2015

Resource Planning in Engineering Services Firms

Vincent Hargaden
University College Dublin, vincent.hargaden@ucd.ie

Jennifer K. Ryan
University of Nebraska - Lincoln, jennifer.ryan@unl.edu

Resource Planning in Engineering Services Firms

Vincent Hargaden and Jennifer K. Ryan

Abstract—**We develop a model to enable engineering professional services firms to improve the management of their competitive resources, i.e., skilled engineers, to be better able to respond to customer demand. The model was informed by semistructured interviews with senior executives from engineering, information technology (IT) services and technical consulting firms. As a result, we capture many of the complexities associated with the resource planning process in the professional engineering services sector. In the resulting model, the key attributes of supply, demand, and operations constraints are identified. Based on information obtained from the interviews, a number of test firms are created. We then use these test firms to study the impact of various resource planning policies on firm performance. These policies include the skill mix and profile of skilled employees, limits on the number of concurrent projects to which an employee can be assigned, and policies governing employee cross-training and hiring. The impact of these policies is evaluated in terms of business metrics, such as the project completion rate and net revenue. Finally, our model is extended to capture a multiphase rolling planning horizon, where projects may span multiple phases with the goal of ensuring consistency in employee assignment to projects.**

Index Terms—**Linear optimization, management of scientists and engineers, project scheduling, resource allocation systems, resource management.**

I. INTRODUCTION

IN RECENT years, the issue of human resource planning in service organizations has gained significant attention and has been identified as an increasingly important topic for research. While the concept of workforce planning is not new, the decline of the discipline since the 1970s has been described and the argument made for increased accuracy in the process, in particular when capturing the attributes of employees and customer demand [1]. To achieve this, many companies have begun to focus on "Talent Analytics" [2], which involves the use of analytical models, rather than a reliance on "gut instinct," as a method for improving competitive advantage, by analyzing employee (i.e., "talent") data in order to maximize productivity.

Much of the existing research on workforce planning has focused on manufacturing supply chains or on mass services, with an emphasis on tactical decision-making, such as the development of schedules for employees and the assignment of employees to specific tasks. A recent call for increased research and development to support *professional services* describes how this sector lags behind manufacturing and other service sectors [3]. One of the key opportunities identified is in the area of resource allocation and human capital management. While there are many different types of firms and service offerings within the professional engineering services sector, our study focuses on identifying the many common workforce planning characteristics across these firms. We develop an integer programming model for the optimization of the human resource planning and workforce allocation process in professional engineering services firms. In developing the model, the attributes of both professional service employees and customer projects are identified. We then use this model to understand the impact of various workforce planning policies on the performance of these types of firms.

II. LITERATURE REVIEW

The workforce planning process can be divided into three decision levels [4]. The first is at the level of the firm, where workforce policies are developed. The second deals with staff planning, including "hiring, discharge, training and reallocation decisions." The third develops short-term schedules or rosters for employees and is dependent on constraints set in the previous two levels. The research presented in this paper is positioned at the first two levels, i.e., the investigation of workforce planning policies and their impact on staff planning. The scheduling of tasks and the rostering of each staff member to those tasks is beyond the scope of this research.

A. Professional Services

Professional service firms have become an important industry group given the shift away from capital-intensive, manufacturing-based industry, toward a knowledge-intensive industry [5]. Professional service firms are regarded as critical players in the 21st century as they "sell expertise" and deliver customized solutions to large corporations and governments [6].

The professional services segment consists of highly unique and customized activities that are heavily dependent on human resources. It is suggested that the most valuable assets of the 21st century organization will be *knowledge workers* and that a critical issue for organizations will be the productivity of these workers [5]. It is also argued that the main challenges organizations will face revolve around personnel policy, the attraction and retention of workers in a multilocation environment and the conversion of knowledge worker productivity into increased capacity [5].

Another key characteristic of professional services firms is that the unit of sales is typically a project contract [3]. Projects by their nature are unique activities with specified time frames, requiring human and nonhuman resources [7]. Most professional service firms operate as multiproject systems [8], consisting of a portfolio of customer projects with different start and end dates, which implies that firms must develop an adequate human

V. Hargaden is with the University College Dublin, Dublin, Ireland (e-mail: vincent.hargaden@ucd.ie).

J. K. Ryan is with the Department of Management, University of Nebraska, Lincoln, NE 68588 USA (e-mail: jennifer.ryan@unl.edu).

resource allocation or capacity management process to make the best use of the firm's resources [9].

B. Workforce Planning

Executives consistently indicate in surveys that human capital is one of their most important assets, but find it difficult to convert that human capital into a source of competitive advantage [10]. However, a number of firms have recently attempted to address this issue. These include IBM, which has developed "OptiMatch" at its Watson Research Center, with the goal of matching people to projects in its Global Business Services organization [11]–[13] and Hewlett Packard, which has developed a new workforce planning and scheduling approach at its Palo Alto Laboratories [14]. These tools were developed to address specific business objectives such as improving the workforce assignment process. This is analogous to staff scheduling and rostering in the mass services and service shop categories.

1) Workforce Planning in Services: Due to the high variability in customer demand, as well as the difficulty of carrying an "inventory" of labor, conventional manufacturing resource planning processes are difficult to apply in a professional services resource planning context [15]. In addition, there is a lack of models and analytic tools for the business service industry compared to the manufacturing sector. In a call for the development of analytic models to support both short term execution decisions and longer term capacity planning decisions, it is suggested that well-established manufacturing models may provide a useful starting point, but that models for professional services must take into account the inherent and unique characteristics of business services [3].

A significant portion of the literature on workforce planning related to services falls within the mass services and service shop segments, where tasks are considered to be homogeneous and where there is a low limit on the types of employee skills or customers. Workforce planning in the services also tends to deal with personnel scheduling or rostering, which is the process of constructing work timetables for staff so that customer demand can be satisfied [16]. This is evident in the transport services literature where issues crew scheduling and pairing are of particular interest [16], [17] and the telephone call centers literature where staff rostering is required [16], [18]. A comprehensive review and classification of the workforce planning literature has been recently carried out, which includes application areas, types of problems, model objective functions, skill modeling and solution techniques [19], [20]. While health care systems can overlap between the service shop (e.g., nurses) and the professional service (e.g., medical doctors) categories, the workforce planning literature has primarily considered nurse scheduling and rostering [16]. The objective is to have sufficient numbers of suitably qualified staff to cover patient demand, subject to regulations, shift distribution and personal choice [21]. In the hospital medical resident scheduling problem, staff must be assigned to various specialties as part of their multiyear training program, rotating between specialties with the rotation frequency dictated by each particular training program [22], [23].

In the financial services sector, the staff scheduling problem relates to audit engagements. Audit firms are concerned with developing a least cost schedule, while completing all the audits by the specified due dates and ensuring that the audit task requirement is performed by an auditor with the appropriate rank and experience [24]. The work allocation process is also considered in engineering consulting organizations. Unlike in mass service or service shop environments, engineering consulting firms "specialize in the production of knowledge by adding value to information in their transformation processes" [25].

