Stochastic project management: multiple projects with multi-skilled human resources

Thomas Felberbauer¹ · Walter J. Gutjahr² · Karl F. Doerner³

Published online: 17 December 2018
© The Author(s) 2018

Abstract

This paper presents two stochastic optimization approaches for simultaneous project scheduling and personnel planning, extending a deterministic model previously developed by Heimerl and Kolisch. For the problem of assigning work packages to multi-skilled human resources with heterogeneous skills, the uncertainty on work package processing times is addressed. In the case where the required capacity exceeds the available capacity of internal resources, external human resources are used. The objective is to minimize the expected external costs. The first solution approach is a “matheuristic” based on a decomposition of the problem into a project scheduling subproblem and a staffing subproblem. An iterated local search procedure determines the project schedules, while the staffing subproblem is solved by means of the Frank–Wolfe algorithm for convex optimization. The second solution approach is sample average approximation where, based on sampled scenarios, the deterministic equivalent problem is solved through mixed-integer programming. Experimental results for synthetically generated test instances inspired by a real-world situation are provided, and some managerial insights are derived.

Keywords Stochastic optimization · Project scheduling · Personnel planning · Heterogeneous skills · Iterated local search

1 Introduction

In project management, the two important tasks of project scheduling on the one hand and of personnel planning on the other hand are usually not faced separately from each other, but rather in a simultaneous or interleaved manner. This interaction considerably increases the complexity of the planning process, with the consequence that project managers frequently wish to get computational support alleviating the cognitive burden of the combined scheduling/staffing decision. Decision support tools addressing the combined

problem have been proposed in the literature, but typically they use deterministic models. Despite the fact that managers and scientists recognize the importance of uncertainty within project management, few works take this aspect into account when proposing methods for quantitative decision support for the combined problem described above, obviously for the reason that, as argued, this problem is already hard in the deterministic case. In this study we highlight the potential of stochastic optimization applied to combined project scheduling and staffing, and compared two different methods of solving the proposed model computationally.

The project scheduling and staffing model proposed by Heimerl and Kolisch (2010) deals with the problem of assigning multi-skilled human resources to work, while taking into account resource-specific and heterogeneous skill efficiencies. The objective is to minimize the costs for internal and external personnel. Heimerl and Kolisch (2010) provide a mixed-integer programming (MIP) formulation with a tight LP bound and compare its performance, using the MIP solver of CPLEX, to that of simple heuristics. Additionally, they investigate the influence of different parameters, such as the time window size, the utilization, and the number of skills per resource, on the personnel costs. In Kolisch and Heimerl (2012) the authors extend their previously presented model

✉ Thomas Felberbauer
thomas.felberbauer@fhstp.ac.at

Walter J. Gutjahr
walter.gutjahr@univie.ac.at

Karl F. Doerner
karl.doerner@univie.ac.at

¹ Department of Media and Digital Technologies, St. Pölten University of Applied Sciences, St. Pölten, Austria

² Department of Statistics and Operations Research, University of Vienna, Vienna, Austria

³ Department of Business Decisions and Analytics, Data Science, University of Vienna, Vienna, Austria

and develop a hybrid metaheuristic as an innovative solution method. The proposed solution method separates the problem into a scheduling and a staffing subproblem, where the staffing problem is solved through a generalized network simplex algorithm. Results show that the hybrid metaheuristic outperforms the MIP solver. In Felberbauer et al. (2016), the models from Heimerl and Kolisch (2010) and Kolisch and Heimerl (2012) are extended by considering labor contracts. For the problem of project scheduling and staffing with labor contracts, Felberbauer et al. (2016) propose a hybrid metaheuristic combining iterated local search and a greedy staffing heuristic.

Van den Bergh et al. (2013) observe that most papers on personnel scheduling problems still appear to feature a deterministic approach and advise researchers to consider stochastic problem versions. They highlight the importance of uncertainty, explicitly mentioning volatile demand, last-minute changes, and rescheduling based on new information, as interesting new research topics. A recent article on workforce planning by De Bruecker et al. (2015) states that the small number of papers in this field taking uncertainty into account is alarming. They refer to the general consensus that uncertainty is ubiquitous in real workforce planning, but uncertainty still remains relegated to the “future research” section of many papers. The authors mention the high complexity of integrating uncertainty within optimization models and techniques as one reason for its inadequate representation in the literature.

Considering uncertainty in project management and its subordinate planning levels, a recent work of Barz and Kolisch (2014) describes resource assignments in the telecommunication industry. The authors investigate a hierarchical, multi-skill resource assignment problem and use a discrete Markov decision process model for incoming jobs over an infinite time horizon. Gutjahr and Froeschl (2013) present a stochastic optimization model for project selection and project staffing. They assume that both the returns and the required efforts of the selected projects are random variables. The problem is decomposed into a project selection problem and a staffing subproblem. For the computational solution of the two problems, an adapted version of variable neighborhood search and a Frank–Wolfe-type algorithm, respectively, are used. Artigues et al. (2013) propose a robust optimization approach to resource-constrained project scheduling with uncertain activity duration, assuming that the decision-maker cannot associate probabilities with possible activity durations. The authors describe how robust optimization can be applied to project scheduling under uncertainty, and they develop a scenario-relaxation algorithm and a heuristic solution method that solves medium-sized instances. Gutjahr (2015) proposes a model for stochastic multi-mode resource-constrained project scheduling under risk aversion with the two objectives of makespan and costs. Activity durations and

costs are modeled as random variables. For the scheduling part of the decision problem, the class of early-start policies is considered. A further decision to be made concerns the assignment of execution modes to activities. To take risk aversion into account, an approach of optimization under multivariate stochastic dominance constraints is adopted. For the resulting biobjective stochastic integer programming problem, the Pareto frontier is determined by means of an exact solution method, incorporating a branch-and-bound technique.

Ingels and Maenhout (2017a,b) investigate the assignment of employees to cover the staffing requirements for specific skills and shifts. Their research focuses on improving the short-term adjustment capability of the shift roster by maximizing the substitutability of employees as a means to increase the robustness of a project plan. The authors propose a three-step methodology including a two-phase preemptive programming approach. Both uncertainty of demand and uncertainty of capacity are considered by stochastic models. A personnel shift roster for a medium-term period offering sufficient flexibility for between-skill substitution and within-skill substitution is provided.

The articles Heimerl and Kolisch (2010), Kolisch and Heimerl (2012), Felberbauer et al. (2016), and Gutjahr and Froeschl (2013) constitute the starting point for the present paper. In line with these articles, we assume that in the case where the work time demand exceeds the available capacity, external capacity is used, e.g., by hiring external personnel. However, we introduce the following new features: (a) Contrary to the deterministic models presented in Heimerl and Kolisch (2010), Kolisch and Heimerl (2012), Felberbauer et al. (2016), we model the required work time demand as *stochastic*. On the other hand, in order to keep our model compact, we simplify the model of Heimerl and Kolisch (2010), Kolisch and Heimerl (2012) by neglecting the possibility of overtime work. (b) Contrary to Gutjahr and Froeschl (2013), we address the *project scheduling* decision. This issue introduces a relevant additional source of computational complexity into the problem, since project scheduling problems (already in a simple, deterministic context) are notoriously hard. In particular, the demand information does not appear here per project, but in a more fine-grained way for each activity of a project. On the other hand, compared to Gutjahr and Froeschl (2013), we do not deal here with the project selection aspect.

To cope with the computational challenge of the proposed model, we develop two solution approaches and compare their performance. The first, “matheuristic” approach uses a metaheuristic for the scheduling part of the problem as well as an exact solution procedure inspired by Gutjahr and Froeschl (2013) for the staffing part. It is noteworthy that the staffing part constitutes a nonlinear (though convex) optimization

problem. The second approach follows a completely different strategy by using sample average approximation to achieve a linearization of the model, and by solving the resulting “deterministic equivalent” problem, which remains a combined scheduling-staffing problem, with the aid of a MIP solver. A comparison of these two ways to address a mixed-integer stochastic optimization problem with a convex “lower level” problem may also be interesting for applications outside the area of project management.

Substantial literature focuses on the deterministic case of project scheduling and staffing. However, to the best of our knowledge, no previous studies have addressed the scheduling and staffing of multiple projects with multi-skilled resources holding heterogeneous skill efficiencies, under consideration of uncertain work time demand. The contribution of our article is:

1. We adapt the model of Kolisch and Heimerl (2012) by considering stochastic demand.
2. We present two solution procedures, namely Sample Average Approximation and a matheuristic, and compare their performance.
3. We analyze deterministic and stochastic planning and discuss cost estimation accuracy and the value of the stochastic solution.
4. We formulate managerial insights from the stochastic planning approach.

The paper is organized as follows: In Sect. 2 the stochastic optimization model is introduced. Section 3 presents the developed solution procedures. The structure of our test instances and the experimental design are described and the results of the computational experiments are discussed in Sect. 4. Finally, in Sect. 5 we formulate some managerial insights that can be drawn from the study of the stochastic optimization problem, and we mention some interesting topics for future research.

2 Problem formulation

In this section the stochastic project scheduling and staffing problem is presented.

Projects P is a set of *projects* that must be completed during the planning horizon T (e.g., one year). All projects are independent of each other but compete for the same resources. The planning horizon T comprises discrete time periods t (e.g., months). Each project p ($p \in P$) has a work schedule defined by a sequence of consecutive *activities* $q = 1, \dots, d_p$, where an activity is defined as that part of a project that is executed during one time period t . In this way, we discretize the overall project into consecutive pieces, each of the same length; as it will be discussed below, interruptions

will be allowed. During activity q , project p demands D_{psq} units of work requiring *skill* s ($s \in S$), where S is a set of skills. The amount of work in a skill s performed in a certain activity q of a certain project p is called a *work package*. The sizes D_{psq} of the work packages, which we can also interpret as work contents or efforts, are considered as random variables. This reflects the very frequently occurring situation that the required efforts D_{psq} are unknown at the beginning of the planning process. Their true realizations become only known during project execution; we assume, however, that the distribution of the variables D_{psq} is known in advance (or can at least be estimated): it is described by the probability density function (PDF) h_{psq} ($p \in P, s \in S, 1 \leq q \leq d_p$). By using a PDF, we suppose that the distribution of D_{psq} is a continuous distribution, such as a triangular or a beta distribution.

