Making Hard Decisions: An Introduction to Decision Analysis*, 2nd ed. Duxbury Press, Belmont, CA.

Grossman, T. A., J. S. Norback, J. R. Hardin, G. A. Forehand. 2008. Managerial communication of analytical work. *INFORMS Trans. Ed.* **8**(3) 125–138.

Law, A. M., W. D. Kelton. 2002. *Simulation Modeling and Analysis*. McGraw-Hill, New York.

Liesiö, J., P. Mild, A. Salo. 2007. Preference programming for robust portfolio modeling and project selection. *Eur. J. Oper. Res.* **181**(3) 1488–1505.

Liesiö, J., P. Mild, A. Salo. 2008. Robust portfolio modeling with incomplete cost information and project interdependencies. *Eur. J. Oper. Res.* **190**(3) 679–695.

Mild, P., A. Salo. 2009. Combining a multiattribute value function with an optimization model: An application to dynamic resource allocation for infrastructure maintenance. *Decision Anal.* **6**(3) 139–152.

Renjish Kumar, K. R., I. Hirvonen, I. Leppänen, T. Kuronen, J. K. Nurminen, A. Salo. 2005. An evaluation model for the end-user utility of office applications. *Proc. Ninth IASTED Conf. Software Engrg. Appl.*, ACTA Press, November 14–16, Phoenix, AZ, USA.

Revans, R. W. 1998. *ABC of Action Learning*. Lemos and Crane, London.

Salo, A., A. Punkka. 2011. Ranking intervals and dominance relations for ratio-based efficiency analysis. *Management Sci.* **57**(1) 200–214.