C. Summary of Contributions

Our research focuses on an area not previously considered in the workforce planning literature: the impact of various managerial-level workforce planning policies on professional engineering services firms. Informed by detailed case study interviews with senior executives in engineering, information technology (IT) services and consulting organizations, we develop a comprehensive workforce planning model for professional service firms. This model takes into account the operational processes and constraints of these entities. It captures the nature of project contracts, the position of employees in a firm's hierarchy, the matrix structure of firms, and the ability of employees to work on more than one project at a time. The model enables these firms to optimize the allocation of their skilled personnel to client projects while at the same time providing valuable insights into the impact of various workforce planning policies commonly employed by such firms during the workforce planning process. In addition, and importantly, we believe the model serves to define a framework for collecting and formatting the necessary data. In other words, the process of building the model and identifying the key input parameters and how they should be defined helps the firm to understand the types of data that are required to support the workforce planning process in this sector. This is a necessary first step that must occur before the firm begins the data collection process.

III. MODEL DEVELOPMENT AND FORMULATION

In order to capture the real-world issues that professional service firms face in the workforce planning process, exploratory case study interviews were conducted with a number of these firms. The exploratory case study approach is a well-established strategy used by empirical researchers [26], [27]. Essentially, if the research study is based on "what" questions, requiring a description of the prevalence of certain outcomes, then some form of exploratory research (e.g., surveys, experiments or case studies) should be undertaken. The questions in our research examine the "how" and the "why" with regards to workforce planning issues, are more open ended, and therefore more explanatory in nature, which suggests the use of case studies as being applicable [27]. When conducting exploratory case studies, it is suggested that the researcher inquires with hunches or conjectures derived from theory, observation or experience [28]. For this research, we derived these from the literature and from observations of how these firms operate. Firms were selected for their exploratory value. In such instances, the question of representativeness of the cases need not be of major concern [28]. It is recommended that theoretically useful cases are selected, rather than employing random sampling [26].

A total of seven firms were investigated. Each provides a broad range of services within the engineering design, consulting and technical services sectors. A semistructured interview protocol was developed and covered topics such as the description of the pool of employees, the nature of customer demand and the types of operational constraints that firms face. In order to ensure consistency between cases, each interviewee was a member of the senior management team in their firm, reporting directly to the chief executive officer (or equivalent title) at that location. This was to ensure a consistent perception of issues across the cases, as it has been shown that perceptions can differ between levels in the organization [29]. Each interview was recorded electronically and was subsequently transcribed and analyzed after each interview. This allowed the researchers in subsequent interviews to probe observations that emerged during prior interviews. The information from the interviews was then compiled and used as inputs to the development of the model's objective function, constraints and parameters.

A. Characteristics of Professional Service Firms

Customer demand in professional engineering service organizations can be represented in terms of projects. Therefore, the workforce planning process in such firms can be described as the matching of the employees of the firm (supply) to the customer projects (demand), subject to certain operating rules and constraints, with the goal of maximizing the firm's net revenue. In professional services, an important feature is that the customer often contributes to the cocreation of the service (e.g., clients have to perform acceptance testing for IT solutions). However, this topic is outside the scope of this paper.

The challenges in developing a model of the workforce planning process include generating an accurate profile of the customer demand (projects) and defining the relevant attributes of the service personnel (employees). During the case study interviews, several distinguishing characteristics of professional engineering services firms emerged. Among the medium to larger firms, which were the focus of this study, employees are divided into different grades or seniority in the organizational hierarchy. There are generally about six levels of seniority, from entry level "Analysts" to senior level "Managing Partners" (the actual labels varied by firm).

The number of projects in the firms' portfolios at any point in time was generally between thirty and seventy. Project duration was typically between three and six months and was often determined by the nature of the firm. The engineering design firms' projects tend to last for three to six months, but the firms could undertake projects of up to a year or longer (such as civil engineering "design and build" projects). In IT services, implementation projects (e.g., IT system deployment) averaged six months in duration. In addition to operating in multiproject environments, where firms execute many projects simultaneously for different customers, professional services firms operate in multilocation environments, with either office locations or customers based across different geographical locations. Employees are expected to agree to be assigned in locations other than their "home" location.

All firms interviewed provided services to multiple sectors and had different lines of business within each sector, so the first attributes of a project to be specified are the industry sector (e.g., power and utility, environmental management, IT) and the line of business (e.g., process reengineering, technology deployment). The next attribute to be specified is the time frame for the project. In a multiproject environment, projects start in different time periods and differ in length. For the purpose of the model development, time periods are assumed to be expressed in weeks, with each weekly time period containing a fixed number of working hours (40 hours). Using the project scope and work breakdown structure, which is usually developed through negotiations between the customer and the firm, it is possible to identify, for every time period of the project, the types of skills that are required to complete each task and the total number of hours for which those skills are required. An underlying assumption is that this data can be obtained or accurately predicted. However, in certain industries, such as software development, where agile project management practices are more common, this may prove difficult as the project deliverable is not clearly defined at the start of the project. A final project attribute that may be relevant is the seniority of the personnel that a customer may request for a project or which may be required for individual tasks within a project.

The attributes of an employee are similar to those identified for a customer project, i.e., an employee is located in the hierarchy of the firm on the basis of their industry sector and line of business expertise, as well as their seniority in the organizational structure (e.g., employee A is at the grade of project manager and is a technology deployment specialist). An employee also has specific skills (e.g., Java programming), which can change over time due to learning (e.g., using those skills on a project assignment or learning new skills through direct training). In addition, each employee has a certain availability of those skills in each time period. Employee availability and utilization (the percentage of available time spent on billable customer projects) varies between grades within a firm. For instance, an entry level engineer may be expected to spend around 80% of available time working directly on customer projects. However, employees at the most senior grades generally only spend 20% of their time directly on customer projects, with the remainder spent on business development.

The nature of professional engineering services organizations is to receive fees from customers for projects executed, while incurring costs as the firm progresses through the life cycle of each project. The contractual arrangement between the client and service provider can take a number of forms, affecting the way in which the client pays the service firm for the project. We consider two types of contracts: 1) fixed price; and 2) billable hour contracts. A firm incurs internal costs as it executes customer projects. It was found that firms tend to incur costs under two main categories: what they refer to as "project overhead cost" and "project assignment cost." For financial reporting and analysis purposes, it was indicated that firms prefer to keep these two costs separate. Project overhead cost includes internal business costs such as a portion of salary and benefits, and a share of general business expenses that can be allocated directly to project work. This cost is assigned to each grade of employee on a per hour basis (other general overhead costs incurred by an employee irrespective of project assignment are not included in