Moreover, we assume that a *time window* $[ES_p, LS_p]$ for the start of each project $p \in P$ is pre-defined. Therein, ES_p and LS_p denote the index t of the earliest time period and of the latest time period, respectively, in which project p can be started. The *latest finish period* of project p , denoted LF_p , is defined as $LF_p = LS_p + d_p - 1$. The *time window size* γ of project p is defined as $\gamma = LS_p - ES_p$. The earliest start time of activity q is denoted ES_{pq} and defined as $ES_{pq} = ES_p + q - 1$. The latest start time of activity q of project p , denoted LS_{pq} , is equal to the latest finish time LF_{pq} of this activity and is defined as $LS_{pq} = LF_{pq} = LS_p + q - 1$. The activities $q = 1, \dots, d_p$ have to be processed in an ascending order. However, it is possible to *interrupt* the project between the discrete activities q once the project has started. Nevertheless, the order of the activities as well as the start time and finish time constraints have to be respected. For example, if $ES_p = 1$, $LS_p = 3$ and $d_p = 3$ for a given project p , then it would be allowed to schedule the three activities of project p in time periods 1, 2, 3 (earliest possible schedule); some other feasible alternatives would be to schedule them in time periods 1, 4, 5, in time periods 2, 3, 5, or in time periods 3, 4, 5 (latest possible schedule). The schedule 2, 4, 6 would not be allowed, because it exceeds the latest finish period $LF_p = 5$.

Resources Our model distinguishes between *internal* and *external* human resources. Each internal resource k taken from the set of all internal human resources K holds a subset of skills $S_k \subseteq S$. Conversely, from the perspective of a skill s , the subset $K_s = \{k \in K \mid s \in S_k\}$ of K contains all resources that can perform skill s . If $s \in S_k$, we take account of different degrees in the performance of resource k in skill s , specified by the *efficiency* η_{sk} of resource k in skill s . The higher the efficiency value η_{sk} , the faster is resource k in executing a work package D_{psq} requiring skill s . It is assumed that η_{sk} is known at the beginning of the time horizon (the time of the decision) and does not change until its end. (In particular, we disregard effects as learning or knowledge depreciation).

The values η_{sk} are strictly positive for all $s \in S_k$. Whenever convenient, we extend the definition of the values η_{sk} to all pairs (s, k) and set then $\eta_{sk} = 0$ for $s \notin S_k$. The quantity η_{sk} represents a factor for the speed in which a certain skill is exerted. To distinguish between work content and actual time needed to perform a work package, the realization d_{psq} of the random variable D_{psq} will also be called *effective work time*. To get the time actually needed by the considered resource, we have to divide the effective work time by η_{sk} , so the *real work time* is d_{psq}/η_{sk} . For example, let us assume that the work package d_{psq} requires 10 units of effective work time to be processed. (This would be the time a “standard” employee would need for it.) Then, if employee k has an efficiency of $\eta_{sk} = 1.25$ in skill s , she or he needs only $d_{psq}/\eta_{sk} = 10/1.25 = 8$ time units of *real work time* to complete this work package.

Capacities Each (internal) resource k has a limited capacity of a_{kt} during time period t ; these values are given in advance. Similar to Gutjahr and Froeschl (2013), we assume that internal resources earn fixed wages per period. On our assumption that there is no overtime work, the costs for these wages sum up to a constant value, so they have no influence on the optimization problem. (Note that for the assumed employment type, even if an employee has idle times, the company has to pay for the hours where she or he is available and not only for the hours of scheduled work.) On the other hand, *external* resources are paid for the effective work they do. Our model assumes that they are available for each skill s to an unlimited extent at the expense of a cost rate c_s^e per unit of effective work time.

Since in our model, the work time demand D_{psq} is stochastic, but the capacities of the internal resources are fixed and the project schedule has to be decided upon in advance as well, it is necessary to perform *recourse actions* after having observed the actual realizations of the D_{psq} in order to be able to stick to the chosen schedule. We assume that whenever the internal capacity turns out as insufficient to cover the demand, the company resorts to external resources (experts in the required skills) and pays them at their given cost rates. The (expected) overall cost for the external resources has to be minimized.

Decision variables For the stochastic optimization model, the decision variables $x_{ptsk} \geq 0$ define the amount of work performed by internal resource k with skill s in period t for project p . Note that x_{ptsk} is measured in effective work time and not in real work time. To define the times when the periods of a project are executed, the binary decision variables $z_{pqt} \in \{0, 1\}$ are introduced: $z_{pqt} = 1$ if in time period t , activity q of project p is executed, and 0 otherwise. Thus, the x_{ptsk} variables define the staffing decision, whereas the z_{pqt} variables define the scheduling decision.

2.1 Stochastic optimization model

In (2)–(7), we present our stochastic optimization model in mathematical terms. For a project p and a time period t , we let

$$\tau_{pt} = \{q \in \{1, \dots, d_p\} \mid ES_{pq} \leq t \leq LS_{pq}\} \quad (1)$$

denote the set of all possible activities q of project p that could lead to a resource demand in time period t .

$$\min_{z, x} \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} c_s^e E \left(\left[\sum_{q \in \tau_{pt}} D_{psq} z_{pqt} - \sum_{k \in K_s} x_{ptsk} \right]^+ \right) \quad (2)$$

subject to

$$\sum_{t=ES_{pq}}^{LS_{pq}} z_{pqt} = 1 \quad \begin{array}{l} p \in P, \\ q = 1, \dots, d_p \end{array} \quad (3)$$

$$\sum_{t=ES_{pq}}^{LS_{pq}} t z_{pqt} < \sum_{t=ES_{p,q+1}}^{LS_{p,q+1}} t z_{p,q+1,t} \quad \begin{array}{l} p \in P, \\ q = 1, \dots, d_p - 1 \end{array} \quad (4)$$

$$\sum_{p \in P} \sum_{s \in S} \frac{1}{\eta_{sk}} x_{ptsk} \leq a_{kt} \quad \begin{array}{l} k \in K, \\ t = 1, \dots, T \end{array} \quad (5)$$

$$x_{ptsk} \geq 0 \quad \begin{array}{l} p \in P, \\ t = ES_p, \dots, LF_p, \\ s \in S, \\ k \in K_s \end{array} \quad (6)$$

$$z_{pqt} \in \{0, 1\} \quad \begin{array}{l} p \in P, \\ q = 1, \dots, d_p, \\ t = ES_{pq}, \dots, LS_{pq} \end{array} \quad (7)$$

The *objective function* in Eq. (2) minimizes the expected external costs. The symbol E denotes the mathematical expectation and x^+ stands for $\max(x, 0)$. The external costs result from that part of the stochastic demand that is not covered by the internally scheduled capacity. The constraints in Eq. (3) ensure that activity q of project p is scheduled exactly once within the pre-defined time window $[ES_{pq}, LS_{pq}]$. The condition that activities q must be processed in an ascending order is guaranteed by the constraints in Eq. (4). Equation (5) formulates the capacity constraints for each resource k and time period t : observe that for a resource with efficiency η_{sk} , an effective work time of x_{ptsk} entails a real work time of x_{ptsk}/η_{sk} . Finally, the decision variables are defined in Eqs. (6) and (7).

Let us remark that in some cases, external costs grow faster than linearly in dependence of the outsourced work, for example because a selected supplier company has limited capacities itself. To generalize the model above to this situation, the given nonlinear external cost function could be approximated by a piecewise linear function, which can be dealt with by a simple extension of our Eq. (2). Numerical

solution techniques for this more complex model are a topic of future research.

As it is seen from the formulation above, our model can be viewed as a tactical two-stage stochastic optimization model with project plan and work assignment as first-stage decisions and outsourcing to external resources as the possible recourse actions. During the ‘real’ project execution, we assume that we get updated information about the realizations of the work package processing times. Our assumed recourse in the case of insufficient internal capacity is the hiring of external resources, where we assume that the information about the work package processing times arrives early enough to hire the necessary external resources. The presented tactical planning approach could be extended including more operative decisions by concerning other recourse actions such as reassignment and rescheduling of internal resources. This could lead to better solutions and new interesting managerial insights but would lead to a more sophisticated planning approach.

Finally, let us recall that our model assumes that the activities q of each project p have to be arranged in a pre-specified linear order analogously as in Kolisch and Heimerl (2012). Nevertheless, due to the requirement of independent work package processing time distributions, an additional assumption, where for each project p and in each time period t not more than one activity q of project p is allowed to be scheduled (no parallel activity processing per project), is needed. If this condition is satisfied, an extension of the developed stochastic solution approach to more general precedence constraints between activities is straightforward. For this extension, the variable z_{pqt} can again be the indicator variable for the decision that activity q of project p is scheduled in time period t . A constraint $\sum_q z_{pqt} \leq 1$ has to be added, and the precedence constraints have to be expressed by a more flexible set of constraints (Artigues et al. 2008) than those given by Eq. (4).

3 Solution procedures

3.1 Problem structure

As mentioned above, the considered project scheduling and staffing problem distinguishes between two decisions: the project scheduling decision where the execution times of the activities q are determined, and the staffing decision where employees are assigned to cover parts of the work packages. Let z and x denote the array of the decision variables z_{pqt} for project scheduling, and the array of the decision variables x_{ptsk} for staffing, respectively. By ξ , we denote the random influence in our stochastic model. Abbreviating the total external cost by $G(z, x, \xi)$ (it depends on ξ since it is a random variable) and its expected value by

$g(z, x) = E[G(z, x, \xi)]$, our optimization problem can be written as $\min_{z, x} g(z, x)$, where (z, x) has to satisfy the constraints (3)–(7). For a given project schedule z , we call $\min_x g(z, x)$ on the constraint that x is feasible with respect to (5)–(6), the *staffing subproblem*.

