PRE-PRINT VERSION

Stummer C., Gutjahr W.J., Denk M., Riedmann H., Froeschl K.A. (2012) Training on the project: A quantifying approach to competence development. Knowledge Management Research & Practice, 10 (1), 64-78. DOI: 10.1057/kmrp.2011.27

The final publication is available at Taylor & Francis via https://doi.org/10.1057/kmrp.2011.27.

Training on the Project: A Quantifying Approach to Competence Development

Christian Stummer[1], Walter J. Gutjahr[2], Michaela Denk[3], Harald Riedmann[4], and Karl A. Fröschl[5]

[1]Bielefeld University, Germany; [2]University of Vienna, Austria; [3]Statistics Department, International Monetary Fund, Washington, U.S.A., [4]Vienna Knowledge Net (vknn), Austria; [5]University of Vienna, Faculty of Informatics, Austria

Correspondence: Karl A. Froeschl, E-mail: karl.anton.froeschl@univie.ac.at

Abstract

In the project "Competence-Driven Project Portfolio Analysis" (CDPPA), an integrated system for supporting R&D project selection, staff assignment and activity scheduling with special consideration of the strategic development of competencies has been designed and implemented prototypically. The system has been field-tested at the Electronic Commerce Competence Center (EC3), a public-private partnership R&D enterprise. Experiences from this trial application are summarised and discussed, particularly concerning data collection and competence measurement, the benefits and limits of the chosen multi-criteria decision analysis approach, the evaluation of introduced changes to the decision making processes, and the transparency of the formal planning model and its components.

Keywords: Decision support, knowledge management practice, knowledge utilisation, measurement, case study

Introduction

Planning of actions in pursuit of predefined goals is deemed the acme of rationality, particularly in the realm of economics aiming at the efficient allocation of scarce resources. In the post-industrial era, time and knowledge are considered the most critical resources, rather than space and energy, which had been the scarce factors dominating industrial production. While (re-)action speed has become a decisive competitive success factor in the short term (as seen in its perhaps currently most extreme occurrence, the "subsecond" management of world-wide electronic finance transactions), *knowledge* increasingly determines the attainment and sustainability of the longer-term, strategic positions of producers in competitive markets [Foray, 2006]. Hence, knowledge management is conceived as a *critical* managerial concern in HRM towards a systematic development of human capital. Traditionally, learning effects in organisations have been related to efficiency gains resulting mainly from ongoing routine production [Arrow, 1969], yet typically subordinated to the priorities of production management [David, 1999] or facilitated in terms of specific vocational training measures [Snell and Dean, 1992; Afiouni, 2007]. Often, also a more thoroughgoing deployment of knowledge management methodologies fostering knowledge-based cultures of communication and interaction within as well as across boundaries of organisations, or organisational divisions, has been promoted for the sake of raising a firm's competitive potential [Snell et al., 2001; James, 2003]. At long last, however, individuals are considered the ultimate knowledge carriers and, as a consequence, the target of developmental measures concerning the formation and increase of human capital [Argyris and Schon, 1978]. Accordingly, in (HR) management emphasis is laid upon "… the formulation and implementation of strategies for developing, acquiring and applying knowledge, and the monitoring and evaluation of knowledge assets and processes for their effective management" [Quintas et al., 1997]. Nevertheless, the perception of knowledge as some kind of commodity that can easily be managed (that is, grasped and passed on) is delusive, as knowledge in its entirety can hardly be captured and codified without ambiguity [Wilson, 2002] and, thus, the ambition of measuring human resources in terms of *formal* competencies ready for disposition remains inherently tentative and limited.

At any rate, recognising that knowledge, without a physical substrate other than symbolic recordings, is bound in persons, the *human resource* takes centre stage in practical knowledge economics in two respects, namely:

- the efficient allocation of available competencies, and
- the strategic development of competencies expected to become required.

In view of the recent unleashed dynamics of markets and fields of competition, especially the focus on *future* competence demands [Jordan, 2005] gains tremendous importance, aggravated by the relative inertness of the formation of intellectual capital (resembling, in this respect, many features of physical fixed capital, notably its immobility). However, in pursuit of strategic goals of competence development, both aspects mentioned necessarily go hand in hand: only through the intentional utilisation of existing competencies does it become possible to advance competencies towards desired competence profiles at an acceptable cost. Implicitly, of course, this is based on the assumptions that while (i) knowledge and competence can be generated deliberately and intentionally *in principle*, (ii) its (re-)production cost is fairly high and response times of developmental measures taken are considerably long; that is, desired qualifications may not only be scarce at a time, but cannot be procured quickly either because of this inexorable path dependency of new capabilities such as knowledge and competencies on a pre-existing stock of such capabilities.

In this contribution, emphasis is placed not primarily on large-scale knowledge production, but rather on the gradual development of competencies within a single R&D organisation; this will be termed *endogenous* competence development as opposed to the (incidentally possible) acquisition of exogenous competencies developed by someone else. Clearly, in order to be able to acquire missing competencies from external sources those have had to breed them before, so there must always be *some* institution bringing forth the requested competencies. Given the specificity of competencies required (or anticipated to become required), it is typically universities, dedicated research institutions, and especially within-corporate research departments that develop such competencies both for their self-preservation and, with respect to educational institutions, as a supply to the market, e.g., in terms of knowledge transfer. In either type of institution, they are considered to perform a *transforming* function implementing an intentional trajectory of knowledge and competence development.

In the following, *production planning* is regarded as the primary means of controlling and governing competence development. From all feasible activity portfolios identified within an organisation, the portfolio that best matches the targeted competence development in terms of learning effects of the organisation's staff is looked for. This approach undoubtedly presupposes a certain R&D orientation of the activities considered, as this seems to be the only performative context that permits an organisation to establish competence development as its core business.

Modelling competence (cf. e.g., [Spencer and Spencer, 1993; Mansfield, 1996; Shippmann et al., 2000]) presupposes some formal conception of what, from a computational point of view, constitutes a *competency*. In a simplistic attempt, 'competence' is assumed to be decomposable into a finite set of elementary competencies, each of which represents a specific, identifiable kind of skill, whether cognitive, tacit, or social, or some combination thereof. In spite of various undeniable shortcomings of such a simplification, a virtue of the ensuing competence model consists in assigning a standardised – yet artificial – measurement scale to each of the (elementary) competencies obtained in this manner. This scale is good for expressing skill levels or competency scores that can be modelled to vary depending on prior training and on the degree of continued competency utilisation; this degree, in turn, can be expressed in terms of the time any such competency is exercised in (research) activities. Moreover, it is conceivable to also reduce scores for competencies left unused over longer time periods, thus taking account of the degradation ("half-life") of skills not practiced regularly. Thus, in the proposed model context, to talk about transforming competencies is tantamount to intentional efforts of raising the individual or collective scores of some isolated competency, incidentally at the cost of decreasing the scores of other competencies expected to be of declining future importance. Computationally, then, competency scores provide a (ratio-scaled) numerical value representing skill capacity.

The natural and, hence, predominant approach towards high-tech competence advancement is to carry out *R&D projects*. A project consists in a temporally limited effort in pursuit of well-delineated targets using a circumscribed set of resources. Practically, therefore, a research organisation manages its day-to-day business by processing a *project portfolio* (that is, a set, or sequence, of individual projects) composed of projects (i) drawing from a – limited – set of (competence) resources and (ii) contributing to strategic goals of the organisation expressed in terms of one or more measurable targets. Taking account of the economic rationality postulate, project portfolio composition amounts to an *optimisation* problem. Obviously, an optimal project portfolio

- maximises the organisation's overall utility, or value gain, according to predefined measurable objectives, and
- is constrained only by the available resources – that is, competencies assignable to projects – effectively determining the size and structure of feasible project portfolios.