TABLE I
MODEL NOTATION

	Input Parameters
C	set of employees i in the firm, $i \in C$
D	set of customer projects p, $p \in D$
S	set of skills j possessed by employees and required by projects, $j \in S$
T	length of planning horizon, divided into t periods (weeks), $t = 1, \ldots, T$
PC_L, PC_U	lower, upper bound on the number of projects assigned to an individual employee in a given period
CP_L, CP_U	lower, upper bound on the number of different employees assigned to a project in a given period
DC_L, DC_U	lower, upper bound on the number of different employees assigned to a project over its duration
PE_L, PE_U	lower, upper bound on the number of projects that can be executed in their entirety over the planning horizon
L_G, U_G	lower, upper bound on gap between employees grade and project grade
Employee Attributes	
OC_{it}	hourly overhead cost of employee i in period t
A_{it}	total number of hours employee i is available to work in period t
CS_i	industry sector in which employee i is based
CB_i	line of business for employee i
CG_{it}	grade of employee i in period t
	Note: CS_i, CB_i, CG_{it} are input in numerical form, e.g., $CS_1 = 5$ may indicate that employee 1 represents the pharmaceutical division
S_{ijt}	1 if employee i has skill j in period t, 0 otherwise
Project Attributes	
R_p	total revenue received from project p (under a fixed price contract)
R_{jpt}	hourly revenue received from project p in period t when an employee is assigned to perform skill j (under a billable hour contract)
AC_{ipt}	assignment cost incurred when employee i is assigned to project p in period t
H_{jpt}	number of hours of skill j required for project p in period t
PS_p	industry sector in which project p is based
PB_p	line of business associated with project p
PG_{pt}	grade or level of project p in period t
Decision Variables	
x_{ipt}	1 if employee i is assigned to project p in period t, 0 otherwise
z_{ip}	1 if employee i is ever assigned to project p, 0 otherwise
w_p	1 if project p is fully completed, 0 otherwise
y_{ijpt}	number of hours employee i uses skill j on project p in period t

the model). For some firms, project overhead cost can form the basis of an employee's performance metric, i.e., how quickly in a particular time period an employee covers their overhead cost. Project assignment costs, on the other hand are not hourly costs but rather external client project-specific costs and are only incurred when an employee is assigned to a particular project. Travel costs to a customer site or project-specific training costs would be the principle items in this cost category.

B. Model Formulation

A mixed integer linear programming model of the workforce planning process for professional services organizations was developed, with the objective of optimally matching employees to client projects in order to maximize the profitability of the firm. The workforce planning problem is an example of the assignment problem (AP) [30]. Applications of the AP include assigning students to project groups [31], the conference reviewer/paper problem [32], assigning surgical cases to hospital operating rooms [33] and grid computing [34]. Integer programming is the most common technique employed for the type of problem outlined in this paper [35].

Table I summarizes the notation used in the model. The model assumes that a firm has a set of employees and a set of customer contracts, in the form of projects to complete. The decision to be made is how to optimally assign the resources of the firm, i.e., the employees, in order to maximize the net revenue earned by the firm across the complete portfolio of projects. Projects that

are approved for execution must have the necessary resources in terms of employee skills and availability to ensure completion (i.e., projects that are started must be completed). We will initially formulate the model for a single stage consisting of multiple periods (e.g., one quarter consisting of 13 weeks). However, in Section V, we will demonstrate how we can extend the model to consider multiple stages using a rolling horizon approach.

(1)–(16) present the model formulation. In the objective function, the first two terms represent the revenue received from customer projects, while the next two terms represent the costs the firm incurs while completing these projects. The first term shows the fixed price contract revenue, R_p, which is multiplied by the w_p decision variable to account for the fact that this revenue is only received if a project is fully completed. The second term captures the case in which the customer pays the service firm based on the number of employee hours that are used on the project. Hourly revenue is defined by skill type, j, project, p, and time period, t. This input parameter is multiplied by the decision variable representing the hours worked by a given employee, using a particular skill, on a given project in a given time period, y_{ijpt}, and summed to yield the total revenue received. Depending on the nature of the project contract, generally either the first or second term in the objective function will be relevant. The final two terms in (1) represent the costs the firm incurs. The third term represents the assignment cost, while the final term represents the overhead costs incurred by the project on an

hourly basis.

$$\max \sum_{p \in D} R_p \, w_p + \sum_{i \in C} \sum_{j \in S} \sum_{p \in D} \sum_{t=1}^{T} R_{ijpt} \, y_{ijpt}$$

$$- \sum_{i \in C} \sum_{p \in D} \sum_{t=1}^{T} \mathrm{AC}_{ipt} \, x_{ipt} - \sum_{i \in C} \sum_{j \in S} \sum_{p \in D} \sum_{t=1}^{T} \mathrm{OC}_{it} \, y_{ijpt} \qquad (1)$$

$$\text{s.t.} \quad \mathrm{CP}_L \le \sum_{i \in C} x_{ipt} \le \mathrm{CP}_U \qquad \forall p \in D, \ t = 1, \ldots, T \qquad (2)$$

$$\mathrm{PC}_L \le \sum_{p \in D} x_{ipt} \le \mathrm{PC}_U \quad \forall i \in C, \ t = 1, \ldots, T \qquad (3)$$

$$\mathrm{DC}_L \le \sum_{i \in C} z_{ip} \le \mathrm{DC}_U \quad \forall p \in D \qquad (4)$$

$$\mathrm{PE}_L \le \sum_{p \in D} w_p \le \mathrm{PE}_U \qquad (5)$$

$$\sum_{i \in C} y_{ijpt} = H_{jpt} \, w_p \quad \forall p \in D, j \in S, \ t = 1, \ldots, T \qquad (6)$$

$$\sum_{p \in D} \sum_{j \in S} y_{ijpt} \le A_{it} \quad \forall i \in C, \ t = 1, \ldots, T \qquad (7)$$

$$y_{ijpt} \le S_{ijt} \, A_{it} \quad \forall i \in C, \ j \in S, \ p \in D, \ t = 1, \ldots, T \qquad (8)$$

$$(\mathrm{CS}_i - \mathrm{PS}_p) \, z_{ip} = 0 \quad \forall i \in C, \ p \in D, \ CS_i, PS_p \in \mathbb{Z}^+ \qquad (9)$$

$$(\mathrm{CB}_i - \mathrm{PB}_p) \, z_{ip} = 0 \quad \forall i \in C, \ p \in D, \ CB_i, PB_p \in \mathbb{Z}^+ \qquad (10)$$

$$(\mathrm{CG}_{it} - \mathrm{PG}_{pt}) \, z_{ip} = 0 \quad \forall i \in C,$$

$$p \in D, \ CG_{it}, PB_{pt} \in \mathbb{Z}^+, \ t = 1, \ldots, T \qquad (11)$$

$$x_{ipt} \le z_{ip} \quad \forall i \in C, \ p \in D, \ t = 1, \ldots, T \qquad (12)$$

$$y_{ijpt} \le x_{ipt} \, A_{it} \quad \forall i \in C, \ j \in S, \ p \in D, \ t = 1, \ldots, T \qquad (13)$$

$$\sum_{j \in S} y_{ijpt} \ge x_{ipt} \quad \forall i \in C, \ p \in D, \ t = 1, \ldots, T \qquad (14)$$

$$z_{ip} \le w_p \quad \forall i \in C, \ p \in D \qquad (15)$$

$$y_{ijpt} \in \{0 \cup [4, \infty)\}, x_{ipt} \in \{0, 1\}, \ z_{ip} \in \{0, 1\},$$

$$w_p \in \{0, 1\}, \forall i \in C, \ j \in S, \ p \in D, \ t = 1, .., T. \qquad (16)$$

The objective function (1) focuses solely on maximizing net revenue, as our interviewees indicated that this was their firm's primary goal. However, it is possible to consider a multicriteria objective function which would be a weighted average of multiple objectives, including net revenue, employee utilization and skill usage. In addition, the issue of project priority can be captured in the objective function by assigning a priority weighting to the project revenue term (R_p), thereby making projects with a higher priority more financially attractive to complete. This approach can also be used to reflect the uncertainty associated with projects at the tender/request for quotation/precontract negotiation stages. Some of these projects may become active in the future and require employees to be assigned. Thus, priorities can be assigned to the project revenue term based on the probability of a firm signing a project contract.