3.2 Expected value problem

The expected value (EV) problem or mean value problem is obtained from the original stochastic problem by replacing each random variable D_{psq} by its expected value $\bar{d}_{psq} = E(D_{psq})$, so that the distribution of D_{psq} collapses to the point mass in $\bar{d}_{psq}$. With the help of the introduction of the auxiliary variables for the external work time required y_{pts} per project p time period t and skill s , the following mixed-integer linear problem is obtained:

$$\min_{z, x} \Theta_{EV}(z, x) = \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} \left(c_s^e y_{pts} \right) \quad (8)$$

subject to constraints (3), (4), (7), and

$$\sum_{q \in \tau_{pt}} \bar{d}_{psq} z_{pqt} \leq y_{pts} + \sum_{k \in K_s} x_{ptsk} \quad \begin{array}{l} p \in P, \\ t = ES_p, \dots, LF_p, \\ s \in S \end{array} \quad (9)$$

$$y_{pts} \geq 0 \quad \begin{array}{l} p \in P, \\ t = ES_p, \dots, LF_p, \\ s \in S \end{array} \quad (10)$$

The EV problem serves as an approximation for cases where the variances of the random variable D_{psq} are very low, or, if the variances are larger, as a means to obtain an initial solution for an iterative search procedure

3.3 Staffing subproblem

For a pre-defined feasible project schedule z , the staffing problem $\min_x g(z, x)$ remains to be solved. Because z is now already fixed,

$$D'_{pst} = \sum_{q \in \tau_{pt}} D_{psq} z_{pqt} \quad (11)$$

is a random variable that does not depend on any decision anymore, i.e., the distribution of D'_{pst} is already known. D'_{pst} represents the effective work time of project p in skill s that has been scheduled for time period t by the given schedule z . Using the random variables D'_{pst} , we can express the staffing problem in the form

$$\min_x \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} c_s^e E \left(\left[D'_{pst} - \sum_{k \in K_s} x_{ptsk} \right]^+ \right) \quad (12)$$

subject to constraints (5) and (6).

Obviously, the objective function (12) is nonlinear, but it is not difficult to see that it is at least convex, since the function $x \mapsto x^+$ is a convex function and the expectation operator, as a linear operator, preserves convexity. Let

$$\varphi_{pst}(\zeta) = E([D'_{pst} - \zeta]^+) = \int_{\zeta}^{\infty} (\theta - \zeta) h_{pst}(\theta) d\theta \quad (13)$$

with $h_{pst}(\theta)$ denoting the probability density function of D'_{pst} , such that (12) can be rewritten as

$$\min_x \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} c_s^e \varphi_{pst} \left(\sum_{k \in K_s} x_{ptsk} \right) \quad (14)$$

subject to (5) and (6). Elementary calculations show that

$$\varphi'_{pst}(\zeta) = \int_{-\infty}^{\zeta} h_{pst}(\theta) d\theta - 1 = H_{pst}(\zeta) - 1, \quad (15)$$

with H_{pst} denoting the cumulative distribution function (CDF) of D'_{pst} .

3.4 Matheuristic

Our first solution method is a *matheuristic* approach, i.e., a combination of a metaheuristic with an exact optimization technique. We solve the staffing subproblem by means of the exact *Frank–Wolfe* algorithm (see, e.g., Clarkson 2010) for convex optimization under linear constraints. This is the topic of Sect. 3.4.1. The scheduling problem is solved by a metaheuristic of Iterated Local Search type, which will be described in Sect. 3.4.2.

3.4.1 Staffing: Frank–Wolfe algorithm

The idea of using the Frank–Wolfe algorithm for staffing problems has already been elaborated in Gutjahr and Froeschl (2013). Our technique used in the present work follows the approach described there rather closely, so we shall keep the presentation short. First, let us introduce the variables $u_{ptsk} = x_{ptsk} / \eta_{sk} a_{kt}$ ($p \in P$, $t = 1, \dots, T$, $s \in S$, $k \in K$). Using this substitution for a fixed project schedule z , the objective function (12) can be rewritten as

$$\min \Theta_z(u) = \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} c_s^e E \left(\left[D'_{pst} - \sum_{k \in K_s} \eta_{sk} a_{kt} u_{ptsk} \right]^+ \right). \quad (16)$$

Moreover, we combine the project p and the skill s to the pair $\sigma = (p, s)$, which we call *project-skill combination*, and we combine the time period t and the resource k to the pair $v = (t, k)$, the *time-employee combination*. Additionally, we relabel the pairs $\sigma = (p, s)$ and the pairs $v = (t, k)$ by introducing the new indices $\sigma = 1, \dots, C = |P| \cdot |S|$ and $v = 1, \dots, L = |T| \cdot K$. The variables u_{ptsk} are relabeled accordingly, i.e., for $\sigma = (p, s)$ and $\mu = (t, k)$, the notation $u_{\sigma v}$ abbreviates u_{ptsk} . In the constraints of Eq. (5), the inequalities can be replaced by equalities, since there is an optimal solution in which each constraint (5) is active. Therefore, (5) can be replaced by $\sum_{\sigma=1}^C u_{\sigma v} = 1$ ($v = 1, \dots, L$). The column vector $u_v = (u_{1v}, \dots, u_{Cv})'$ is an element of the standard simplex in $\mathbb{R}^C$, hence the feasible set is a Cartesian product of L standard simplices.

The idea of the iterative Frank–Wolfe algorithm is to replace in each iteration i the convex function Θ_z by its linear approximation at a current feasible solution $u^{[i]}$. The approximating linear function has a minimizer $g^{[i]}$ on the feasible set which can be easily determined. Next, the minimum of the convex function Θ_z restricted to the line segment between $u^{[i]}$ and $g^{[i]}$ is identified. This can be done by line search. The minimizer found in this way is used as the new current solution for the next iteration.

For our application, the linear approximation in point u produces the optimization problem

$$\min \left\{ \sum_{\sigma=1}^C \sum_{v=1}^L \frac{\partial \Theta_z(u)}{\partial u_{\sigma v}} r_{\sigma v} \mid r = (r_{\sigma v}) \in S_C^L \right\}. \quad (17)$$

This problem decomposes into L partial problems

$$\min \left\{ \sum_{\sigma=1}^C \frac{\partial \Theta_z(u)}{\partial u_{\sigma v}} r_{\sigma v} \mid r_v = (r_{1v}, \dots, r_{Cv})' \in S_C \right\}, \quad v = 1, \dots, L. \quad (18)$$

The solution of the v th problem in (18) is an extremal point of the simplex S_C . Therefore, the point $g^{[i]}$ is of the form $(e_{\sigma^*(1)}, \dots, e_{\sigma^*(L)})$, where e_{σ} is the σ th unit vector, and $\sigma^*(v)$ is the index of the optimal extremal point for time-employee combination v . It is easily seen that $\sigma^*(v)$ is given by the index σ of the smallest value among the partial derivatives in (18). One finds

$$\frac{\partial \Theta_z(u)}{\partial u_{ptsk}} = c_s^e \eta_{sk} a_{kt} \varphi'_{pst} \left(\sum_{k' \in K_s} \eta_{sk'} a_{k't} u_{ptsk'} \right)$$

which, by (15), leads to the maximization problem

$$\max_{p,s} c_s^e \eta_{sk} a_{kt} \left(1 - H_{pst} \left(\sum_{k' \in K_s} \eta_{sk'} a_{k't} u_{ptsk'} \right) \right). \quad (19)$$

Solving (19) by enumeration, we find for each time–employee combination $v = (t, k)$ the project–skill combination $\sigma^*(v) = (p^*(v), s^*(v))$ providing the highest cost reduction potential achievable by employee k in time period t .

Algorithm 1 presents the basic version of the Frank–Wolfe algorithm applied to the stochastic staffing subproblem. We also slightly modify this basic procedure by doing the line search for each column v separately, followed by an immediate change of column v of the current matrix $u^{[i]}$. This gives Algorithm 2.

3.4.2 Project scheduling: matheuristic solution method

After the description of the solution of the staffing subproblem in the previous section, the current section presents the search procedure that optimizes the project schedule z . The implemented metaheuristic (which is also used for solving a deterministic version of our project scheduling problem in Felberbauer et al. (2016)) is a modification of iterated local search (ILS) (e.g., Lourenço et al. 2010) using variable neighborhood descent (VND) as the local search component (e.g., Mladenović and Hansen 1997). The metaheuristic is com-

Algorithm 1 Basic Frank-Wolfe algorithm for the staffing problem

```

1:  $i \leftarrow 1$ 
2:  $u^{[i]} \leftarrow \text{InitialSolution}()$  ▷ //Compute a feasible initial solution
3: repeat
4:    $i \leftarrow i + 1$ 
5:   for all  $v = (t, k)$  do
6:     for all  $\sigma = (p, s)$  with  $\sum_q z_{pqt} > 0$  and  $s \in S_k$  do
7:        $\chi(p, s) \leftarrow c_s^e a_{kt} \eta_{sk} [1 - H_{pst}(\sum_{k' \in K_s} \eta_{sk'} a_{k't} u_{ptsk'}^{[i]})]$ 
8:     end for
9:      $\sigma^*(v) = (p^*(v), s^*(v)) \leftarrow \arg \max_{p,s} \chi(p, s)$  ▷ //Optimal index combination
10:     $g_v^{[i]} \leftarrow e_{\sigma^*(v)}$  ▷ //Set the target column vector
11:  end for
12:   $g^{[i]} \leftarrow (g_1^{[i]}, \dots, g_L^{[i]})$ 
13:   $\vartheta^* \leftarrow \arg \min_{0 \leq \vartheta \leq 1} \Theta_z((1 - \vartheta)u^{[i]} + \vartheta g^{[i]})$  ▷ //Apply line search
14:   $u^{[i+1]} \leftarrow (1 - \vartheta^*)u^{[i]} + \vartheta^* g^{[i]}$  ▷ //New  $u^{[i]}$  according to  $\vartheta^*$ 
15: until  $i = i_{\max}$ 