While this plainly restates the standard set-up of optimisation methodology [Churchman et al., 1959], in the problem instance considered it must however be borne in mind that the optimisation objectives defined typically include the objective of endogenous competence development. In a sense, it might be argued, this optimisation problem turns – almost – circular: the objective function links back to the resource constraints, if only in affecting the feasibility of *intended* future project portfolios. Exactly by implementing the projects of a selected portfolio, new skills are attained, contributing to the (then extended) available competence resources. Thus, a research organisation conducts research mainly to improve on its activity, that is, to being better able to continue conducting (better) research – seemingly a concept of self-sufficiency.

An obvious problem associated with project portfolio selection originates from the inherent complexity of the planning task. It is for this reason primarily that computational approaches are considered worthwhile in spite of incurring, in turn, the complexity of modelling the planning domain first. Acknowledging the notoriously intricate management of human resources, advanced models of multi-criteria optimisation promise tractable solutions to this planning task, once the tuning parameters and heuristics of the optimisation procedure are calibrated adequately. It has been the ambition of the CDPPA (short for "competence-driven project portfolio analysis") project, reported on in this paper, to explore the planning domain and devise a practical and computationally tractable approach towards project portfolio selection, capitalising on the virtues of (i) optimisation meta-heuristics to reduce the search space to Pareto-efficient project portfolios, and (ii) graphical user interfaces to analyse this Pareto frontier of project portfolio candidates in view of generally conflicting yet simultaneous managerial objectives such as competence advancement, aspired utility levels, budgetary targets, and so on (see [Gutjahr et al., 2007, 2008, 2010; Stummer et al. 2009]).

During the CDPPA project, a practical application scenario has been referred to throughout, namely the Vienna-based "Electronic Commerce Competence Center" (EC3), a public/private partnership R&D organisation committed to applied research and experimental development in advanced business models and solutions in the domain of digital commerce. In particular, the experiences documented in the present paper originate from sample computations and interpretations of results obtained from this field trial.

The remaining sections of the paper are organised as follows: Section 2 presents an outline of the CDPPA approach towards both the formal competence model employed and the empirical acquisition of competency scores as embedded in the socio-cultural context of a research organisation. This section also covers the implemented optimisation methodology and sketches the interactive decision support tools provided by the implemented CDPPA software, including a brief review of the input data structures describing both the objective functions and the project candidates. Section 3, then, is devoted to an in-depth review of the empirical findings of the EC3 field trial using CDPPA, focusing also on the wider implications of the proposed model and approach with respect to human resource management in scientific institutions. The concluding section restates the main findings of the CDPPA trial application.

The CDPPA Framework

Modelling and *quantifying* the competence structure of a (research) organisation is a fundamental prerequisite for the rational management of the organisation's (R&D) project portfolio in order to achieve its strategic objectives in terms of competence development. Besides the available competence resources, the competence requirements of potential projects, as well as the contribution of projects to strategic competence objectives must also be accounted for.

A formal competence model

Approaches to measuring competencies basically emerged from two fields of research, viz. management science and cognitive science. In contrast to management science and its primarily economic focus on human resources, other disciplines, such as psychology, educational science, or neurosciences, emphasize the individual, often with a "biographic" view of competence and its personal development. In the context of the CDPPA project, competence measurement takes place at the level of the individual employees since, aside from the selection of Pareto-optimal project portfolios, an appropriate assignment of researchers to projects (staffing) is required. To this end, a human-centric approach that focuses on the recording and assessment of individual competencies, emphasize (explicit) knowledge, experience, personal characteristics, and the application of (explicit) knowledge and experience as constituents of competence (cf., for instance, [Sundberg, 2001]) has been opted for.

Pursuant to the HR-XML standard [HR-XML Consortium, 2007], the CDPPA competence model is based on the concepts of competencies and evidences as significant proofs or indicators of competencies: a 'competency' denotes the basic entry in a – manageable – *competencies catalogue* established with regard to human resource planning and development. Within a competency considerable flexibility, extensibility, and summability are supposed; the focus is on fundamental understanding that allows quick adaptations to closely related problems and tasks rather than on specific technical skills. For instance, a person who is competent in one high-level programming language is assumed to rapidly familiarise with another high-level programming language (that is usually conceived in a similar way). The competencies catalogue encompasses competencies from all four principal competence classes according to the categorization by Erpenbeck and Heyse [1999].

With respect to evidences, material (or objective) and subjective evidences are discerned. Material evidences particularly encompass formal qualifications such as university degrees, certificates, diplomas, letters of reference, patents, as well as professional experience and "idle times" of competencies, the latter as a kind of negative evidence. This is in line with the distinction between training/education and experience, as well as the concepts of growth and (in-)stability of competencies in Sveiby's Intangible Asset Monitor [1997]. In contrast, subjective evidences comprise assessments of the level of a competency by a person herself, by colleagues ("peer assessment"), and by supervisors. An *evidences catalogue* has been developed that appropriately systematises acceptable proofs of competencies relevant in view of the competencies catalogue underlying the model.

The relevance of each evidence item for each competence has to be specified according to the basic assumption of the CDPPA approach that the degree to which a person possesses a certain competency is quantifiable in terms of a *competency score*. This includes the specification of apt evidence and competence measurement scales, as well as of the joint contribution of a set of evidences to a particular competence (embracing knowledge depreciation rates and learning rates) also taking into account the substitutability of evidence (e.g. more professional experience versus a superior level of education). Objective and subjective evidences are treated differently.

The *objective* competency score of an employee and a particular competency is defined as the sum of the contributions of all objective evidences of the employee constrained to the interval [0, 100]. The contributions of the objective evidences are specified as an evidence-by-competency score matrix based on background information such as curricula or journal citation indices and the rule of thumb that the score contribution of a master's degree should be approximately the same as the score contribution of three to four years of research experience in the same competency. Intramural experience is rated higher than external professional experience. The depreciation rate is set at a rather optimistic level. Both, the learning as well as the depreciation rate, are provided as score per period and – for the time being – are kept constant for most of the competencies, except for several methodological competencies that are supposed to grow and diminish more slowly.

The *subjective* competency score is defined as the weighted average of the scores obtained from the underlying subjective evidences, i.e. the self (50%), peer (20%), and management (30%) rating of the competence. In analogy to Dreyfus and Dreyfus' skill acquisition model [1986] these assessments are collected at a six-item ordinal scale discerning the categories 'no competence', 'novice' (requiring permanent guidance and supervision), 'advanced beginner' (requiring temporary guidance and supervision), ‘competent performer’ (rather independent task fulfilment), ‘expert’ (absolutely independent problem solving, decision-making, result responsibility), and ‘mentor’ (normative, guidance and supervision of others). All ratings are transformed to the interval [0, 100]. The ratings of all peers and the ratings of all supervisors, respectively, are then aggregated to one ‘peer score’ and one ‘management score’ by means of a weighted arithmetic average with the weights determined by the self-ratings of the respective peer/supervisor. The overall competency score is then computed as weighted average of the objective (75%) and subjective (25%) competency scores.

In order to determine the *competence efficiency* of an employee with respect to a particular competency, the competency score is regarded in relation to the time it takes an employee to carry out a task that requires *only* the competency concerned. The competence efficiency can be interpreted as the share of work performed in one time unit by a person on a task merely necessitating one particular competence compared to a “norm-efficient” person, that is, possessing perfect knowledge and skills in the respective competency who would perform the task in a single time unit. This implies the substitutability of a higher competency score by a higher temporal effort and reflects the appropriate assignment of competence holders [Sveiby, 1997]. For personal, activity-oriented, and social-communicative competencies, the substitutability assumption does not seem appropriate; it appears only meaningful for the professional and methodological type of competency. The competence efficiency is calculated by a monotonous transformation of the competency score to the interval [0, 1]. The class of logistic functions that is frequently used for modelling organisational learning (e.g., [Chen and Edgington, 2005; Ngwenyama et al., 2007]) represents a promising choice for this purpose. The parameters of the logistic function are chosen as identical for all competencies and their values are defined based on the specification of the input value domain and the assumption of a high increase of efficiency for medial competency scores in contrast to small gains for rather low and rather high competency scores. In this context 'rather low' is tantamount to a competency score below the score that is assigned to graduation (approximately 30 to 40), i.e. competencies trained at university level are located at the lower bound of the range of competency scores with high gains in efficiency. On the other hand experts with a long record of formal qualifications and/or research experience and, thus, a high level of efficiency do not gain much in efficiency anymore.