Constraint (2) limits the number of employees assigned to a given project in each time period. The firm may specify a lower bound, CP_L (e.g., $\mathrm{CP}_L = 1$ would ensure that at least one employee is assigned to a project in each period) and an upper

bound, CP_U. The upper bound will be influenced by the total number of skills required during the project.

Constraint (3) limits the number of different projects that can be assigned to an individual employee in a given time period. A lower bound of $\mathrm{PC}_L = 1$ would indicate that at least one project must be assigned to each employee in each period. However, to prevent infeasibility when there is excess employee capacity compared to project demand, the lower bound PC_L may be set equal to zero. Such a scenario can help identify where there is excess employee capacity in certain skills. The upper bound PC_U will depend on the firm's policy regarding the number of distinct projects an employee can be involved with at a time. Based on our interviews, we find that most firms prefer to limit each employee to one project at a time to ensure high levels of customer satisfaction. There are occasions, however, where employees may work on two projects simultaneously, but rarely three or more. The upper bounds in constraints (2) and (3) can be removed if the firm chooses not to limit the number of projects assigned to an employee. These upper bounds can have a significant impact on the project completion rate, as will be seen.

Constraint (4) sets bounds on the number of different employees that are permitted to work on any project over that project's life cycle. The lower bound DC_L of this constraint can be set greater than or equal to 1. It is often unrealistic to expect one employee to be able to execute a complete project in isolation. On the other hand, having a large number of employees assigned to complete a particular project may negatively affect the operation of the project. The goal of this constraint is to balance these concerns. To prevent a large number of employees being assigned to a project, the upper bound DC_U could be initially set equal to the number of different skills required for the project, which assumes that each employee possesses only one skill, and then to reduced the bound to the lowest value at which the project is completed.

While constraints (2) to (5) describe the nature of the firms that were the focus of our interviews, they can be eliminated if the policies they reflect overconstrain, or are not relevant to, the firm using this model for workforce planning.

As part of the project tendering/request for proposal stage, a firm must demonstrate that it has the resources to execute the project in its entirety. In other words, projects cannot commence and be partially completed, at least from a human resource perspective (there may be reasons that emerge during the execution of a project which prevent the project from being completed, e.g., due to material, technical or budget issues). Thus, it is assumed that a project is only awarded to the firm when it can confirm it has the workforce available for the duration of that project. With this in mind, the w_p decision variable determines whether or not a project is taken on by the firm.[1] A firm may wish to limit the number of projects it has in its portfolio, which is captured in constraint (5). The bounds specify the minimum and maximum number of active projects the firm can have in its portfolio. An alternative approach is to set no upper bound. In this way, the firm could to determine the maximum number

---

[1]Model presented in this section is static. In reality, the workforce planning process is dynamic, across a rolling time horizon. We extend the model to account for this in Section V.

of projects that could be completed given a particular supply of resources and customer demand profile.

Constraint (6) uses the w_p decision variable to ensure that the number of employee skill hours assigned to a project during each time period (y_{ijpt}) matches the hours of skills required (H_{jpt}) on that project during each time period. If w_p equals 1, then it is possible to execute a project in its entirety and the total employee skill hours assigned in each time period must meet the project skills and time requirements for each period.

Constraint (7) states that the total number of hours an employee works, across all projects, in a given period, cannot exceed the total number of working hours that employee is available.

Constraint (8) is designed to ensure that an employee who is assigned to use a certain skill on a project in a given period actually possesses that skill in the time period in question. Constraints (9) and (10) ensure that employees with particular industry sector and line of business experience are matched to projects in their own industry sector and line of business. We assume the firm has a well-defined set of industry sectors (e.g., power and utility, environmental management, IT). For the purposes of preparing the input data, these sectors need to be codified and labeled numerically. For example, we might have power and utility = 1, environmental management = 2, IT = 3. Then, if employee i specializes in environmental management, $CS_i = 2$. Similarly, if project p is in the environmental management domain, $PS_p = 2$. The inputs CB_i, PB_p, CG_{it} and PG_{pt} are similarly defined and coded.

For firms that have a number of levels of seniority, it may be desirable to match the seniority (often referred to as the grade) of the employee with what has been identified as the grade of the project. This may prevent situations in which very senior employees are assigned to relatively simple projects. This requirement is reflected in constraint (11). While not shown here, both the employee grade and the project grade can change over time. A change in employee grade reflects the reality of promotions and long-term service. A change in the project grade could occur, for example, if a more senior employee is required for a particular phase of the project. On the other hand, companies may be forced to relax this constraint if there is not a suitably qualified person available. For example, a firm may need to assign an over qualified individual. To allow this, the firm could change the "=" in (11) to "≥." If the project and employee grades, PC_i and CG_i, are defined and ordered numerically, with their values increasing with grade (e.g., an entry level analyst has grade = 1, while a managing partner has grade = 6), this modification of (11) would allow a higher grade employee to be assigned to a project requiring a lower grade.

Constraints (12)–(15) provide the linkages between the decision variables. Constraint (12) is also required for constraints (9)–(11), i.e., to ensure consistency in the values of x_{ipt} and z_{ip}. Constraints (16) define the set of allowable values for each of the decision variables. The y_{ijpt} decision variable is classified as a semicontinuous decision variable, the value of which can be zero (if no employee is assigned to a particular project in a time period), but whose minimum value if an employee is assigned

to a project must be four hours. The software we used for our computational study (FICO XpressMP) allows for the explicit definition of semicontinuous variables. Based on our interviews with professional service firms, if an employee is assigned to a project in any time period, it must be for at least four hours (or half a working day).

Constraints (2) to (16) are specified as "hard" constraints, which could make it difficult to find the perfect assignment of employees to projects. To overcome this, it is possible to make some of the constraints softer, by reducing the number of different categories within the classifications of levels of skill, hierarchical grades, business domain and geographical location (e.g., three levels of skill instead of five or clustering locations by region rather than city).

IV. ANALYSIS AND RESULTS

The model presented above can be used as a starting point in the development of a tool to assist professional services firms in work force planning. We are interested in how the workforce planning process can be modeled and what insights such a model can provide into workforce planning policies. The model can also be used to identify shortages of skills which prevent projects from being completed, as well as excesses of skills that are not being fully utilized. This information can be used to improve training, hiring and project selection decisions. The development of this planning model was informed by the interviews conducted with a number of professional services firms. During the course of these interviews, indicative upper and lower bounds for various input parameters were identified. This data were used to generate a number of "test firms" which were used to develop insights into the impact of various workforce planning policies on firm performance. We first describe the methodology used to generate the input data for these test firms. We then discuss the various workforce planning policies that we examined and describe their impact on firm performance.