```

Algorithm 2 Modified Frank-Wolfe algorithm for the staffing problem

```

1:  $i \leftarrow 1$ 
2:  $u^{[i]} \leftarrow \text{InitialSolution}()$  ▷ //Compute a feasible initial solution
3: repeat
4:    $i \leftarrow i + 1$ 
5:   for all  $v = (t, k)$  do
6:     for all  $\sigma = (p, s)$  with  $\sum_q z_{pqt} > 0$  and  $s \in S_k$  do
7:        $\chi(p, s) \leftarrow c_s^e a_{kt} \eta_{sk} [1 - H_{pst}(\sum_{k' \in K_s} \eta_{sk'} a_{k't} u_{ptsk'}^{[i]})]$ 
8:     end for
9:      $\sigma^*(v) = (p^*(v), s^*(v)) \leftarrow \arg \max_{p,s} \chi(p, s)$  ▷ //Optimal index combination
10:     $g_v^{[i]} \leftarrow e_{\sigma^*(v)}$  ▷ //Set the target column vector
11:    obtain  $g^{[i]}$  from  $u^{[i]}$  by replacing column  $u_v^{[i]}$  with  $g_v^{[i]}$ 
12:     $\vartheta^* \leftarrow \arg \min_{0 \leq \vartheta \leq 1} \Theta_z((1 - \vartheta)u^{[i]} + \vartheta g^{[i]})$  ▷ //Apply line search
13:     $u^{[i]} \leftarrow (1 - \vartheta^*)u^{[i]} + \vartheta^* g^{[i]}$  ▷ //New  $u^{[i]}$  according to  $\vartheta^*$ 
14:  end for
15:   $u^{[i+1]} \leftarrow u^{[i]}$ 
16: until  $i = i_{\max}$ 

```

Fig. 1 Matheuristic: framework

combined with the exact Frank–Wolfe algorithm to a matheuristic (MH), schematically depicted in Fig. 1.

A pseudocode of the MH is given in Algorithm 3. Therein, instead of z , an alternative representation of a project schedule by a two-dimensional scheme Z is used. The rows of Z correspond to the projects p ($1 \leq p \leq |P|$), whereas the columns correspond to the activity indices $q = 1, \dots, d_p$ (since the numbers d_p need not to be identical, the number of entries per row can vary). The entry Z_{pq} in row p and column q indicates the time period t in which activity q of project p is processed. A project schedule Z must satisfy the requirements on the time windows and on the starting time relationships, as defined in terms of the representation z by Eqs. (3) and (4), respectively. In Algorithm 3, some design parameters have already been set to fixed numerical values; these values resulted from numerous pretests as the most successful parametrizations.

search to a solution Z' . The local search is based on a neighborhood defined by the current value of the parameter k (see below), and it is continued until (i) the neighbor solution is better than the current solution Z , (ii) the neighbor solution is better than the current incumbent Z^* , or (iii) a local optimum is reached (no better neighbor solution exists). The name “FirstImprovement” indicates that we already terminate the search as soon as for the first time, an improving neighbor has been found; we do not necessarily explore the entire k -neighborhood. The neighborhood definition, controlled by parameter k , is based on a local move operator (“ k -move”) that works as follows: First of all, for the given solution Z , a capacity profile is determined. This capacity profile indicates the costs of external capacity needed for each time period. Next, we identify that time period t for which the capacity costs are maximal, we consider the set of activities (p, q) that are executed in time period t , and we sort these activities in descending order according to their contribution to the external capacity costs in period t . A k -move to a neighbor solution consists in selecting k consecutive activities (p, q) from the list and in shifting the execution time periods of the selected activities either one period forward or backward. Afterward, in a subprocedure *RepairSolution()*, all predecessor periods and successor periods of an affected project are checked for feasibility (see constraints in Eq. (4)) and, if necessary, repaired. Finally, the objective function of the modified project schedule is evaluated, which requires a call of the Frank–Wolfe algorithm as explained in Sect. 3.4.1.

Algorithm 3 MH search scheme

```

1:  $n_r \leftarrow 0$ 
2:  $i_{\min} \leftarrow 50, \tilde{i} \leftarrow 1000$ 
3:  $Z, Z^* \leftarrow \text{InitialSolution}()$ 
4: while  $\text{time} < t_{\max}$  do
5:    $k \leftarrow 1$ 
6:   repeat ▷ //GLS-VND
7:      $n_r \leftarrow n_r + 1$ 
8:      $i_{\max} \leftarrow n_r^{0.5} \cdot i_{\min}$ 
9:      $Z' \leftarrow \text{FirstImprovement}(Z, Z^*, k, i_{\max})$  ▷ //Find the best neighbor
10:     $(Z^*, k) \leftarrow \text{neighborhoodChangeS}(Z^*, Z', k, \tilde{i})$  ▷ //Neighborhood change
11:     $Z \leftarrow Z^*$ 
12:   until  $k > k_{\max}$ 
13:    $Z \leftarrow \text{Perturbation}(Z, \lfloor \pi \cdot |P| \rfloor)$  ▷ //Perturbation
14:    $\text{time} \leftarrow \text{CpuTime}()$ 
15: end while
16: return  $Z^*$ 

```

Initial solution In the first step of MH, an initial solution Z for the project schedule is generated by solving the Expected value problem presented in Sect. 3.2. The solution Z is also used as the initialization of the incumbent solution Z^* .

First improvement The procedure *FirstImprovement()* takes a current project schedule Z and improves it by a local

The numerical accuracy of the solution value determination by the Frank–Wolfe algorithm is controlled through a parameter $i_{\max}$, which will be explained in the remark at the end of this subsection.

Neighborhood change In the procedure *neighborhood-ChangeS()*, the candidate solution Z' that resulted from the local search in *FirstImprovement()* is either accepted as the new incumbent Z^* or rejected, depending on its objective function value and on the outcome of a random event: If Z' is better than the current Z^* , then Z' is accepted in any case, i.e., Z^* is replaced by Z' , and the neighborhood size parameter k is reset to its initial value 1. Otherwise, we increase k by 1; moreover, with probability β , solution Z' is accepted though it is worse than Z^* (i.e., Z^* is replaced again by Z'), and with probability $1 - \beta$, the current incumbent solution Z^* is preserved. In the experiments, a value $\beta = 0.1$ turned out as a good choice.

The *repeat* loop in the algorithm performs a local search with a varying neighborhood definition, specified by the value of the parameter k . Such a procedure is usually called Variable Neighborhood Descent (VND). For the maximal neighborhood size k_{max} , we chose $k_{max} = 3$.

Perturbation The procedure *Perturbation()* generates a new solution from the current solution Z by performing random swaps, using a parameter that indicates the number of swaps. The procedure is triggered if all neighborhoods up to $k = k_{max}$ have failed to improve the incumbent solution. The number of swaps is chosen as equal to the product of a perturbation factor π and the number of projects $|P|$. Thus, we let the number of swaps depend on the problem size. A swap consists in exchanging the schedules (starting times of activities) of two randomly chosen projects and repairing the solution afterward, if the new schedule is infeasible. In our experiments, it turned out that $\pi = 0.15$ was a suitable value. Note that the perturbation procedure only works for projects with an identical number of activities.

Remark When applying the Frank–Wolfe algorithm as a subprocedure of the heuristic Algorithm 3, we adopt two strategies from Gutjahr and Froeschl (2013). First, in order to use computation time economically, the solution accuracy of the Frank–Wolfe algorithm can be increased gradually during the search process as a function of the number n_r of conducted neighborhood searches. It turned out that using $i_{max} = i_{min} \cdot n_r^{0.5}$ iterations in the Frank–Wolfe algorithm, where i_{min} is a fixed initial value, produced good results. For the evaluation in the procedure *neighborhood-ChangeS()*, we always use a comparably large number $\tilde{i}$ iterations. In the second strategy, we suppose that the current upper bound $B(z, i)$ of the expected external costs after i iterations of the Frank–Wolfe algorithm converges with (approximately) exponential speed to the true value $a = \min_x g(z, x)$ as $i \rightarrow \infty$. That is, we assume $B(z, i) \sim a + b \cdot \exp(-ci)$, and estimate the parameters a , b , and c from the three observations B_1 , B_2 and B_3 of the bound in iteration $i = 0.5 \cdot i_{max}$, $i = 0.75 \cdot i_{max}$ and $i = i_{max}$, respectively. This leads to the following estimate of the expected

external costs: $\tilde{\Theta}_z^* = a = (B_0 B_2 - B_1^2)/(B_0 - 2B_1 + B_2)$.

3.5 Sample average approximation

Our second solution approach for the stochastic project scheduling and staffing problem is the method of sample average approximation (SAA), see Kleywegt et al. (2002). The SAA method samples scenarios from the given distribution of the random events and approximates the given stochastic problem by the so-called *deterministic equivalent*, the problem resulting as the average over the scenarios. The deterministic equivalent can be solved by methods from deterministic optimization.