Based on the CDPPA competence scoring model, each employee gets attributed a *competence profile* that consists of the vector of competence efficiencies indexed by the used competencies catalogue. To this end the input data for the model, i.e. the specified objective and subjective evidences, have to be collected for all employees. In addition, the parameters of the model (in particular, the score matrix and the parameters of the logistic function) have to be specified as outlined in the previous section. Moreover, as a prerequisite for the entire data collection process, a competencies catalogue and an evidences catalogue – for use also in expressing project requirements formally – must be devised suitably which, as a matter of fact, is a rather time-consuming task, at least when carried out for the first time.

Project portfolio selection

Proponents of the resource-based view of the firm provide arguments for the "usefulness of analysing firms from the resource side rather than from the product side" [Wernerfelt, 1984]. Yet, in a production function resources are assigned to activities implementing the (economic) goals of an institution; in the case considered, these activities refer to R&D projects, to be chosen from a candidate set of project opportunities in favour of predefined objectives.

Projects, project opportunities, and the planning frame

A *project* is considered a structured endeavour composed of different tasks, each of them requiring a special "profile" of competencies. The task set generally bears sequence dependencies, including temporal constraints such as ready times, (estimated) durations, and due dates, for each task. These constraints may be of stochastic nature, or known only partially. For analytical purposes, tasks are split further into distinct *jobs* each of which gets assigned a specific competency listed in the competencies catalogue; to fulfil a job, the competency must be provided with a particular quantity expressed in terms of norm-efficient competency time units. Tasks, in turn, may be grouped into work packages representing substantive subgoals of a project. In addition to competency requirements, projects may feature further attributes relevant for the portfolio selection process, such as monetary volumes, project realization probabilities, a general utility value the project might generate when implemented, etc.

At planning time, a set of *project opportunities* from which to choose a portfolio – that is, a subset – is assumed as given. Each project opportunity represents an option to implement with some probability, provided sufficient resources are available in the opportunity window (that is, the time interval in which the project would be feasible in principle) of such a project candidate. In formal terms, the set of project opportunities generates a combinatorial space to be scrutinized in the planning effort. Obviously, while the identification of appropriate project opportunities is both burdensome and critical, their explicit specification is by no means trivial.

In addition to the competencies catalogue defined above, the *planning frame* is further structured by

- a discrete time scale resolving the temporal axis into equidistant time slices, or planning periods (such as months), respectively;

- limiting the planning horizon to some finite value (depending on the maximally tolerated project make-span, such as 24 to 48 planning periods);
- assuming a fixed stock of resources within the planning horizon (in particular, that the workforce remains constant);
- considering as constant the individual employees' competence efficiencies during a project assignment (that is, competency scores are updated after termination of a project only), and
- assuming some idle resources left for assignment to further project candidates without forcing the cancellation or interruption of projects already underway.

By definition, *planning time* always (re-)sets $t = 0$, implying a shift of the (current) planning horizon and, because of ongoing projects, a recalculation of resource levels actually available at planning time. Clearly, in a dynamic environment, after a project portfolio has been dispatched to the shop floor, new project opportunities may arise occasionally, suggesting portfolio revisions. To provide the updated candidate set, the new project candidates are pooled with those portfolio projects that have not yet commenced at the reset planning time. By convention, the withdrawal of projects, once started, at planning time is considered infeasible.

Defining the solution space

To complete the definition of the CDPPA solution space, decision criteria have to be introduced as a further model component for the sake of ranking alternative project portfolios selected from the pool of project opportunities available at planning time, provided that these portfolio alternatives meet all task and resource constraints. Basically, two classes of criteria are discernible, viz.

- *portfolio performance objectives* ('ppo', for short): these are functions of the project portfolio alone and deliberately do not consider portfolio impacts on competence development;

- *competence performance objectives* ('cpo', for short): these represent formal expressions explicitly measuring (aggregate) changes in competency scores, or efficiencies, respectively, achieved through implementing portfolio alternatives.

Typically, ppos are used to maximize expected utilities, involving parameters such as monetary project volumes, third-party funding raised, overall resource utilization, etc., relating to routine business goals. However, the decision maker's interest in purposively changing the stock of competencies through appropriate portfolio selection is covered by adding cpos. In particular, the choice of cpos reflects – longer-term – strategic goals such as a gradual adaptation of the stock of competencies to enhance an organization's fitness for upcoming project opportunities. Conceptually, cpos are framed in terms of *competency baskets*, resolving the devised competencies catalogue into a weighted composition of elementary competencies; thus, a competency basket represents a kind of "macro-competency" integrating a relevant combination of constituent competencies (much like a university degree or a vocational qualification). Used as a formal objective, a competency basket then measures the "progress" of competence development restricted to the set of elementary competencies included. If, for example, a research organisation intends to engage in advanced ground-based astronomy instrumentation, a competency basket comprising a range of competencies highlighting knowledge and experience in adaptive optics might be a reasonable choice for a pertinent cpo.

For quite obvious reasons, ppos and cpos tend to oppose each other, since investments into competence development strategies – in other words, opportunity costs of training – are traded off against shorter term economic gains achievable by more opportunistic (that is, less strategy-driven) portfolio selections ignoring longer-lasting effects on competence structures. As ppos and cpos, in general, are incommensurable, *multi-criteria* decision making is entailed.

Project portfolio selection is linked to the constrained resources (that is, the stock of available competencies at planning time) through a staffing procedure determining the allocation of available competency quantities – in terms of employee work time – to project tasks according to the competency requirements of the projects of the selected portfolio. Provided that the staffing procedure terminates with a complete schedule, or work plan, consistent with all imposed problem and resource constraints, the portfolio is *feasible*. Due to the tremendous complexity of determining the exact feasibility of work plans, the staffing procedure [Gutjahr et al., 2008] resorts to various assignment heuristics and constraint approximations (such as an even allocation of resources within scheduled task intervals).

In terms of competency accounting, the choice of time for evaluating objective functions, especially for cpos, is critical: for each (future) evaluation time, the net effect of a project portfolio on competency scores depends, according to the work plan assignments, linearly on the degree of fulfilment of projects in a considered portfolio.

Interactive decision support

In multi-criteria decision making, instead of a unique optimum only a *set* of Pareto-optimal solutions can be identified. Accordingly, a rational decision maker will always choose a Pareto-optimal solution. Presuming that, in general, the size of the set of feasible solutions for a real-world application scenario may range in the millions (or even billions), the restriction of the solution space to its *Pareto frontier* entails a significant reduction of decision complexity. However, the Pareto frontier may still remain impressively large, impeding a direct comparison of its solution candidates. As an effective means of exploring entangled Pareto frontiers, a combination of

- visual inspection based on *graphical displays*, and

- an *interactive mode* of exploration, directly linked to the visual information display,

seems to provide a reasonable approach in managing cognitively heavy-loaded decision making environments.

In the CDPPA approach, from amongst a variety of display formats, ranging from exhibits of multivariate aggregates (glyphs, faces, boxes, trees or castles; cf. [Chernoff, 1973; Kleiner and Hartigan, 1981; Korhonen, 1991]) to presentations facilitating direct manipulation in setting aspiration levels (e.g., the spider web diagram proposed by Kasanen et al. [1991], Lotfi et al. [1992], or Stummer and Heidenberger [2003]), five modes of interactive graphic interfaces – all of them supporting the decision maker in exploring the solution space of Pareto-efficient project portfolios from different points of view and/or with different means of interaction – are drawn, viz.:

- *Heat maps* (also known as "coloured matrix plots" [Cook et al., 2007]) enable a quick sorting of portfolios for any selected decision criterion and thus help to gain a rough overview of portfolios worth closer examination; however, heat maps lack topological interpretability.