A. Development of Test Firms

Based on the interviews with firms, we were able to develop the characteristics by which employees, customer demand and operational constraints could be described. These characteristics were then used as the inputs to the development of the model's objective function, constraints and parameters. The firms indicated that that the key characteristics used to describe employees were business domain, hierarchical grade, skills, cost and capacity. A project was described based on its business domain, start time, total duration, hours of each skill required per time period and revenue. Typical operational constraints included business domain matching between project and employee, skill matching, capacity limit of each employee (e.g., hours per week), maximum number of projects per employee and maximum number of employees per project. During the interviews, the firms provided indicative data (sample lower and upper bound information only) for employee, project and constraint characteristics. We then used this data to generate the profiles of four test firms used to demonstrate the effectiveness of our model. The nature of these profiles is described below. The firms differed in the total number of potential projects, the typical duration of the

---
[2]In reality, this matching may also be performed at the skill level.

projects, and the number of skills required per project. Each firm consisted of 60 employees, spread across four grades. For each of the firms, six different employee skill scenarios were considered, resulting in a total of 24 firm-skill scenario combinations. The planning horizon for each firm was six months. An approach that allows for longer planning horizons will be presented in Section V. The size of the four test firms is representative of medium-sized organizations in the sector, or of a typical business unit which would conduct its own workforce planning in a larger organization. Tables V–VIII in Appendix provide detailed descriptions of these test firms and the skill scenarios.

1) Employee Profile: The distribution of employees across different grades in the firm's hierarchy is not even. The firms interviewed suggested that there are fewer employees at the bottom and top levels in the organizational hierarchy, with a peak in the middle grades. Thus, a multinomial distribution was used to assign the employees across four grades. The characteristics of this distribution are as follows: grade one (lowest grade) has $p_1 = 0.2$, grade two has $p_2 = 0.35$, grade three has $p_3 = 0.35$, and grade four (highest grade) has $p_4 = 0.1$. See Table V.

2) Employee Skill Scenarios: Every employee is characterized by the number and types of skills possessed. For the test firms, we used a total of thirty different skills. These skills were divided into subsets associated with each grade, in proportion to the number of employees in each grade, so that more skills are associated with middle grades. See Table VI. Six different skill scenarios were created. In scenarios 1, 2 and 3, each employee possesses one, two and three skills, respectively, from within their current grade. In scenario 4, employees carry over skills from their earlier grades, e.g., grade two employees will possess their grade one skills in addition to their grade two skills. In each case, the skills associated with each employee were randomly selected from the skills associated with their grade. See Table VII.

In the first four skill scenarios, every employee has the same level of proficiency across all of his/her skills. In reality, employees have "primary" skills for which they are deemed to be highly proficient and "secondary" skills in which they are less proficient. To capture this behavior, we assigned "capability values" to workers to reflect different proficiency levels in skills [36]. Capability values close to 0 indicate very poor capability in that skill, values close to 1 represent normal capability, and values greater than 1 reflect unusually high capability. In skill scenario 5, the skills generated in scenario 4 are assigned a capability value of 1.0 for the current skill and a capability value of less than 1.0 for older skills. An example of this scenario would be an employee who is hired at grade one as a software engineer with a primary skill of Java programming, for assignment to technology deployment projects. This individual's role is to write software programs on a daily basis. Over time, the individual is promoted to grade three, where they assume a project management role. While they never completely forget how to write code, they become less proficient at that skill. Finally, in scenario 6, scenario 5 is extended to allow some employees to be "highly proficient" at their current skill, by assigning a skill capability value of up to 1.2 to their current skill. For scenarios 5 and 6, constraint (7) can be adjusted to incorporate this skill capability value as a rating factor (similar to the approach

used in time studies), whereby the number of hours required to complete a task by a highly proficient employee is reduced.

3) Number of Projects: The case study interviews indicated that the number of projects in a firm's portfolio can vary between thirty and seventy at any point in time. Bearing in mind that the firms in question would be classified as large multinational, for the creation of the test firms we chose to create medium-sized firms, with ten, twenty and thirty live projects in their respective portfolios. Firms One and Two each had twenty average duration projects, Firm Three had thirty short projects and Firm Four had ten long, skill intensive projects. A summary of each firm's project and skill parameters is contained in Table VIII.

4) Project Duration: For our test firms, all projects have a duration between three and six months (between twelve and twenty six weeks). The length of each project was drawn randomly from an appropriate Beta distribution, which is commonly used in project management to model task duration [7]. The start date for each project was drawn from a uniform distribution. In Firm One, which has twenty projects, the start dates for projects 1–10 (projects 11–20) were generated from a uniform [1,12] (a uniform [8,20]) distribution.

5) Project Skill Requirements: A uniform distribution was used to determine how many skills were required for each project in the portfolio according to the ranges shown in Table VIII, e.g., for Firm One, a uniform [4,7] distribution was used to determine the number of skills per project. To determine which specific skills were required for each project, we assumed that firms seek to win business from customers that are closely aligned with the capabilities of the firm. In other words, we assumed that the profile of the skills in each set of projects would be similar to the profile of the skills across the pool of employees. For example, 20% of the employees are in grade one, which has skills [1–6]. If a project required ten skills, then 20% of those skills were randomly drawn from the set of grade one skills.

6) Project Skill Hours: Each skill is generally not required for the entire duration of a project. An individual "project life cycle" is comprised of a number of phases, with low resource requirements at the start of a project (the "definition phase"), peak resource requirements at the execution phase, and decreasing resource requirements during the closure/sign-off phase [7]. This observation was used to determine which specific skills were be required in each time period over the life cycle of every project. The next step was to determine how many hours of each skill were needed. While employees are available for 40 hours per time period (i.e., per week), projects may not require 40 hours of a given skill in each time period. For our test firms, skills hours were randomly assigned in time buckets of eight hours.

B. Analysis of Workforce Planning Policies

Given these test firms and skill scenarios, a set of experiments was conducted with the goal of obtaining practical insights into the impact of common workforce planning policies on various business metrics. These metrics include project completion rates, firm profitability, employee utilization, and shortages and excesses of skills (which we will refer to as "gaps" and "gluts").

1) Employee Skills: As expected, we observed an increase in the project completion rate as the number of skills possessed by

TABLE II
PROJECTS COMPLETED IN SKILL SCENARIOS 1–6

	Firm One	Firm Two	Firm Three	Firm Four
Skill Scenario 1	12	14	15	5
Skill Scenario 2	17	16	20	7
Skill Scenario 3	20	19	25	9
Skill Scenario 4	18	18	21	7
Skill Scenario 5	12	2	8	6
Skill Scenario 6	18	14	22	7
Maximum Possible	20	20	30	10

each employee increased (see Table II). The number of projects completed increased by about 50% between skill scenario 1 and skill scenario 3. In scenario 5, i.e., when employees acquire new skills as they gain seniority, but lose proficiency in previous skills, it was found that the decline in proficiency had a significant negative impact on the number of projects that could be completed, and thus on the firms' revenue. This finding presents challenges for firms, who seek to use promotion as a strategy to motivate and retain staff, and who generally try to assign employees to projects that use their current or most senior skills. Successful firms must put into place the necessary recruitment, promotion, training and monitoring mechanisms to ensure that employees maintain multiple skills.

2) Cross-Training: We also considered the impact of cross-training on the project completion rate and net revenue. Under cross-training, employees gain new skills over time [37]–[39]. Firms will often use cross-training, which ensures that employees have multiple skills, to provide the flexibility necessary to cope with uncertainty, which can occur in project opportunities, project requirements and employee availability. Optimization models for workforce training have been developed previously [40], [41]. Four static optimization models that focus on training to support production planning based on a forecast of future skills needs were developed [40]. An optimization model was developed for a pool of construction workers that seeks to minimize the cost of hiring, switching to another activity (assuming the workers possess multiple skills) and firing [41].