In the case of our problem, we draw a set of N random scenarios, described by the realizations $d_{psq}^{(1)}, \dots, d_{psq}^{(N)}$ of the random variables D_{psq} ($p \in P$, $s \in S$, $1 \leq q \leq d_p$). Thus, $d_{psq}^{(n)}$ denotes the demand of work package (p, s, q) in scenario n . The external work time required under scenario n for project p and skill s in time period t will be denoted by the new variable $y_{pts}^{(n)}$. The deterministic equivalent is then the problem (20)–(22), which is obviously a mixed-integer linear program. Note that $\Theta_{SAA}(z, x)$ represents the average of the external costs over all scenarios, which approximates the expected external costs, and that the values of the variables $y_{pts}^{(n)}$ result in (21) by a formula analogous to (9). The mixed-integer linear program (20)–(22) can be solved by standard solvers such as CPLEX. However, for realistic instance sizes and an appropriate number of samples, the problem (20)–(22) can become rather large, such that it may become difficult for the solver even to find a feasible solution within a reasonable time budget.

$$\min_{z, x} \Theta_{SAA}(z, x) = \frac{1}{N} \sum_{n \in N} \sum_{p \in P} \sum_{t=ES_p}^{LF_p} \sum_{s \in S} c_s^e y_{pts}^{(n)} \quad (20)$$

subject to constraints (3)–(7) and

$$\sum_{q \in \tau_{pt}} d_{psq}^{(n)} z_{pqt} \leq y_{pts}^{(n)} + \sum_{k \in K_s} x_{ptsk} \quad \begin{array}{l} p \in P, \\ t = ES_p, \dots, LF_p, \\ s \in S, \\ n \in N \end{array} \quad (21)$$

$$y_{pts}^{(n)} \geq 0 \quad \begin{array}{l} p \in P, \\ t = ES_p, \dots, LF_p, \\ s \in S, \\ n \in N \end{array} \quad (22)$$

4 Experimental results

4.1 Test instance generation

For our computational experiments, we generated test instances by the test instance generator proposed in Heimerl and Kolisch (2010). This test instance generator is inspired by data from the IT department of a large semiconductor manufacturer. However, we had to extend the test instances since in the present paper, the information about the demand of work packages is assumed to be uncertain. In particular, in addition to the parameters number of projects $|P|$, time window size $\gamma = LS_p - ES_p$, and number of skills per resource $|S_k|$, also a further parameter representing the degree of uncertainty (which will be explained later) had to be varied across the test instances. The database with the instance set is available for download at the website <http://phaidra.fhstp.ac.at/o:2529>.

Table 1 lists the parameter values for the basic instance structure. For this group of instances, the number of projects is 10, and the earliest start period ES_p of each project p is drawn from a uniform distribution between 1 and 7. The project length d_p is set to six periods for each project, and the planning horizon is defined as $T = 12$, which represents annual strategic project scheduling and personnel planning. We choose a time window size $\gamma = 1$ which means that for each activity q , there are two possible start times available, considering precedence constraints. With $S^{(p)}$ denoting the set of skills required by project p , and with $S^{(p,q)}$ denoting the set of skills required by activity q of project p , the test instance generator limits the number $|S^{(p)}|$ of skills per project by a pre-defined bound; moreover, for each activity q of a project p , the test instance generator specifies the number $|S^{(p,q)}|$ of required skills. For all skills s not contained in $S^{(p,q)}$, the values d_{psq} are set to zero. [For details concerning this aspect of the test instance generation, see Heimerl and Kolisch (2010)]. In our case, we chose $|S^{(p,q)}| = 2$ and limited the total number of skills per project by $|S^{(p)}| \leq 3$. Ten internal resources are assumed ($|K| = 10$), each owning $|S_k| = 2$ out of $|S| = 10$ skills. The resources have different efficiency values η_{sk} for each skill k they own; these efficiency values are drawn from a truncated normal distribution with an expected value of $\mu = 1$, a standard deviation of $\sigma = 0.25$, and minimum and maximum threshold values of 0.5 and 1.5, respectively. The available capacity per internal resource a_{kt} is 20 per time period. Note that we do not assume that a time period has a length of one time unit: In our above-mentioned interpretation of the test instances as referring to annual planning, the time unit is a day, and a time period extends over a month, so $a_{kt} = 20$ means that an employee works 20 days per month.

The external cost rates c_s^e differ for different skills and are drawn from a truncated normal distribution $TN_{a,b}(\mu, \sigma)$

Table 1 Basic instance structure

$ P = 10$	$ S^{(p)} \leq 3$
$ES_p \sim U(1, 7)$	$ K = 10$
$d_p = 6$	$ S_k = 2$
$T = 12$	$\eta_{sk} \sim TN_{0.5,1.5}(1, 0.25)$
$\gamma = LS_p - ES_p = 1$	$a_{kt} = 20$
$ S = 10$	$c_s^e \sim TN_{600,1000}(800, 100)$
$ S^{(p,q)} = 2$	$\rho = 1.0$
$[c_{min}, c_{max}] = [0.7, 1.3]$	

with $\mu = 800$, $\sigma = 100$, $a = 600$ and $b = 1000$. In test instance generation, (planned) utilization ρ is defined as the ratio of the overall expected resource demand to available internal resource capacity. In the basic instance structure, we set $\rho = 1$. For each work package, the test instance generator computes an initial value $[E(D_{psq})]^{init}$ of the expected resource demand from the utilization ρ and the resource supply values a_{kt} . The actual expected value of the resource demand $E(D_{psq})$ is then drawn from a normal distribution, with mean $\mu = [E(D_{psq})]^{init}$ and a coefficient of variation $CV = 0.1$. For the basic instance structure, we assume a symmetric triangular distribution of the *actual* demand with parameters $(D_{psq}^{min}, D_{psq}^{mod}, D_{psq}^{max})$, where $D_{psq}^{mod} = E(D_{psq})$, $D_{psq}^{max} = D_{psq}^{mod} \cdot c_{max}$, and $D_{psq}^{min} = D_{psq}^{mod} \cdot c_{min}$. The interval $[c_{min}, c_{max}]$ controls the level of uncertainty; if $c_{min} = c_{max}$, we get the deterministic boundary case. For the basic instance structure, a moderate level of uncertainty is assumed by setting $c_{min} = 0.7$ and $c_{max} = 1.3$. Notice that because we use *symmetric* triangular distributions in our basic test instances, the expected value $E(D_{psq})$ is identical to the modal value D_{psq}^{mod} of the distribution.¹

The basic instance structure is varied then to obtain other instance structures, according to Table 2 which lists the parameters with their used values. For test instance generation, we use a *ceteris paribus design*, which means that we fix all parameters on the value of the basic instance structure and vary the value of one investigated parameter. This yields $4 + (4 - 1) + (6 - 1) + (3 - 1) = 14$ different instance structures. For each instance structure, we generate 10 instances, which leads to 140 test instances. Additionally, for each of these 150 test instances, four different levels of the degree of

¹ The choice of the triangular distribution for related models is discussed, e.g., in Law et al. (1991) and Wing Chau (1995). Law et al. (1991) see the triangular distribution as a rough model suitable for cases where there is limited data available or the costs of data collection are high. Also the Beta distribution would be a candidate for modeling the distributions of the random variables D_{psq} . However, as shown in Gutjahr and Froeschl (2013) for a related model, using a Beta distribution within the Frank–Wolfe framework comes at the price of a distinct increase in computation time, and the results are typically similar to those obtained by the triangular distribution.

Table 2 Parameters and experimental values for the test instance generation

Parameter	Experimental value
$ P $	10, 15, 20, 25
$\gamma = LS_p - ES_p$	0, 1, 2, 3
$ S_k $	1, 2, 4, 6, 8, 10
ρ	0.8, 1.0, 1.2

Table 3 Investigated degree of uncertainty

Parameter	Experimental values
$[c_{min}, c_{max}]$	[0.9, 1.1], [0.7, 1.3], [0.5, 1.5], [0.2, 1.8]

uncertainty (see Table 3), are investigated. This produces a total of 560 instances.

All tests are performed on a standard PC with an Intel Quad Core Processor. In detail, we used an Intel Xeon E3-1271 v3 processor (Frequency: 3,60GHz) with eight kernels and 32 Gigabytes working memory. All presented algorithms are implemented in Eclipse using Java version 1.7. We use the Java API of ILOG CPLEX version 12.4 for the SAA and the EV model formulations.

4.2 Parameter setting for the Frank–Wolfe algorithm

There are several design decisions that have to be made when using the Frank–Wolfe algorithm within our matheuristic framework procedure: First of all, two slightly different implementation variants of this algorithm were presented in Sect. 3.4.1. Secondly, the algorithm requires an initial solution; we consider two options for its choice. Finally, it has to be specified how the line search is done; again, two different options will be investigated.

Algorithmic variants We shall compare the basic Algorithm 1 (“old” in Table 4) to the modified Algorithm 2 (“new” in Table 4).

Initial solution For the determination of the initial solution, two alternative approaches are tested. First, we use the MILP solver CPLEX (“lp” in Table 4) to solve the deterministic linear staffing problem defined by the parameters $d'_{pst} = E(D'_{pst}) = \sum_q \bar{d}_{psq} z_{pqt}$. Secondly, we apply a greedy staffing heuristic (Felberbauer et al. 2016) to solve the same deterministic counterpart problem heuristically (“gh” in Table 4).

Line search method For the line search step of the algorithm, two methods are analyzed: The first method (“gs” in Table 4) uses Golden Section Search according to Kiefer (1953). The second method (“fs” in Table 4) follows the suggestion in Clarkson (2010): it refrains from determining the

arg min in line 13 of Algorithm 1 or line 12 of Algorithm 2, respectively, but uses instead in each iteration i a pre-defined step size $\vartheta^* = \vartheta^*(i)$ depending on the iteration index. The value of ϑ^* is calculated as $\vartheta^* = 2/(i + 2)$ ($i = 1, 2, \dots$). It is clear that the value of ϑ^* determined in this way does not produce the minimizer on the line segment between $u^{[i]}$ and $g^{[i]}$, but by the special choice of the step sizes (convergence to zero and finiteness of the sum of the squares), the convergence property of the Frank–Wolfe algorithm to the exact overall minimizer of Θ_z is preserved (for details, see Clarkson 2010). The advantage of the fixed step sizes scheme is that it does not require an evaluation of the function values Θ_z during the execution of the algorithm; it suffices to evaluate the derivatives of Θ_z .