- *Parallel coordination plots* juxtapose the decision criteria used, representing portfolio performance by "profile lines" across the different criteria exhibited; thus, patterns of positive or negative interrelationships are quite easily detected.

- *Interactive column charts* provide similar functionalities, used preferably for directly comparing a rather small number (7±2) of portfolios, displaying columns, for each portfolio, side-by-side for each of the selected decision criteria.

- *Line charts* help visualize the competence development implied by candidate portfolios plotting competency scores against time. Using drop-boxes, the decision maker can vary the scope of information displayed from aggregate levels down individual employees and competencies.

- *Competency maps* are cross-tabulations of colour-coded cells expressing the competence efficiencies by employee and competency, respectively. The maps show the distribution of accumulated competency scores at a glance, and may also be animated to view the dynamics of this distribution in terms of either absolute or difference competency scores.

Interactivity implies that a decision maker is free to use any combination or sequencing of graphical displays to gain insight into the structure of the Pareto frontier of a given problem space. A more detailed description of the proposed interfaces, as well as an illustrative prototypical use case is provided in [Stummer et al., 2009].

Learning by Doing

A prototype version of the CDPPA system was implemented during 2006–2007 in joint work of the University of Vienna with the EC3, and tested based on (close to) real-world data. This section of the paper summarizes the most salient findings of the field trial, including both data acquisition and decision making phases; for a more detailed account of the socio-economic cross-evaluation of CDPPA; cf. [Riedmann, 2008ab].

Transparency of the planning model

Routine utilisation of a formalised planning model like CDPPA strongly depends on the belief in the model's reliability: all presentations must apparently relate to the set of inputs, and the findings displayed have to meet intuitive expectations, at least after providing clues to resolve initially counter-intuitive perceptions. In other words, the tool is judged *transparent* if it establishes (or, even better, increases) confidence in the decisions made with its help. However, transparency is fairly hard to assess, not the least because of the inherent complexity of the CDPPA model.

Three major components determine the cognitive tractability of system behaviour: two of them refer directly to problem parameters such as competency scores and optimisation criteria, that is, data quality issues, while a further source of opacity is the scheduling algorithm used for project staffing based on recorded staff competencies.

Competence model and calibration

The decomposition of competence into elementary competencies – quite in the spirit of Taylorism (e.g., cf. [Gilbreth, 1912]) – probably represents the strongest methodological presupposition of the CDPPA approach. While lending itself favourably to a formal modelling of competence dynamics, the deliberate and simplifying ignorance of competency *interactions* may have rather detrimental effects: it is pretty obvious that positive externalities prevail on both individual and collective levels. More specifically, depending on the chosen decomposition of competence, learning effects almost certainly correlate amongst elementary competencies. This places the choice of elementary competencies, or the design of the competencies catalogue, respectively, at the centre-stage of CDPPA: while a more detailed catalogue of competencies suggests a better reflection of individual competence profiles, doing so increases the risks of both introducing quantification artefacts and underrating competence correlations (that is, generating inconsistent profiles); conversely, a more condensed set of competencies may fall short of differentiating among salient job assignment criteria required for project staffing crisply reflecting portfolio selection goals.

Most probably, this dilemma cannot be resolved on theoretical grounds, but rather needs practical experience and calibration. Unfortunately, the EC3-based field trial had to remain limited to a small R&D staff using a restricted set of project candidates and lacked the resources to extend over several successive planning iterations, all of which together might provide a more supportive environment for model evaluation and calibration. In particular, choosing a competencies catalogue with 80 elementary competencies certainly overshot the mark in the EC3 case; with hindsight, this led just to an undue computational burden in the portfolio selection stage, since many "core" competencies required by project candidates happened to be very scarce resources that could hardly be substituted. Nevertheless, in other circumstances, an even larger number of competencies may well apply. As a tentative conclusion, the *scope* of the competencies considered – rather than simply their number – matters; presumably, this scope depends critically on the respective application scenario and, thus, needs case-based tuning.

Besides the reductionist competence decomposition accomplished in whatever way, there obviously remains the common quantification problem. In the field trial, although mapping evidences to the competencies discerned represents a best effort, little confidence in the general validity of the established score conversions is justified as yet. This caution applies to both absolute scores and relative score differentials assigned from the evidences considered; by the way, the types of evidences (to be) taken into account themselves are, of course, another source of uncertainty. In this respect, external effects on a person's competence profile emanating from informal ("off duty") activities are particularly hard to come by, although personal intellectual or skill development may generate considerable impacts on occupational performance as well.

A further crucial influence on the planning model is exerted by the model parameters controlling competence *dynamics*. This refers to both the initial conversion of evidences to numerical scores expressing competence efficiencies, as well as subsequent score updates in the wake of project assignments. Obviously, because of their cumulative net effects, an appropriate choice of multiplicative and shape parameters of competence scoring and update functions, respectively, is essential for model transparency. For one thing, an easy check of plausibility involves the stationarity of average competency scores attained under the condition of implementing "routine" project portfolios: while no growth or fitness effects are to be expected then, neither should a serious drop in overall levels take place. However, to run such a check validly, the model must have been in use for a while, as otherwise the observable effects may be superimposed by initialisation conditions. Generally, of course, due to the restriction to only endogenous competence development, the intended increase of score levels for some selected competencies is almost by necessity accompanied by losses for other competencies (as these receive less project assignments).

For the reasons indicated, the limited CDPPA field trial could not provide a final answer to the question of the best choice of values for the competence model parameters, nor do modelling experiences warrant generalisation in any direction. Clearly, both improvement and depreciation of competency scores may vary over competencies as well as individuals, and may additionally depend on further organisational and social features of the respective R&D environment.

Optimisation criteria

A major element of CDPPA towards increasing decision rationality consists in the formal expression of planning targets. Clearly, the decision maker's ability to state decision criteria faithfully representing those planning targets is presupposed. However, it has to be taken into account that "decision makers do not always know what they want" [Selten, 2001: 24] and, moreover, explicating decision criteria incurs an additional effort which – according to rationality standards – has to be offset by improved decision performance, at least in the longer run.

In past planning practice, at least in the EC3 context, decision criteria were rarely expressed in terms of objective functions ready for numerical evaluation. Rather, criteria have been matched implicitly to the centre's production organisation consisting of rather stable teams bundling specific competencies (for natural language processing, workflow modelling, data mining, etc.), providing a structural planning frame removing most of the combinatorial complexity of project elicitation and resource allocation. Accordingly, stipulating a switch to the ex ante specification of planning targets requires a deep and abstract reflection particularly of cpos (competence performance objectives) in terms of both, quantitative demands of competencies for each project under scrutiny and the eventual contribution of assigned competencies to strategically set competence growth and fitness objectives:

- *resource quantification* may quickly turn into quite a tedious (if not unmanageable) task, mainly because of accuracy concerns in anticipating project structures and resource allocations, while
- matching initially qualitative strategic development goals of an organisation to *competency baskets* (that is, weighted mixtures of competencies) presumes an articulated proficiency of the decision maker as to the goal-impact of each of the competencies mapped.

In either case, the emerging quantification demands challenge traditional practices, which used to emphasize figuring out sensitive clues of a planning scenario as inputs to problem-fitted decision heuristics helping to efficiently discriminate satisfactory project portfolios from poor ones – in this way implementing a kind of "ecological rationality" [Gigerenzer, 2001] typical of coping with complex decisions lacking enough structure to admit formal analysis, and shortcutting intricate and costly quantitative analysis by often nonetheless highly effective rules of thumb.