In contrast, we focus on cross-training as an approach to improve flexibility. Our approach was motivated by cross-training in serial production lines, where it has been found that targeted cross-training at identified bottlenecks leads to superior performance than cross-training every worker on the line in every skill [42]. Using this idea, rather than cross-training every employee in every skill, targeted cross-training was carried out in specific skills for which shortages were preventing projects from being completed. As part of this process, the model described in Section III-B was used to determine where there were shortages of skills ("gaps") which prevented a firm from executing a project and where there was an excess of skills ("gluts"), which resulted in employees having low utilization. The employees with low utilization were identified, along with their corresponding skills. The list of projects that could not be executed was analyzed to determine the skills that were required but not currently available. Targeted cross-training of the low utilization employees was performed to address these "gaps." We find that the marginal benefit of cross-training single skilled

employees in a second skill was greater than marginal benefit of cross-training two-skilled employees in a third skill. This is similar to [36], [43] and [41], who found a diminishing positive effect of expanding the levels of labor flexibility. In general, we find that it is possible for firms to achieve their maximum project completion rates through targeted cross-training of a portion of the workforce. The impact of the cross-training lead time, and reduced employee availability during this lead time, was also analyzed. In order to prevent any reduction in the number of projects completed, the training lead time must be very short, i.e., two weeks or less for our test firms. If this lead time becomes too long, there can be a significant impact on project completion rates.

3) Minimum Skill Usage: When employees possess multiple skills, it may be desirable, for skill proficiency and employee motivation reasons, to establish a minimum number of hours for which each skill must be used over a particular planning horizon. If employees do not use their skills on a regular basis, they run the risk of becoming less proficient in those skills. In order to capture this issue, an optional constraint can be added to the model stating that for each employee the total number of hours in which each skill is used, over a particular time horizon, must be at least ρ. We envision that such a constraint would be added only in cases where maintaining skill proficiency is of critical importance to the firm. The value of ρ may be dependent on the specific skill, with some (or even most) skills having $\rho = 0$. In other words, a positive value of ρ may only be assigned to those skills which experience significant diminished proficiency due to lack of use. Or, ρ could be dependent on the specific skill, with some skills having $\rho_j = 0$. For our experiments, we found that the firms with a project portfolio containing longer, more skill intensive, projects were better able to satisfy this constraint without negative impacts on the performance metrics.

4) Revenue Versus Utilization: From our interviews, we found that resource allocation managers in firms often perform project assignment with the short-term goal of maximizing employee utilization. Thus, in addition to solving the net revenue maximization problem, the model formulation was modified to consider utilization maximization. The overall net revenue earned by the firm under the two approaches were compared in order to understand the magnitude of the potential lost revenue when resource allocation is performed under the short term objective of "keeping everyone as busy as possible." Workforce planning managers prefer to keep employees busy because this can reduce attrition. Attrition rates are an important factor to consider in long-term workforce planning, since loss of skill workers can negatively impact net revenue. Across the test firms and scenarios, there was a consistent 5% to 10% reduction in net revenue when workforce planning was done under the maximize utilization objective. The implication is that the policy of keeping all employees as busy as possible can have a significant detrimental impact on the profits of the firm. Thus, it may be preferable to incorporate both revenue and utilization into a weighted objective function.

5) Employee Hiring: Our interviews indicated that the hiring of new employees occurs for two main reasons: to replace skilled
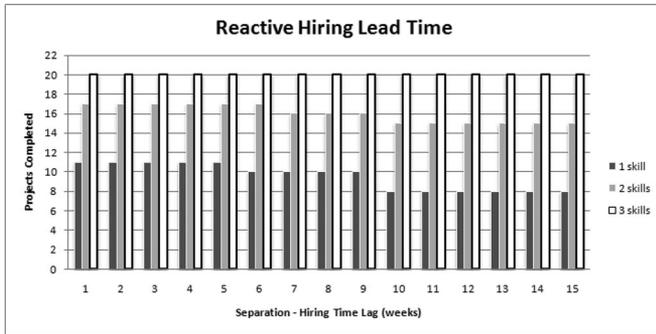
Fig. 1.   Impact of the reactive hiring lead time.

TABLE III
NUMBER OF PROJECTS COMPLETED (FIRM ONE)

Skill Scenario	UB on Constraint (3)		
	≤ 1	≤ 2	≤ 3
1	9	13	13
2	10	13	13
3	14	14	18

employees lost through voluntary separation for a variety of reasons [44] and to increase the capacity of the firm when there is a forecast increase in demand. Two main approaches to hiring were discussed. The first approach is to be proactive, hiring additional staff in anticipation of a certain level of voluntary separation. In this case, firms will estimate an average separation rate on an annual basis and hire new staff in anticipation of this separation. However, predicting the actual staff members (and skills) that will be lost is difficult. Thus, firms may end up with a mismatch between the skills lost and the skills hired. The second option is for the firm to hire only when the increase in demand has occurred or when a specific employee has left the firm. The advantage of this reactive approach is that the firm knows the specific skills required and can hire accordingly. However, there is often a lead time (ten weeks on average [45]) for hiring new employees.

We find that the reactive policy will only be preferred when the hiring lead time is sufficiently short (e.g., less that six weeks). If the firm can replace a departing staff member relatively quickly with a similarly skilled new staff member, there is minimal impact on business metrics. The challenge for the firm is to be able to reduce the hiring lead time. Fig. 1 demonstrates the impact of the hiring lead time under reactive hiring. The figure considers Firm One and assumes a 10% separation rate of employees at each grade during the planning horizon. Departing employees were randomly selected and the timing of each departure, i.e., the week in which the employee left the firm, was chosen at random. We considered employee skill scenarios 1 through 3, and varied the hiring lead time from one to fifteen weeks. We can see that reducing the hiring lead time is most effective for firms whose employees possess a limited number of skills.

6) Number of Projects per Employee: The case study interviews indicated that the preferred practice is to set the upper bound in constraint (3) equal to one, especially for employees in lower grades. In other words, each employee is allowed to work on one project in any time period. As expected, as shown in Table III, this policy leads to reduced employee utilization and fewer projects completed. For skill scenarios 1, 2 and 3, experiments consistently indicated that increasing the upper bound on constraint (3) from one to two, or in some cases to three, leads to the highest project completion rate, and that increasing the upper bound beyond three does not improve the project completion rate. Thus, there is a conflict between the firms'

practical policy, which is motivated by the desire to be able to dedicate employees to customers in each time period, in order to "not spread resources too thinly," and the net revenue maximization objective, which pushes firms to be more flexible in their employee assignments.

7) Project Uncertainty: The operations literature identifies two general approaches for coping with uncertainty: flexibility and redundancy. Firms will often use cross-training, which ensures that employees have multiple skills, in order to provide the flexibility necessary to cope with uncertainty in project opportunities, in project requirements and in employee availability. The other approach to cope with uncertainty is redundancy. In a professional services context, a firm may maintain excess capacity in order to ensure that projects can be completed even if the project requirements change or the available resources change. Such excess capacity can be ensured by scaling down the availability of some or all employees in the availability constraint (7). For example, the firm may initially assume that each employee is available for 35 hours per week. However, if needed, the firm can use each employee for 40 or more hours per week.