Combining the two alternative options for each of the three design decisions indicated above, we get $2^3 = 8$ different design variants of the Frank–Wolfe algorithm. The following results compare the performance of these eight design variants. For each design variant, we shall report its average solution value sv and its average computation time ct at 1000 randomly selected time schedules $z = (z_{pqt})$ for a single fixed problem instance generated according to the basic instance structure. In a pretest, it turned out that a number $i_{max} = 1000$ of iterations was sufficient to get close enough to the value Θ_z achieved by a much higher number 10^6 of iterations. Therefore, we used $i_{max} = 1000$ for all design variants.

The results are shown in Table 4. The comparison between the two algorithmic variants Algorithm 1 and the new Algorithm 2 (immediately applying the best partial derivative for the update of the project plan) shows a slight superiority of Algorithm 1. For the decision on the used initial solution method, the test shows that the computation time for solving the deterministic staffing problem takes in the average ≈ 3.73 ms using the LP solver and ≈ 0.26 ms using the greedy heuristic. On the other hand, the expected cost of the initial staffing plan x in iteration $i = 1$ according to the greedy heuristic is $\approx 9\%$ higher than the one obtained by the exact LP solver. Nevertheless, applying a two-tailed sign test to the final results for solution values and computation times shows no significant difference between the performance of the LP solver and the greedy heuristic (significance level $\alpha = 0.05$). For larger instances, where the LP solving time increases rapidly, the greedy heuristic can become the only feasible alternative, so that in total, we may give a preference to the greedy heuristic. Concerning the line search method, finally, it can be seen that the pre-defined step sizes scheme clearly outperforms the Golden Section Search: Although Golden Section Search provides faster improvements in the first few iterations than the step sizes scheme, the solution quality after 1000 iteration is not better, and the computation time per iteration is ≈ 200 times higher.

Table 4 Average solution value (sv), standard deviation of the solution value, average computation time (ct), and standard deviation of the computation time for the eight Frank–Wolfe design variants. Last two columns: relative values compared to those of the best design variants. Based on $i_{max} = 1000$ iterations and 1000 schedules z of a fixed instance

FW design variant	Absolute				Relative	
	sv [CU]	sv-stdev [CU]	ct [ms]	ct-stdev [ms]	Δsv [%]	Δct [%]
1. lp-fs-new	286, 341.3	0.00	102	6	0.00	3.15
2. lp-fs-old	286, 341.4	0.00	213	14	0.00	115.00
3. lp-gs-new	286, 342.1	0.00	22, 439	118	0.00	22, 515.00
4. lp-gs-old	286, 354.5	0.00	14, 448	112	0.00	14, 462.00
5. gh-fs-new	286, 341.2	0.00	99	53	0.00	0.00
6. gh-fs-old	286, 341.3	0.00	209	9	0.00	110.00
7. gh-gs-new	286, 498.4	0.00	30, 015	200	0.05	30, 151.00
8. gh-gs-old	286, 554.8	0.00	15, 095	88	0.07	15, 113.00

Summarizing, the best-performing Frank–Wolfe variant, i.e., variant 5 in Table 4, uses the immediate application of each best partial derivative for the update of the project plan, the greedy heuristic for the calculation of the initial staffing solution, and the pre-defined step sizes scheme.

4.3 Parameter setting for sample average approximation

The crucial parameter of the sample average approximation procedure is the number N of sampled random scenarios. Therefore, the following pretests have been conducted to find an appropriate value of N . It is clear that the objective function of (20) is only an approximation to the true objective function, such that even if the SAA problem is solved exactly, we do not necessarily obtain the exact solution of the original problem. To explore the tradeoff between the two effects of increasing the value of N , namely to improve the accuracy of the objective function estimation on the one hand, and to increase the computation time on the other hand, a subset of our instances has been investigated. We chose the instances of those instances structures where the number of projects is varied as $|P| \in \{10, 15, 20, 25\}$, the time window size is $\gamma \in \{0, 1, 2, 3\}$, and the number of skills per resource as well as the utilization are fixed to the values of the basic instance structure, i.e., $|S_k| = 2$ and $\rho = 1$. The degree of uncertainty was varied as $[c_{min}, c_{max}] \in \{[0.9, 1.1], [0.7, 1.3], [0.5, 1.5], [0.2, 1.8]\}$. For these instances, we varied the sample size as $N \in \{10, 20, 30, \dots, 100\}$ and analyzed the solution time and the achieved solution quality. By the SAA model from Sect. 3.5 with sample size N , we compute the solution (z_{SAA}^*, x_{SAA}^*) , where z_{SAA}^* is the optimal project plan and x_{SAA}^* is the optimal staffing plan. Now, we compare the obtained objective function value $\Theta_{SAA}(z_{SAA}^*, x_{SAA}^*)$ of (20) to the true evaluation $E[G(z_{SAA}^*, x_{SAA}^*, \xi)]$ of the solution (z_{SAA}^*, x_{SAA}^*) according to the underlying exact probability model (cf. the notation in Sect. 3.1). The relative gap between the two eval-

Fig. 2 Performance of solution time and solution gap in dependence of sample size N for $|P| = 20$, $\gamma = 1$, $\rho = 1.0$, $|S_k| = 2$, and $[c_{min}, c_{max}] = [0.7, 1.3]$

uations is described by

$$\frac{|E[G(z_{SAA}^*, x_{SAA}^*, \xi)] - \Theta_{SAA}(z_{SAA}^*, x_{SAA}^*)|}{E[G(z_{SAA}^*, x_{SAA}^*, \xi)]}. \quad (23)$$

In Fig. 2, the solution time and the relative solution gap according to Eq. (23) as well as their 95% confidence intervals are depicted in dependence of the sample size N . We show here the special case of $|P| = 20$ projects and the other parameters as in the basic instance structure. It can be observed that for a sample size of $N = 100$, the solution gap is $\approx 0.5\%$, i.e., the average objective function value over the scenarios can be considered as a good estimate for the expected external costs. A further observation is that the solution time of the SAA model varies to a considerable extent. This behavior points out a first drawback of relying on the exact solution of the SAA model to solve our problem.

These results were extended by solving all test instances of the instance subset specified at the beginning of this subsection. With an appropriate sample size of $N = 100$ and a computation time limit of 360 sec for the CPLEX solver, we observed that CPLEX was able to solve the SAA problems for all considered test instances. In the average over all considered test instances, a relative gap of 0.71% according to Eq. (23) was obtained.

Fig. 3 Solution value performance of the sample average approximation approach (SAA, circles) and the matheuristic (MH, diamonds) with respect to the given time budget for the instance structure with $|P| = 20$, $\gamma = 1$, $\rho = 1.0$, $|S_k| = 2$, and $[c_{min}, c_{max}] = [0.7, 1.3]$

4.4 Matheuristic vs. sample average approximation

This section reports on the numerical comparison of the two presented solution methods, i.e., the developed matheuristic (MH) and the sample average approximation (SAA) approach. Note that by fixing a sample size N , the computation time consumed by the SAA approach is already defined. To ensure a fair comparison between SAA and MH, we computed, for a given test instance structure, the average computation times of the SAA approach for sample sizes $N = 10, 20, \dots, 100$ and used each of these ten values as the time budget (termination criterion) for a corresponding run of the MH approach. For each problem instance structure, ten random instances were generated, and for each of these generated instances, ten optimization runs (with different seed values for the random number generator) were executed. This led to 100 solution values for MH and 100 solution values for SAA per time budget. The averages of these solution values for a special instance structure are depicted in Fig. 3. The solution values have been computed based on the determination of the exact objective function value of the solution (z, x) provided by the respective approach, i.e., the value $E[G(z, x, \xi)]$. The reader will observe that Fig. 3 only contains 9 solution time values instead of 10, as one would expect. This is because sample sizes $N = 80$ and $N = 90$ led to identical solution times, cf. Fig. 2.

The results show that the solution quality of MH is less sensitive to the time budget than that of SAA. Applying a two-tailed sign test to the results of each of the ten different time budgets, we found that the MH results were significantly superior, at significance level $\alpha = 0.05$, for the smallest time budget 14 sec which was the time budget produced by sample size $N = 10$. For larger time budgets, the sign test could not confirm statistically significant superiority at level $\alpha = 0.05$ of either SAA over MH or vice versa. However, this lack of significance may be due to our small sample size of ten instances (in order to ensure indepen-

dence, we had to take the average over the ten runs of each instance). Summarizing, in the considered instances structure, MH and SAA produce results of a comparable quality in the medium computation time range, with a slight but non-significant advantage for SAA. For small computation times, MH is superior.

Next, we compare SAA, applying sample size $N = 100$, to MH, using for both the same time limit of 360 s. Let (z_{MH}^*, x_{MH}^*) and (z_{SAA}^*, x_{SAA}^*) denote the optimal project schedule and staffing plan according to MH and SAA, respectively. To do the comparison, we compute the normalized difference

$$\begin{aligned} & (E[G(z_{SAA}^*, x_{SAA}^*, \xi)] \\ & - E[G(z_{MH}^*, x_{MH}^*, \xi)]) / E[G(z_{SAA}^*, x_{SAA}^*, \xi)]. \end{aligned} \quad (24)$$

of the expected external costs. Averaging this measure over all test instances, we obtained a value of 0.1288, indicating that in the average, the SAA solutions are by 12.88% worse than the MH solutions. However, this result should be interpreted cautiously. The considerable difference can mainly be attributed to test instances where the expected external costs of the optimal solutions are almost zero. Such a situation easily occurs in instances where the number $|S_k|$ of skills per resource is high. Investigating the same measure only for the basic instance structure, we find that there is no significant difference between the two solution methods.