As a consequence, any direct comparison of planning and decision modes tends to be misleading: the application of the CDPPA model entails a shift of effort from making first-level decisions towards the preparation of quantified model inputs which, at any rate, cannot be simply predicted to benefit the eventual decision outcomes. Rather, the direct utility materializes in terms of a dramatic gain in power for exploring visually and interactively the structure of the ensuing problem space, typically comprising many more project candidates (and, thus, Pareto-efficient portfolios) than available for scrutiny without the preceding formalisation effort. However, the *process* of decision making becomes entirely redefined on the condition that critical problem inputs (such as the set of project candidates and their resource demands, the decision criteria, performance evaluation time, etc.) are indeed easily accessible for adaptation. Accordingly, the real potential of CDPPA must not be seen in replacing the decision maker, but rather in replacing the traditional mode of making decisions in that CDPPA facilitates the search for alternatives and comparative clues in terms of a highly supportive (that is, a learning) environment of enhanced ecological rationality.

Scheduling/staffing

One of the most difficult subtasks of the project portfolio selection process concerns the validation of portfolio *feasibility*. Without effective calculation support, a manual validation is entirely out of scope, as the task is known to be computationally (NP-) hard and, hence, intractable beyond very moderate problem sizes even for mechanised discrete optimisation (e.g., cf. [Blazewicz et al., 1983]). Therefore, heuristic approaches are taken recourse to. From the planner's point of view, the scheduling heuristic simply operates as a black box. Still, to understand the implications of the – fairly conservative – CDPPA scheduling meta-heuristics, the simplified constraint management of the scheduler has to be kept in mind. In particular, sometimes otherwise Pareto-optimal project portfolios may be excluded because of almost negligible (real or approximated) resource violations – but, in project definition, there is always some leeway to argue. Evidently, the feasibility of even a single project may have a tremendous effect on the Pareto frontier, hence it is of great importance to scrutinize projects *excluded* from portfolios: based on the identified cause of an exclusion, project descriptions may be (reasonably) tuned until a coherent picture results. Accordingly, the generated solution space needs critical checking that so far has not been supported particularly well in the interactive user interface, in spite of the apparent potential in improving problem comprehension and planning confidence.

A serious limitation of the scheduler, as implemented in the prototype system of CDPPA, regards the composition of *project teams*. For quite obvious reasons, the production function of creative outputs is multiplicative [Caves, 2000], generally implying a labour division scheme mixing together the competencies of different researchers teaming up for a project. Accordingly, feasible project portfolios are composed of projects with job bundles assigned to each project, albeit CDPPA is only dealing with the internal structure of these job bundles in a rather crude way [Gutjahr et al., 2008: subsection 2.3]: in practice, however, there is a penalty to highly fragmented competency distributions in project teams (entailing increased intra-project coordination efforts). Because of these non-negligible coordination cost (and, perhaps, other reasons like social preference etc.), comparably small-sized “core teams” that neatly integrate quite self-contained competence patterns matching typical project requirements tend to evolve. In a R&D context, such core teams likely coincide with an organisation’s core competences. Thus, a sound resource allocation heuristic would first try to assign established project teams to candidate projects, augmenting these teams with additional staff contributing the required project competencies the allocated core teams cannot cover by themselves. As a matter of fact, team building is yet another useful tactics towards competence development, but currently not accessible – through specific cpos, for instance – in implemented CDPPA scheduling heuristics. While this is in perfect concordance with the competence summability assumption maintained in the CDPPA competence model, it reflects the tacit CDPPA concept of competencies as *functional* inputs (that is, production resources with limited availability) as opposed to a more behavioural conception better reflecting the aspects of the social construction of competence, which relates competence to the organisational and social setting where (team-) work is actually conducted (cf. [Delamare Le Deist and Winterton, 2005: 31] and the literature cited therein).

Data Procurement

Contrary to more informal project portfolio selection approaches, the CDPPA model rests on a fairly extensive set of input data, comprising (i) the competence profiles of all R&D staff, (ii) a set of formalized decision criteria (objective functions), and (iii) a sufficiently broad set of project opportunities from which to select Pareto-optimal portfolios. Apparently, there is a trade-off between the gains in structured decision making and the effort in collecting all data necessary to make use of the model. While a link to established HR management tools is advisable (particularly for the assessment of R&D staff competencies), other input data generally have to be generated from scratch, perhaps with modest workflow and electronic data management support.

Empirical acquisition of staff competence profiles

For the EC3 case study, both competencies and evidences catalogues were developed in a multi-stage social process in order to reach a consensus on the catalogues prior to the actual measurement of competencies. A first draft of both catalogues was validated by way of an employee survey. The availability of the competencies and evidences, the comprehensibility of the items, potential additional items, and the willingness to be assessed with regard to human factors were questioned. In addition to the electronic questionnaire, a handout summarising both the objectives and the main ideas of the project was provided.

A set of personal, activity-oriented, and social-communicative competencies was subsumed under the heading 'human factors', contrasting the set of professional competencies. Since the latter constitute the major precondition for a reasonable assignment of employees to projects or even the feasibility of projects, this set was specified at a much greater level of detail, with approximately 90 competencies. Based on the results of the survey among EC3's researchers, the draft underwent several adaptations. In particular, only the professional competencies were kept, mainly due to the problems concerning the 'human factors', especially the missing consent of the majority of the employees to be rated with respect to these competencies and related questions of privacy. Eventually, the competencies catalogue still consisted of 80 competencies from the triptych of empiricism, economics, and technology characteristic of EC3.

Essentially, apart from minor rewordings, the evidences catalogue remained stable after the employee survey. The final version encompasses 56 evidence items divided into objective and subjective evidences. The objective evidences contain (i) formal qualifications, in particular secondary (vocational) education, master's and doctoral degrees in different disciplines relevant at the EC3, academic theses and scientific publications, major field(s) of study, patents, trade certificates, vocational entitlements, and (ii) professional experience, structured into entire professional experience in person months, relevant professional experience (at the EC3, external, international), and duration of disuse (i.e. idle times as indicator for a negative development). The subjective evidences encompass the assessment of the competence level by the competence holder herself, by the peers from the same research group, and by the scientific director.

After having finalized the competencies and evidences catalogues, the objective and subjective evidences had to be gathered from all employees to enable the computation of each researcher's initial competency score, as well as the respective efficiency values. Furthermore, the researchers' disposable capacities for the planning horizon and the definite gains in competency score and efficiency from the projects that had already commenced (or were approved at least) for each period in the planning horizon had to be estimated. The objective and subjective evidences were collected from the 28 members of EC3's scientific staff, including the six heads of the research groups into which EC3 was structured, as well as the scientific director and six allied freelancers, in addition to the 15 full- or part-time permanent researchers. Seven persons refused to provide the objective evidences and approx. 4% of the subjective scores were missing. Formal qualifications were supplemented if possible. Subjective scores as well as the objective scores still missing were imputed by arithmetic means. The acquired data were then simply fed into the competence model and transformed to competency scores and efficiencies. The disposable capacities, as well as the definite gains in competency score and efficiency from yet ongoing projects, were estimated from data of previous research periods.

In practice, objective evidences might be directly extracted from a human resource database instead of being collected tediously from the personnel. In any case, this part of the survey is required once only. In future periods, updates account for recent formal qualifications. Additional professional experience should be computed from the assignment to project tasks and, thus, from labour time recording systems. Subjective evidences should be gathered at regular intervals. Overall, the routine effort to keep the employees' competence profiles up-to-date is assumed to be considerably lower than the start-up effort.

Project Data Collection

For obvious reasons, the portfolio selection process brings to bear its virtues only if it is fed with reasonably large or well-compatible sets of project candidates, lest poor capacity utilization is the likely consequence. While, traditionally, project opportunities are screened quite intentionally (that is, candidates are chosen in the first place with an expectedly good fit into rather small sets of prospective projects already at hand or on the verge of implementation), the CDPPA approach suggests gathering as many and as varied project candidates as possible before applying the

Pareto-selection process based on the specified decision criteria. However, identifying and working out promising project candidates is itself quite costly in terms of resource contention, and is further complicated by the inherent difficulty of providing and assessing many of the formal and quantitative parameters, such as

- the envisaged project implementation structure (that is, the work plan of tasks, task dependencies, time estimates, etc.);
- competency requirements for *each* task of the work plan;
- various project performance data (such as turnover, raised third-party funding shares, etc.) depending on the respective objective functions used; and
- project "stochastics" (probability to start and finish the project, time and cost estimates, assessments for almost every project parameter considered).