V. ROLLING TIME HORIZON IMPLEMENTATION

Workforce planning in most professional services firms is a dynamic process, with each project having a different status. Projects may be at the tendering/request for quotation stage, have a contract signed but not yet commenced, or be underway, delayed, on-hold and canceled. The model described in Section III is static and assumes only new projects. To overcome this limitation and to capture the dynamic nature of the project life cycle, we incorporate a rolling time horizon perspective to the workforce planning process for an extended planning horizon.

Through the interviews conducted as part of this research, we found that while the base unit of time for projects is generally measured in weeks, aggregate planning tends to occur on a rolling quarterly basis, with each planning cycle looking at least two quarters into the future (Fig. 2). This means that firms are always planning two quarters into the future, but update the plan at the start of every new quarter. By updating the plan at the start of every quarter, firms can capture the different elements of the project life cycle, such as continuing projects (which need employees who were already assigned to these projects in an earlier quarter to be prioritized for these projects if their skills are required), winning and starting new projects (available employees are assigned to these), projects being put on hold (therefore requiring no hours for a number of time periods), projects canceled (no hours required beyond a certain time period, so employees assigned are now available for other
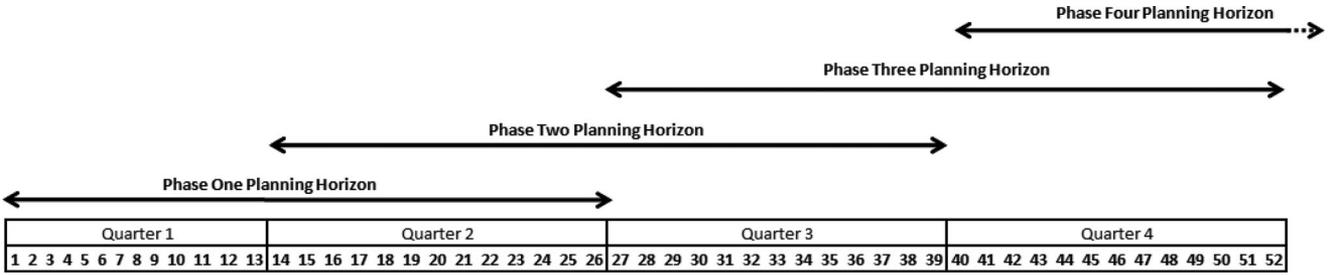
Fig. 2. Extended planning horizons.

assignments). In addition, long customer projects can overlap several of these planning phases. This process also reflects the reality that firms cannot complete their workforce planning for a twelve month time horizon with absolute certainty. Unforeseen changes are common, e.g., clients may change the scope of their projects or new projects may appear suddenly.

For the purpose of project team efficiency and customer satisfaction, it is desirable for projects that overlap two or more planning phases to have consistency in the employees assigned to the skills required on those projects. In other words, if an employee performs a particular skill on a particular project in one time period, and if that same skill is needed in the same project at a later time period, the same employee should be assigned, to the extent possible. Therefore, the proposed rolling horizon approach must be capable of capturing the assignment of employees to skills from earlier planning phases. This is particularly critical in a multiskilled employee firm, since the high number of employees possessing each skill type creates the potential for more employee switching between projects.

To achieve consistency in employee-skill assignments between planning phases, we modify the model developed in Section III-B. A binary decision variable, u_{ijp}, is introduced, with $u_{ijp} = 1$ if employee i uses skill j on project p, and $u_{ijp} = 0$ otherwise. This decision variable is linked to the other variables through the additional constraint $y_{ijpt} \leq u_{ijp} \times M$ for all i, j, p, t, where M is a large number. The u_{ijp} decision variable is used to keep track of which employees use which skills on which projects. The proposed rolling horizon approach can now be described.

1) *Phase One*: Solve the model for phase one of the planning horizon (quarters 1 and 2). From the output, determine the set $I(j,p)$, which is the set of employees that used skill j on project p in phase one, i.e., $I(j,p)$ is the set of all i that have $u_{ijp} = 1$, for a given j, p.

2) *Phase Two*: Solve the model for phase two of the planning horizon (quarters 2 and 3), with the following additional constraints:

$$\sum_t y_{ijpt} \geq \sum_t y_{i'jpt} \quad \forall \quad i \in I(j,p), i' \notin I(j,p), j, p.$$

(17)

Employee i, with $i \in I(j,p)$, is an employee that had $u_{ijp} = 1$ in phase one, i.e., used skill j on project p.

Employee i', with $i' \notin I(j,p)$, is an employee that had $u_{i'jp} = 0$ in phase one, i.e., did not use skill j on project p. The constraint thus says that if employee i' has positive hours assigned using skill j on project p in phase two, i.e., if $\sum_t y_{i'jpt} > 0$ in phase two, then employee i must use skill j on project p for at least as many hours. If no other employees use skill j on project p in the phase two, i.e., if $\sum_t y_{i'jpt} = 0 \quad \forall i' \notin I(j,p)$, then employee i may or may not use skill j on project p in phase two. Thus, the phase one and phase two planning horizon results are linked through the additional constraints and the definition of the set $I(j,p)$.

3) *Phase Three*: Using the output from phase two, determine the set $I(j,p)$ for phase two. Solve the model for phase three of the planning horizon (quarters 3 and 4) with the addition of constraint (17) to ensure consistency in assigning employees to skills within a project.

4) *Phase Four*: Repeat the process for phase four of the planning horizon.