Real-world project scheduling problems are often rather large. To investigate which effect an increasing size of the problem instance exerts on the comparison between MH and SAA, we generated test instances where both the number of projects and the number of resources were increased simultaneously as $|P| = |K| = 50, 100, 150, 200, 250$, and the parameter values $\gamma = 2$, $|S_k| = 2$ and $\rho = 1$ from Table 1 were used. For the SAA model, we applied CPLEX with default values and set the time limit to 10 h, which is an appropriate time budget for a tactical planning problem. We found that the solver could not provide a feasible integer solution for 40% of these instances; the share of solvable instances rapidly drops as $|P|$ and $|K|$ become larger than 150. However, even for the instances for which the SAA model can be feasibly solved, the SAA solution is in the average by 45.57% worse than the MH solution. Additionally, MH needs only a fraction of the SAA solution time to find a good solution.

We conclude that the developed matheuristic is a robust solution procedure that performs well for small and medium-sized test instances, and offers good solutions also where the SAA formulation fails to return a feasible solution. The main drawback of the SAA approach is its poor reliability with respect to the identification of a feasible solution and its volatile solution time. Nevertheless, for not too large instances, the SAA approach is a promising alternative to

Table 5 Relative accuracy BB^{rel} of the EV problem, measured in %, (i) for the basic instance structure and (ii) for all instances, in dependence of the degree of uncertainty

Instances	$[c_{min}, c_{max}]$ [0.9, 1.1]	[0.7, 1.3]	[0.5, 1.5]	[0.2, 1.8]	$\emptyset$
Basic instance structure	3.96%	11.32%	17.94%	26.68%	14.97%
Avg. over all instances	33.33%	41.76%	47.58%	54.20%	44.22%

the application of a (partially) heuristics-based method, especially in cases where powerful hardware is available.

4.5 Deterministic vs. stochastic planning

Is the advantage of treating the given project scheduling and staffing method as a stochastic optimization problem substantial enough to justify the increased computational effort, compared to a simplified, deterministic formulation? To shed light on this question, we deepen our experimental analysis in the following two subsections. In Sect. 4.5.1, the accuracy of using the solution value of the expected value problem as a forecast for budget planning is investigated, whereas in Sect. 4.5.2, the value of the stochastic solution is discussed.

4.5.1 Accuracy of the expected value problem

The solution of the expected value (EV) problem (8)–(10) provides a manager with a project schedule z_{EV}^* , a staffing plan x_{EV}^* , and a corresponding solution value $\Theta_{EV}(z_{EV}^*, x_{EV}^*)$. The obtained forecast $\Theta_{EV}(z_{EV}^*, x_{EV}^*)$ of the external costs could be used as an input for the budget planning process. If the demand information is actually stochastic, this forecast will be rather rough, and it will tend to underestimate the true costs. Evaluating the obtained project and staffing plan based on the stochastic model by computing $E[G(z_{EV}^*, x_{EV}^*, \xi)]$ gives a hint of how the forecast provided by solution of the EV problem will perform in the stochastic environment. The difference

$$BB = E[G(z_{EV}^*, x_{EV}^*, \xi)] - \Theta_{EV}(z_{EV}^*, x_{EV}^*) \quad (25)$$

could be considered as a *budget bias* caused by a deterministic solution approach. Table 5 shows the relative bias BB^{rel} of the EV problem solution value, i.e., the quotient

$$BB^{rel} = (E[G(z_{EV}^*, x_{EV}^*, \xi)] - \Theta_{EV}(z_{EV}^*, x_{EV}^*)) / E[G(z_{EV}^*, x_{EV}^*, \xi)], \quad (26)$$

for the basic instance structure and for all instances, in dependence of the degree of uncertainty. We see that with increasing uncertainty, the bias increases, and that it reaches fairly large values. For the basic instance structure, already under a moderate level of uncertainty of $[c_{min}, c_{max}] = [0.7, 1.3]$, the use of a deterministic planning approach leads

to an underestimation of the external costs by 11%. Averaged over the instances from all instance structures (with ten random instances from each instance structure), this deviation is even distinctly higher (42%). According to these results, we conclude that the deterministic EV approach to the considered project scheduling and staffing problem leads to a systematic underestimation of external costs and can, as a consequence, seriously threaten the budget plan.

4.5.2 Value of the stochastic solution

Whereas the last subsection investigated the difference between the predicted and the true costs of the EV solution, we turn now to the question of how much worse the EV solution is in comparison with the solution of the stochastic optimization problem. This latter difference indicates the value of taking the stochasticity of the demand into account instead of using the simplified deterministic EV model for the planning process. As the solution procedure for the stochastic optimization problem, we choose the MH approach. As before, let (z_{EV}^*, x_{EV}^*) and $E[G(z_{EV}^*, x_{EV}^*, \xi)]$ denote the EV solution and its expected external cost, respectively, and let (z_{MH}^*, x_{MH}^*) and $E[G(z_{MH}^*, x_{MH}^*, \xi)]$ denote the MH solution and its expected external cost, respectively. The difference

$$VSS = E[G(z_{EV}^*, x_{EV}^*, \xi)] - E[G(z_{MH}^*, x_{MH}^*, \xi)] \quad (27)$$

is called the *value of the stochastic solution* (VSS); it describes the cost savings achieved by applying the stochastic solution approach instead of the deterministic one. We define the relative value of the stochastic solution as the quotient

$$VSS^{rel} = (E[G(z_{EV}^*, x_{EV}^*, \xi)] - E[G(z_{MH}^*, x_{MH}^*, \xi)]) / E[G(z_{EV}^*, x_{EV}^*, \xi)]. \quad (28)$$

In Fig. 4, the absolute and the relative VSS are depicted in dependence of the degree of uncertainty for the instances of the basic instance structure. Unsurprisingly, both the absolute and the relative VSS increase as uncertainty increases. The VSS rapidly grows with increasing degree of uncertainty, and it reaches about 16% for $[c_{min}, c_{max}] = [0.2, 1.8]$. Applying a linear regression, we statistically confirm ($\alpha = 0.001$)

Fig. 4 Absolute and relative value of the stochastic solution for the instances of the basic instance structure, in dependence of the degree of uncertainty

the intuitive conjecture that both the absolute and the relative VSS are positively correlated with the degree of uncertainty (for VSS and relative VSS, we get correlation coefficients of $R = 0.936$ and $R = 0.753$, respectively). In some other instance structures, the potential of the stochastic solution approach is even higher. For example, for large values of the number of projects and of the number of skills per resource, this gain reaches values between 30 and 100%. The considerable size of the VSS confirms the need of applying stochastic optimization techniques in project scheduling and staffing.

4.6 Costs of uncertainty

In Fig. 5, we plot (i) the expected external costs and (ii) the relative expected external cost increments, both in dependence of different degrees of uncertainty, for the instances of the basic scenario structure. In this figure, the relative expected external cost increments indicate by which percentage the expected external costs increase if the situation $[c_{min}, c_{max}] = [0.9, 1.1]$ of low uncertainty is replaced by a higher uncertainty interval $[c_{min}, c_{max}]$ for the demand distributions. As one would anticipate, the expected external costs increase as uncertainty increases. A linear regression for the absolute value of the expected external cost shows that the coefficient of correlation is $R = 0.119$, with a p value of 0.017 (significance at level $\alpha = 0.05$). Comparing the “almost deterministic” situation $[c_{min}, c_{max}] = [0.9, 1.1]$ to the situation $[c_{min}, c_{max}] = [0.2, 1.8]$ of poor information on demands, we find a gap of 12%. We would like to emphasize that a degree of uncertainty represented by $[c_{min}, c_{max}] = [0.2, 1.8]$ (i.e., a distribution allowing real efforts of work packages to exceed estimated efforts by up to 80%) is not extreme from an applied point of view: In areas as software engineering or in the construction industry, even higher overruns occur. Therefore, also the costs of uncertainty can be considerable in practice.

Fig. 5 Plot of (i) expected external costs and (ii) relative expected external cost increments compared to the situation with low stochasticity, $[c_{min}, c_{max}] = [0.9, 1.1]$, as a function of $c_{max} - c_{min}$, for the instances of the basic instance structure

4.7 Influence of parameters on external costs

In Heimerl and Kolisch (2010), the authors investigate in a deterministic context how parameters as the number of projects, the time window size etc. influence the optimal costs. We shall extend now their results to the stochastic context of the present paper. In the present subsection, we use the *ceteris paribus design* explained in Sect. 4.1. That is, the parameters of the basic instance structure of Table 1 are applied, with the exception of modifying one single parameter among the parameters in Table 2.

Influence of the number of projects. First, we investigate the influence of increasing the number of projects on the resulting expected external costs. As Heimerl and Kolisch (2010), we keep the total resource demand constant while increasing the number of projects, which means that for a larger number of projects, the work packages become smaller. This change increases the flexibility of the planner, so one expects decreasing costs of the optimal solutions. This was confirmed indeed in Heimerl and Kolisch (2010) for the deterministic context. We obtained similar results in the stochastic context: In Fig. 6, we depict the expected external costs for four levels of uncertainty as a function of the number of projects (with fixed total demand). Additionally to the mean values, we plot the 95% confidence interval to account for the randomness in test instance generation. Apart from the already known effect that higher degree of uncertainty leads to higher external costs, one can see that all levels of uncertainty show the same behavior for an increasing number of projects: The conjecture that a larger number of smaller work packages is easier to balance across the planning horizon than bigger and fewer work packages is confirmed. Also as expected, we observe that from a certain value on, the potential of the flexibility achieved by reducing work packages sizes diminishes.

Influence of the time window size. In a similar way as it was done for the influence of the number of projects, we

Fig. 6 Plot of the expected external costs as a function of the number of projects $|P|$

Fig. 7 Plot of the expected external costs as a function of the number of skills per resource $|S_k|$

investigated also the influence of the time window size on the costs. The results showed that in our test instances, an additional degree of freedom, i.e., a time window size change from zero to one, led to cost savings of around 20%. A further cost reduction of up to 50% was achieved by increasing the time window size γ from 1 to 2. For time window sizes larger than two time periods, the costs did not decrease further.