Concerning these abundant demands of project specification, it certainly comes as no surprise that a trade-off emerges between the accuracy of project descriptions and the attainable size of candidate sets. Moreover, in many cases, such a detailed ex ante specification of projects may turn out next to impossible without introducing unreasonable guesses spoiling all the aspired benefits of formalization. In particular, the tight selectiveness of decision criteria suffers if project parameters are set rather arbitrarily or only on very coarse (for instance, work package rather than task) levels of description.

Actually, this trade-off shifts the burden of optimisation to a "second order" decision making [Klein, 2001], namely that of *stopping rules* determining the limits of the search process for candidates. Even worse, the *focus* of search for candidates may interfere with the formal decision criteria used in such a way that an ill-defined portfolio selection problem might ensue: the defined objective functions simply may fail to discriminate between feasible portfolios, since all of them yield very similar scores, resulting in a rather non-robust Pareto frontier. This interference is quite likely to happen given that, in practice, many projects are proposed at short notice (for various operational reasons) without having in mind the longer-term strategic objectives of the organization mapped to the set of agreed-upon decision criteria.

Benchmarking the pilot

The assessment of achievements rests on (i) a baseline and (ii) an alternative. Unfortunately, no direct comparison between the exercised practice and CDPPA as a contrasting future practice can be demonstrated, as there is no obvious common point of reference in the sense of a *tertium comparationis*. For one thing, in spite of the efforts to model the planning framework as accurate as possible, many salient contingencies are hardly known, difficult to grasp, or simply too complex to admit high fidelity representations: as a consequence, the powerful machinery of optimisation tends to solve problems somewhat different from those originally posed. As such, this is no disadvantage, though, as the alternate mode may perform superiorly. What we are left with, hence, is to assess (i) the compromises to put up with in problem representations and their eventual effects, and (ii) the gains, or losses, one is likely to incur by switching to the alternative mode.

For cogent reasons, the CDPPA model is embedded in a wider decision support context; to consider CDPPA from a purely formal or technical perspective would disregard the broader implications of integrating a mechanical procedure into a complex managerial environment. To judge the validity and merits of the proposed decision support tool, its relation to existing planning practices and the various “interfaces” ensuing must be scrutinized. As part of such a comparison, however, the existing practices themselves receive particular attention in exploring the borderline between experience, intuition, and in fact deductive elements of the decision making process.

Model fidelity and usability

The difficulties in appropriately mapping decision-relevant problem features to model parameters have been described already; what remains to be discussed is the question in how far the interactive accessibility of the solution space (as outlined in subsection 2.2.3) lives up to the claim that decision quality improves with the degree of coincidence between task representation and problem visualization. Apparently, in offering a whole range of displays useful in highlighting discriminative criteria, CDPPA literally offers unprecedented viewing angles on the solution space, providing the decision maker a deep and variegated exploration of its topology. Because of the abstractness of this topology, though, the interpretability of visual clues strongly depends on their fidelity in representing meaningful features of solution candidates (that is, project portfolios). Lacking prior experience and easy means of comparison, trust in the fidelity of the model mainly builds on (i) the presence of crucial and versant decision parameters, (ii) the concordance of model assumptions (about competencies and their dynamics, staff assignment, etc.) and acquired problem intuition, and (iii) a reasonably realistic and accurate representation of candidate projects generating the solution space.

The experimental application of the CDPPA pilot testified to the usefulness of the visual presentation of results; particularly, offering *different views* cannot only be judged very helpful but, in addition, the option to change perspectives indeed seems to facilitate *intuitive* insight into the often rather intricate underlying interrelationships of problem parameters. This links to the observation of expert decision making relying more on the perception of relevant patterns rather than meticulous fact analysis [Shanteau, 1992]. Conversely, there is no doubt the real danger of permanently overwhelming the decision makers' intuition just because of the wealth of information delivered: it takes quite some time to familiarise oneself with all the inspection tools and to align the various clues provided into a coherent picture.

In a field trial, general CDPPA model characteristics and the properties of specific candidate project sets indispensably superimpose, implying a methodological difficulty in sorting out the peculiarities of a single application instance from invariant model traits. Regarding the effort it takes to prepare a consistent set of CDPPA input data, a natural way of differentiation consists in varying a given set of candidate projects and project description data, respectively, while keeping model parameters and the competence database fixed. Then, by tracking back selected portfolios to related input projects, as well as visually analysing the ensuing Pareto frontier, the CDPPA portfolio selection process itself becomes more understandable. Methodologically speaking, the creation of *what-if scenarios* provides a useful augmentation of the CDPPA decision support framework, as the experimental variation of input parameters and matching them to observed differences in the output helps to gain subjective confidence in the fidelity and reliability of the portfolio selection mechanics of the system, and provides an effective means of assessing the robustness of solutions in the face of the inherent inaccuracy and tentativeness of many input project parameters. In fact, the CDPPA pilot application quickly revealed implausible input data several times, suggesting repeated revisions of original project specifications. The same reasoning, essentially, applies to the decision criteria: again, the tangible effect of each of the criteria is explored easiest by stepwise variation. As a consequence, the extension of CDPPA core functionality to support the management of problem variations is to be highly recommended.

In principle, the CDPPA user interface can fairly comfortably handle problem sizes involving up to about 40 project candidates, and can probably cope well with even larger sets. However, it is not known how well the (meta-)heuristics combining multi-objective portfolio selection with staff assignment and task scheduling scale up to larger problem instances of, say, 200 candidate projects, as they may occur in some industrial applications. For combinatorial reasons, the tractability bottleneck is caused by the staff assignment and scheduling component. Tests indicate that the implemented planning heuristics have a hard time in dealing with as many as 80 single staff competencies, giving rise to unreasonably fragmented job assignments and, accordingly, quite unrealistic schedules. Little wonder that the ensuing project portfolios deviate, sometimes even grossly, from what one might expect of an intuitive decision. On the other hand, it should be kept in mind that in practice, project selection rarely happens from scratch (cf. the Conclusions). Instead, it is done in a rolling-horizon fashion, with only a small set of new project candidates to be taken account of at the time of a current planning decision, and the consideration of already scheduled projects does not cause a relevant computational burden. Thus, the algorithmic capacity of the CDPPA approach may be completely sufficient for coping with practical requirements.

Impacts on decision making

The essential claim of the CDPPA methodology concerns the gains it makes possible in decision making in terms of better outcomes achievable through a thorough formalisation of the planning problem. In the end, of course, a proof of this claim is aspired. Quite naturally, rationality imperatives suggest an optimisation framework to maximize the economic performance of explicit key criteria subject to constrained resources. It is well known, however, that optimisation is rarely captured straightforwardly in case of complex decision making environments, bearing in mind the effort it takes to meet all of the formal requirements of such an endeavour and conceding that the decision maker herself is neither omniscient nor in possession of unlimited calculating power [Klein, 2001]. In other words, the decision maker by necessity behaves *boundedly* rational; so, as a consequence, CDPPA has set out to provide effective decision support rather than deriving peculiarly optimal solutions on its own. Correspondingly, the question of optimisation superiority shifts towards the actual contributions of the CDPPA model to R&D project portfolio selection in practice. In this respect, decision support results from (i) highlighting the aspects of a planning problem critically influencing the eventual outcome, and (ii) relieving the decision maker from subordinate calculations such as feasibility checking of particular solution candidates – in either way effectively reducing the cognitive workload of the decision maker.

What became quickly apparent in the EC3 field trial was the *partial* replacement and enhancement of an intuitive approach towards project portfolio selection, that is, one based largely on perceived patterns among sets of projects indicating previously successful combinations, with now more *analytical* elements exploiting the various clues of a solution space provided by interactive modes of graphical exploration. In other words, the availability of a range of structural properties of a solution space could indeed be used to *sharpen* intuition, this way contributing tangibly to the rationality of the decision maker. However, at the same time, of course, the crucial question arises if and to which degree the exploratory guidance offered introduces its own unnoticed bias. At any rate, in relation to concepts like the aspiration adaptation theory [Selten, 1998] (which captures the traditional EC3 project portfolio selection "strategy" quite reasonably), the CDPPA decision analysis obviously facilitates a quick detection of general performance levels achievable by given project candidates (by inspecting the Pareto frontier) and, thus, simplifies the decision about whether or not (and into which direction, respectively) to continue spotting further project candidates. Moreover, intense examination of a solution space triggers *learning* effects as to which clues provide decisive information in certain circumstances: the (graphical and numerical) presentation of the solution space delivers a *language* to express detected differences among solution candidates that indeed contributes to a deeper reflected portfolio selection. Likewise, and no less important, is the arguable linkage of selected project portfolios to explicit performance criteria – both in terms of self-assurance of the decision maker, as well as in justifying choices of the management against the interests of EC3 owners and stakeholders. Clearly, the enhanced transparency of decision making also helps to negotiate internal criticism concerning project priorities, once the competency baskets and performance criteria have been broadly accepted. From an expert's point of view [Shanteau, 1992: 18], CDPPA could be judged as an apt support particularly in avoiding grossly inferior decisions while the choice amongst the set of Pareto-efficient solution candidates may still depend predominantly on experience, qualitative comparative judgement, and a (large) remainder of intuitive insight.

An interesting idea emerging from the field trial, the CDPPA approach suggests the splitting of the planning effort into a managerial role and a team role: while management defines the strategic competence development goals of the organisation expressed in terms of ppos and cpos, it is largely up to the R&D staff to contribute pertinent project candidates for consideration in the portfolio formation process. This may establish a very effective labour division scheme, reaping the benefits of team creativity and employee engagement, as well as directing the collective planning process, provided that access to the planning model is granted to all team members (or, in case of organisations larger than the EC3, team leaders) pari passu. Furthermore, a "dual" project definition mode is conceivable, seeking to augment partial project portfolios – composed, for instance, of projects the organisation has to implement at any rate for external reasons – with project candidates modelled after *ideal projects* in the sense of optimally complementing the partial portfolios with respect to the set decision criteria. While, apparently, ideal projects are fictitious ones most of the time, they highlight preferable competency mixtures and quantities as a screening, or design, pattern of realizable project candidates.

Finally, to state it once more, model fidelity is the Achilles heel of the CDPPA approach: the experienced advantages of the visually and interactively enhanced decision processes really come true only inasmuch the competency database, the project descriptions, and the specified decision criteria faithfully capture the actual planning conditions. Because of the massive data requirements of the model, the structure of the entire process of project portfolio selection changes markedly, substituting much of the cognitive burden of direct project assessment with increased efforts in project elicitation and quantification of resource demands of elicited project candidates. Even if a more participatory data procurement and planning approach is adopted, this is still no patent panacea to the imminent problem of the collaborative coordination of scarce planning resources: for example, while collaborative approaches to the elicitation of project opportunities may ensure the necessary breadth of candidate sets through labour division and collective review, this per se does not yet provide any solution to implement reasonable stopping rules for the search process or warrant thematically balanced project candidate sets – rather, the converse may hold.

Conclusions

The CDPPA project has tried hard to capture the essence of R&D planning as a multi-criteria decision support system that takes into account as many problem determinants as possible with realistic modelling effort. In an attempt to augment traditional project portfolio selection based on mainly qualitative decision criteria and intuitive comparisons of solution candidates with a formalised representation of the solution space akin to rational modes of optimisation, CDPPA seeks to challenge the expertise of R&D program managers through a computational tool that implies that, on the basis of calculation, better outcomes and, thus, a higher degree of rationality in decision making is achievable. Realistically, and acknowledging that intuitive decision making often performs surprisingly well, CDPPA aims at providing *rationalistic support* in a very complex decision making environment like R&D planning, in which reliance on experience and intuition never can be dispensed with entirely, simply because of the difficulty, or undue expense, encountered in formalising *all* of the salient problem determinants, many of which actually belong to meta- and meta-meta planning levels. So, eventually, the task confines to an "optimisation" of naturalistic decision making [Zsambok and Klein, 1997] based on the (visual) presentation of decision-critical structural properties of a high-dimensional solution space.

The CDPPA project resulted in a prototype implementation of an interactive multi-criteria optimisation model incorporating an elaborate representation of the cardinal productive R&D resource: human competencies. This representation is augmented with an accounting mechanism that takes into consideration competence dynamics resulting from the assignment of labour force to implemented R&D projects, selected on account of their contributions to specified decision criteria expressing the organisation's goals as to economic performance and, notably, strategic competence development. In principle, assuming a faithful representation of all decision-critical problem parameters and modelled competence data, the generated project portfolios are expected to be practically feasible.

Regardless of model fidelity considerations, it turns out very difficult to assess the real contributions of the CDPPA approach to the quality of the decision making practice. Basically, as highlighted in an extensive field trial, the decision making process itself is heavily affected, and there are some indications that these changes benefit decision quality by increasing particularly the transparency and accountability of the decisions made. What could not be validated so far, however, is any gain in performance as measured by the very decision criteria, for a variety of reasons:

- The CDPPA planning scenario abstracts from several contingencies influencing practical project selection: many projects are obligatory but resource assignments may be modified fairly flexible even at short notice, the interests of established project teams have to be allowed for, long-term thematic or research commitments restrict the decision maker's degrees of freedom, projects indicating signs of failure or under-performance may be stopped or re-configured, etc.

- To be effective, CDPPA planning must be fed with considerably larger sets of project candidates than have ordinarily been prepared, mainly because of the effort it takes to define candidate projects of varying selection probability with the accuracy required for CDPPA to process the data reasonably.
- The field trial started portfolio selection from scratch, a situation hardly reflecting the traditional planning mode of a more continuous update of the project schedule implying both, a strong path dependency of schedule adaptation and a fairly restricted project candidate set to decide upon each time. It turns out very intricate to phase a system like CDPPA into an ongoing decision making process, while still maintaining the lab condition one wishes for in order to accomplish fair performance comparisons.

As a natural consequence of these remarks, the practical contributions of the formalized CDPPA approach to real decision making must be judged rather cautiously: in addition to addressing apparent shortcomings of more traditional, elliptic portfolio selection procedures, it introduces its own intricacies that are hard to control or work around. At any rate, of course, in representing a sample case of naturalistic decision making in a fairly complex and dynamic environment, work on CDPPA contributes a lot to the meta-analysis of practical decision making and thus in itself improves the practice of decision making in the field. In particular, with respect to management style, two deviations from the initial CDPPA sketch have emerged, viz. (i) a focus shift from exclusive support for planning decisions of upper management to supporting collaborative decision making processes across multiple levels of the organizational hierarchy, and (ii) a view on strategic project planning as a rather permanent, partly experimental, multi-faceted process embedded in the regular workflows of the organisation.

So far, the CDPPA field trial did not cover a range of issues concerning the preconditions that need to be met for its practical introduction into the routine planning services of R&D organisations, notably larger research centres, academic institutions, or R&D planning and funding agencies, all of which in principle have to deal with R&D project portfolio formation in one way or another. These issues address both questions of interfaces to other management and planning subsystems, as well as the scalability needed to fit a variety of organisation schemes and modes. Despite the considerable modelling effort borne, CDPPA may have failed to cope with the more subtle cultural and social connotations of competence, with quite tangible but fairly hard to resolve consequences with regard to both decision quality and team-sociological effects.

This is to say that, while the CDPPA prototype system could be considered a first attempt to address a number of methodological issues in R&D project planning and human resource development by means of a novel approach, no conclusive evidence about the proposal's merits has been reached as yet. However, since the prototype is amenable to more extensive experimentation and evaluation, a discussion that might otherwise remain on an abstract level can be turned into a more precise and empirically underscored reasoning. No doubt, there is lot of further insight and evidence to be covered in the field.