Influence of the number of skills per resource. In Fig. 7, we plot the expected external costs in dependence of the number of skills per resource. Please note that the situation where the number of skills per resource is $|S_k| = 10$ represents a situation when all resources possess all skills. It can be observed that the external costs decrease monotonically with increasing $|S_k|$. Two special findings may be particularly important. First, in the context of the considered set of instances (basic instance structure, with modifications only with respect to $|S_k|$), a situation where all employees are extremely specialized, i.e., possess only one skill per person, leads to about the double expected external costs in comparison with a situation where the employees have two skills per person. A further investment in training that causes an increase in the number skills per person from two to four leads to another significant decrease in the expected external costs. Secondly, from a value of about four skills per person on, the investment in additional skills achieves only very limited cost savings. Of course, the quantitative amount of cost reduction depends on

Fig. 8 Plot of the expected external costs as a function of the utilization factor ρ

Fig. 9 Plot of (i) expected external costs and (ii) relative expected external cost increments compared to the situation with low stochasticity, $c_{\max} - c_{\min} = 0.2$, as a function of $c_{\max} - c_{\min}$, for the instances with right-skewed distribution

the specific characteristics of the test instance set (especially on the chosen parameter value for utilization, here: $\rho = 1$). Nevertheless, the results suggest the managerial insight that (i) multi-skilled resources can lead to significant cost savings, but (ii) companies should be aware that over-qualification allows no return on investment.

Influence of the utilization factor. Finally, in Fig. 8, the expected external costs are visualized as a function of the utilization factor ρ , the ratio of the expected demand to the available time capacities of the resources. As expected, Fig. 8 shows that the external costs increase as the utilization increases. In more detail, we find that when starting at 100% utilization, a 20% decrease in utilization leads to a 57% decrease in the expected external costs, whereas a 20% increase in the utilization leads to a rise of costs by 23%.

4.8 Asymmetric work time distributions

In the previous tests, we assumed that the triangular distribution of the variables D_{psq} was *symmetric*. In practice, distributions of work times are often right-skewed. Therefore, we checked whether or not substantially different results were obtained by a replacement of the symmetric distributions by right-skewed triangular distributions. For this

purpose, we took the basic instance structure and changed it as follows: The modal value was defined as $c_{mode} = c_{min} + (c_{max} - c_{min}) \cdot 0.25$, where c_{min} and c_{max} are the minimum and the maximum value of the distribution, respectively. For a series of four test instances, we fixed the expected value $c_{expected} = (c_{min} + c_{max} + c_{mode})/3$ of the triangular distribution to 1. On this constraint, each of the four instances was constructed in such a way that the length $c_{max} - c_{min}$ of the support interval of the distribution was gradually increased, taking the values 0.2, 0.6, 1.0 and 1.6, respectively, which corresponds to increasing uncertainty. Figure 9 shows a plot of the results. We see that the plot is very similar to Fig. 5. We conclude that the skewedness of the distribution has little impact on the outcomes.

5 Conclusion

We developed two solution approaches for a stochastic project scheduling and staffing problem under uncertainty on required efforts. The scheduling decision is described by the choice of the start times of activities during the planning horizon. Each activity can consist of several work packages, where each work package requires a stochastic amount of effort exerted in one specific skill. The staff assignment decision matches available human resources (with heterogeneous skills) with the work package requirements of the activities. Demand for work time that is not covered by the assigned internal resources has to be satisfied by paying external work.

For our first solution approach, a “matheuristic” combination of a metaheuristic with a convex optimization procedure, we decompose the problem into a project scheduling subproblem and staffing subproblem. The project schedules are optimized by an iterated local search heuristic, using variable neighborhood descent as a solution component. The search is guided to initially tackle time periods leading to high external cost. The staffing subproblem is a convex optimization problem which we solve by the Frank–Wolfe algorithm. Different design variants of this algorithm are investigated in application to our problem. Our second solution approach uses a sample average approximation model of the stochastic optimization problem. This yields mixed-integer programming formulation that can be solved by CPLEX. For this second approach, a crucial parameter is the chosen sample size.

Experimental results for synthetically generated test instances show that the matheuristic is a robust solution procedure that performs well for small as well as medium-sized test instances and provides solutions even in cases where the SAA model fails to return a feasible solution. Nevertheless, up to medium-sized instances, the sample average approximation is also a good choice. We find that deterministic planning replacing the stochastic optimization problem by the corresponding expected value problem bears the risk

of drastically underestimating external costs. Moreover, we demonstrate that the value of the stochastic solution, which is a measure for the cost savings achievable by a stochastic solution approach instead of using the deterministic expected value problem, is considerable, especially for moderate and high levels of uncertainty.

Our experiments confirmed some of the managerial insights that have been found (in a deterministic context) in Heimerl and Kolisch (2010) to be valid also in the context of uncertainty on required efforts. This holds especially for the potential as well as the limitations of multi-skilled resources, of small work packages and of large time windows with respect to possible cost savings. The main difference to the deterministic context of Heimerl and Kolisch (2010) is that a much more conservative planning strategy is necessary in presence of uncertainty in order to avoid excessive additional costs by required external work.

Finally, let us point out an important topic of future research: Whereas our approach only assumes substitutability of internal by external resources, the articles Ingels and Maenhout (2017a, b) focus on (within-skill or between-skill) substitutions of internal resources by other internal resources, which is another type of recourse action that is frequently deployed for absorbing work load peaks. It would be very interesting to extend the model presented here to these forms of reactive planning and to develop suitable solution techniques for such an extension.

Another topic of future research might be risk aversion. Our model is risk-neutral, which fits well to a situation of multi-project management with several small or medium-sized projects and a long-term perspective, but may be not appropriate anymore for a situation where very big problems cause specific risks. For such situations, extensions to risk-averse optimization approaches should be studied.

Acknowledgements Open access funding provided by St. Pölten - University of Applied Sciences.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Artigues, C., Demassey, S., & Neron, E. (Eds.). (2008). *Resource constrained project scheduling: Models, algorithms, extensions and applications*. Control systems, robotics and manufacturing series. Wiley-ISTE.
- Artigues, C., Leus, R., & Nobibon, F. T. (2013). Robust optimization for resource-constrained project scheduling with uncertain activity durations. *Flexible Services and Manufacturing Journal*, 25(1–2), 175–205.

- Barz, C., & Kolisch, R. (2014). Hierarchical multi-skill resource assignment in the telecommunications industry. *Production and Operations Management*, 23(3), 489–503.
- Clarkson, K. L. (2010). Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Transactions on Algorithms (TALG)*, 6(4), 63.
- De Bruecker, P., Van den Bergh, J., Beliën, J., & Demeulemeester, E. (2015). Workforce planning incorporating skills: state of the art. *European Journal of Operational Research*, 243(1), 1–16.
- Felberbauer, T., Dörner, K. F., & Gutjahr, W. J. (2016). Hybrid metaheuristics for project scheduling and staffing, considering project interruptions and labor contracts. In H. Dawid, K. Dörner, G. Feichtinger, P. M. Kort, & A. Seidl (Eds.), *Dynamic perspectives on managerial decision making, dynamic modelling and econometrics in economics and finance* (Vol. 22). Cham: Springer.
- Gutjahr, W. J. (2015). Bi-objective multi-mode project scheduling under risk aversion. *European Journal of Operational Research*, 246, 421–434.
- Gutjahr, W. J., & Froeschl, K. A. (2013). Project portfolio selection under uncertainty with outsourcing opportunities. *Flexible Services and Manufacturing Journal*, 25(1–2), 255–281.
- Heimerl, C., & Kolisch, R. (2010). Scheduling and staffing multiple projects with a multi-skilled workforce. *OR Spectrum*, 32(2), 343–368.
- Ingels, J., & Maenhout, B. (2017a). Employee substitutability as a tool to improve the robustness in personnel scheduling. *OR Spectrum*, 39(3), 623–658. <https://doi.org/10.1007/s00291-017-0476-0>.
- Ingels, J., & Maenhout, B. (2017b). A memetic algorithm to maximise the employee substitutability in personnel shift scheduling. In B. L. I. M. Hu (Ed.), *Evolutionary computation in combinatorial optimization (EvoCOP 2017)* (Vol. 10197, pp. 44–59). Lecture Notes in Computer Science Cham: Springer.
- Kiefer, J. (1953). Sequential minimax search for a maximum. *Proceedings of the American Mathematical Society*, 4(3), 502–506.
- Kleywegt, A. J., Shapiro, A., & Homem-de Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2), 479–502.
- Kolisch, R., & Heimerl, C. (2012). An efficient metaheuristic for integrated scheduling and staffing it projects based on a generalized minimum cost flow network. *Naval Research Logistics (NRL)*, 59(2), 111–127. <https://doi.org/10.1002/nav.21476>.
- Law, A. M., Kelton, W. D., & Kelton, W. D. (1991). *Simulation modeling and analysis* (Vol. 2). New York: McGraw-Hill.
- Lourenço, H. R., Martin, O. C., & Stützle, T. (2010). Iterated local search: framework and applications. In M. Gendreau & J. Y. Potvin (Eds.), *Handbook of metaheuristics* (Vol. 2, pp. 363–397). New York: Springer.
- Mladenović, N., & Hansen, P. (1997). Variable neighborhood search. *Computers & Operations Research*, 24(11), 1097–1100.
- Van den Bergh, J., Beliën, J., De Bruecker, P., Demeulemeester, E., & De Boeck, L. (2013). Personnel scheduling: a literature review. *European Journal of Operational Research*, 226(3), 367–385.
- Wing Chau, K. (1995). The validity of the triangular distribution assumption in monte carlo simulation of construction costs: empirical evidence from hong kong. *Construction Management and Economics*, 13(1), 15–21.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.