Acknowledgment. Financial support from the Austrian Science Fund (FWF) through grant L264-N13 is gratefully acknowledged.

References

Afiouni F (2007) Human Resource Management and Knowledge Management: A Road Map Toward Improving Organizational Performance. *Journal of American Academy of Business* 11(2), 124-130.

Argyris C and Schon DA (1978) *Organizational Learning: A Theory of Action Perspective.* Addison-Wesley, Reading, MA.
Arrow KJ (1969) Classificatory Notes on the Production and Transmission of Technological Knowledge. *American Economic Review* 59(2), 29-35.
Blazewicz J, Lenstra JK and Rinnooy Kan AHG (1983) Scheduling Subject to Resource Constraints: Classification and Complexity. *Discrete Applied Mathematics* 5, 11-24.
Caves RE (2000) *Creative Industries – Contracts Between Art and Commerce.* Harvard University Press, Cambridge, MA, London.
Chen ANK and Edgington TM (2005) Assessing Value in Organizational Knowledge Creation: Considerations for Knowledge Workers. *MIS Quarterly* 29(2), 279-309.
Chernoff H (1973) The Use of Faces to Represent Points in k-dimensional Space Graphically. *Journal of the American Statistical Association* 68(342), 361-368.
Churchman CW, Ackoff RL and Arnoff EL (1959) *Introduction to Operations Research.* Wiley, New York, London, Sydney.
Cook D, Hofmann H, Lee EK, Yang H and Nokolau B (2007) Exploring Gene Expression Data, Using Plots. *Journal of Data Science* 5(2), 151-182.
David PA (1999) Path Dependence and Varieties of Learning in the Evolution of Technological Practice. In *Technological Innovation as an Evolutionary Process* (Ziman J, Ed), pp 118-133, Cambridge University Press, Cambridge (UK).
Delamare Le Deist F and Winterton J (2005) What Is Competence? *Human Resource Development International* 8(1), 27-46.
Dreyfus H and Dreyfus S (1986) *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer.* Free Press, New York.
Erpenbeck J and Heyse V (1999) Kompetenzbiographie – Kompetenzmilieu – Kompetenztransfer (in German). QUEM-Report 62, Berlin.
Foray D (2006) *The Economics of Knowledge.* MIT Press, Cambridge, MA, London.
Gigerenzer G (2001) The Adaptive Toolbox. In *Bounded Rationality – The Adaptive Toolbox* (Gigerenzer G and Selten R, Eds), pp 37-50, MIT Press, Cambridge, MA, London.
Gilbreth FB (1912) *Primer of Scientific Management.* Van Nostrand, New York.
Gutjahr WJ, Katzensteiner S and Reiter P (2007) A VNS Algorithm for Noisy Problems and Its Application to Project Portfolio Analysis. In *Proc. SAGA 2007 – Stochastic Algorithms: Foundations and Applications* (Hromkovic J et al., Eds), pp 93-104, Springer Lecture Notes in Computer Science 4665.
Gutjahr WJ, Katzensteiner S, Reiter P, Stummer Ch and Denk M (2008) Competence-driven Project Portfolio Selection, Scheduling and Staff Assignment. *Central European Journal of Operations Research* 16(3), 281-306.
Gutjahr WJ, Katzensteiner S, Reiter P, Stummer Ch and Denk M (2010) Multi-objective Decision Analysis for Competence-oriented Project Portfolio Selection. *European Journal of Operational Research* 205(3), 670-679.
HR-XML Consortium (2007) Competencies (Measurable Characteristics) Recommendation 2007-04-15, [WWW document] http://ns.hr-xml.org/2_5/HR-XML-2_5/CPO/Competencies.html (accessed 12 March 2010).
James CJ (2003) Designing Learning Organizations. *Organizational Dynamics* 32(1), 46-61.
Jordan GB (2005) What Matters to R&D Workers. *Research Technology Management* 48(3), 23-32.
Kasanen E, Östermark R and Zeleny M (1991) Gestalt System of Holistic Graphics: New Management Support View of MCDM. *Computers & Operations Research* 18(2), 233-239.
Klein G (2001) The Fiction of Optimization. In *Bounded Rationality – The Adaptive Toolbox* (Gigerenzer G and Selten R, Eds), pp 103-121, MIT Press, Cambridge, MA, London.

Kleiner B and Hartigan JA (1981) Representing Points in Many Dimensions by Trees and Castles. *Journal of the American Statistical Association* 76(3/4), 260-269.
Korhonen P (1991) Using Harmonious Houses for Visual Pairwise Comparisons of Multiple Criteria Alternatives. *Decision Support Systems* 7(1), 47-54.
Lotfi V, Stewart TJ and Zionts S (1992) An Aspiration-level Interactive Model for Multiple Criteria Decision Making. *Computers & Operations Research* 19(7), 671-681.
Mansfield RS (1996) Building Competency Models: Approaches for HR Professionals. *Human Resource Management* 35(1), 7-18.
Ngwenyama O, Guergachi A and McLaren T (2007) Using the Learning Curve to Maximize IT Productivity: A Decision Analysis Model for Timing Software Upgrades. *International Journal of Production Economics* 105, 524-535.
Quintas, P, Lefrere P and Jones G (1997) Knowledge Management: A Strategic Agenda. *Long Range Planning* 30(3), 385-391.
Riedmann H (2008a) CDPPA: Contextualization of the Optimization Tool and Analysis of the Data Collection Process, Technical Report.
Riedmann H (2008b) CDPPA: Meta-Analysis of the First Test Phase Without Regard to Risk and Uncertainty, Technical Report.
Selten R (1998) Aspiration Adaptation Theory. *J. Math. Psychology* 42, 191-214.
Selten R (2001) What Is Bounded Rationality? In *Bounded Rationality – The Adaptive Toolbox* (Gigerenzer G and Selten R, Eds), pp 13-36, MIT Press, Cambridge, MA, London.
Shanteau J (1992) The Psychology of Experts: An Alternative View. In *Expertise and Decision Support* (Wright G and Bolger F, Eds), pp 11-23, Plenum Press, New York.
Shippmann JS, Ash RA, Battista M, Carr L, Eyde LD, Hesketh B, Kehoe J, Pearlman K and Sanchez JI (2000) The Practice of Competency Modeling. *Personnel Psychology* 53(3), 703-740.
Snell SA and Dean JW (1992) Integrated Manufacturing and Human Resource Management: A Human Capital Perspective. *Academy of Management Journal* 35(3), 467-504.
Snell SA, Shadur MA and Wright MA (2001) Human Resources Strategy: The Era of Our Ways. In *The Blackwell Handbook of Strategic Management* (Hitt MA, Freeman E and Harrison JS, Eds), pp 627-629, Blackwell Publishers, Oxford (UK).
Spencer LM and Spencer SM (1993) *Competence at Work*. Wiley, New York.
Stummer Ch and Heidenberger K (2003) Interactive R&D Portfolio Analysis with Project Interdependencies and Time Profiles of Multiple Objectives. *IEEE Transactions on Engineering Management* 50(2), 175-183.
Stummer Ch, Kiesling E and Gutjahr WJ (2009) A Multicriteria Decision Support System for Competence-driven Project Portfolio Selection. *International Journal of Information Technology and Decision Making* 8(2), 379-401.
Sundberg L (2001) A Holistic Approach to Competence Development. *Systems Research* 18, 103-114.
Sveiby KE (1997) *The New Organizational Wealth. Managing and Measuring Knowledge-Based Assets*. Berrett-Koehler, San Francisco.
Wernerfelt B (1984) A Resource-based View of the Firm. *Strategic Management Journal* 5, 171-180.
Wilson TD (2002) The Nonsense of 'Knowledge Management'. *Information Research* 8(1), paper no. 144.
Zsambok CE and Klein G, Eds (1997) *Naturalistic Decision Making.* Lawrence Erlbaum Ass, Mahwah NJ.