With this process, a firm can use a rolling horizon approach to the workforce planning problem while ensuring consistency in the assignment of employees to projects across the time horizon. However, enforcing this consistency reduces flexibility in workforce planning since this consistency policy has the effect of (essentially) preallocating a worker to a project in a later planning phase. As a result of this reduced flexibility, the firm may see a decrease in the number of projects completed. This reduced flexibility is particularly likely to occur if an employees availability decreases from one period to the next. This effect can be mitigated by relaxing constraint (3), i.e., by increasing the number of simultaneous projects to which an employee can be assigned. Thus, when dealing with extended planning horizons and projects that span multiple phases, firms face a tradeoff. While there is a general preference that employees work on just one project at a time, in order to maintain consistency in skill assignments for projects spanning multiple planning phases, this requirement may need to be relaxed. Otherwise, the project completion rate, and the firm's profits, will be reduced. A firm can use the model to create an efficient frontier to demonstrate the tradeoff and then allow managers to decide where on that curve they want to operate. The firm can solve the model for various values of the upper bound in (3) and use the results to create a graph showing number of projects completed (or net revenue) versus the upper bound (or average number of projects per employee).
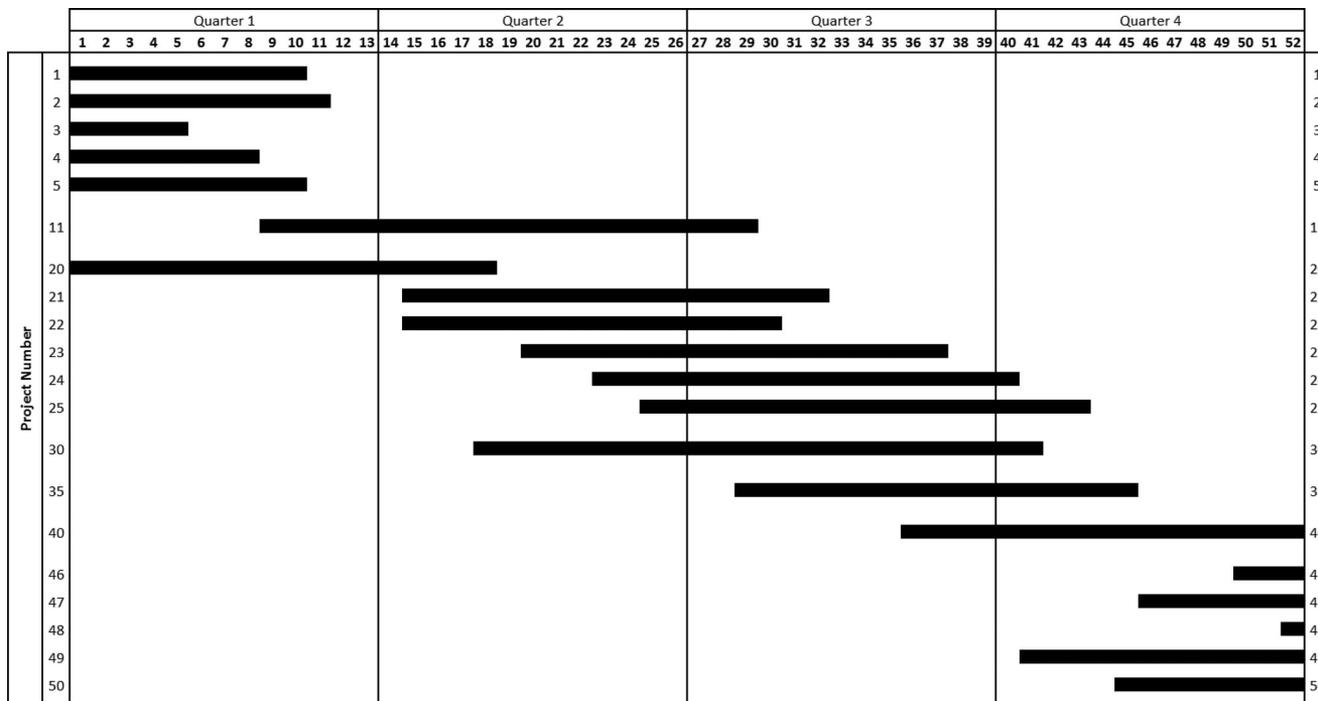
Fig. 3. Firm Five summary Gantt chart.

TABLE IV
FIRM FIVE - PROJECT 11 ASSIGNMENT - BEFORE AND AFTER CONSISTENCY CONSTRAINT

Skill Required	No Consistency Employee # (Phase One)	Employee # (Phase Two)	Consistency Applied Employee # (Phase One)	Employee # (Phase Two)
4	10	10	4, 10	4, 10
15	16, 17	33	28, 33	28, 33
16	16, 27	28, 33	15, 28	15, 28
24	48, 50	50, 51	50	50
29	60	60	60	60

To demonstrate the rolling horizon approach, a new model test firm (denoted Firm Five) was created which had a project portfolio spanning a full year. This planning horizon was divided into four quarters of thirteen weeks each. Firm Five has 50 projects in its portfolio over this one year planning horizon. In each quarter, ten new projects commence. The start date, duration and skill requirements for each project were determined using the process employed during the creation of Firms One through Four. Projects 1–10 represent those that are continuing from the previous planning cycle and are due for completion during quarter 1. Projects 41–50 represent those that commence during quarter 4 and continue into the subsequent planning cycle. A number of the projects that commence in quarter 3 also extend into the next planning cycle. A truncated Gantt chart for Firm Five is provided in Fig. 3 in Appendix.

We first solved Firm Five's problem without the constraints that ensure consistency in employee assignments across quarters. Table IV show the results for project 11, assuming skill

scenario 3 (i.e., assuming each employee possesses three skills). Project 11 takes place during phases 1 and 2 of the planning horizons. If these phases are managed separately, with no consideration given to consistency in employee assignments, we see significant switching, with different employees assigned to perform the same skill in different periods (columns two and three). Clearly, such switching could be quite disruptive to project team management, while also potentially harming customer satisfaction, particularly when employees work on site at the customer. We then resolved Firm Five's problem incorporating the constraints that ensure consistency in employee assignments to projects. The results for project 11 are shown in columns four and five of Table IV. As desired, consistency in employee-skill assignments has been achieved across the planning phases.

VI. MANAGERIAL IMPLICATIONS AND CONCLUSION

This research has focused on understanding the complex human resource workforce planning problem faced by professional services firms in which customer demand (projects) and the large number of labor attributes must be captured and modeled, including indicators of employee skills, knowledge and experience.

Through our model, engineering services firms can now begin to use analytical methods, rather than relying on intuition, for workforce planning [46]. Our research is beneficial to such firms since it demonstrates some of the critical steps that engineering managers must take to enable the application of analytical techniques to workforce planning, including: 1) developing a system to characterise employees and customer projects along

multiple dimensions, 2) articulating the operational constraints faced by the firm, and 3) collecting accurate data in the appropriate format. This research also demonstrates the implications of various work force planning policies, which can then be extended to rostering and scheduling [4]. Looking forward, to demonstrate the benefit of an analytical approach to workforce planning, firms can compare performance measures (e.g., employee utilization) under their previous intuition-based workforce planning approaches to those obtained when analytical approaches are applied. The long term implications of this research should be better decision-making, improved deployment of professional service employees, greater transparency [5] and an ability to carry out workforce planning across multiple planning horizons.

There are opportunities for further analysis in a number of areas which will provide additional insights into workforce planning in professional service firms. We would like to extend our work to consider policies regarding staff reduction and cross-training versus subcontracting. We would also like to investigate how demand management could improve project completion rates by adjusting the overall schedule of projects in the firm's portfolio, and to consider the impact of demand uncertainty over an extended planning horizon.

APPENDIX
DESCRIPTION OF TEST FIRMS

TABLE V
EMPLOYEE-GRADE DISTRIBUTION

Grade	Number of Employees
1	12
2	23
3	19
4	6
Total	60

TABLE VI
SKILLS ASSIGNED TO EMPLOYEE GRADES

Grade	Skills
1	1, 2, 3, 4, 5, 6
2	7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18
3	19, 20, 21, 22, 23, 24, 25, 26, 27
4	28, 29, 30

TABLE VII
SUMMARY OF SKILL SCENARIOS

Skill Scenario	Description
1	One skill per employee
2	Two skills per employee
3	Three skills per employee
4	Employees carry skills over from their earlier grades
5	As in Scenario 4, but a loss in capability of earlier skills
6	As in Scenario 5, but some employees highly capable in current skill

TABLE VIII
PROJECT AND SKILL PARAMETERS FOR TEST FIRMS

	Firm One	Firm Two	Firm Three	Firm Four
Total number of projects	20	20	30	10
Project duration (weeks)	8–16	8–16	7–12	20–24
Skills per project	4–7	5–8	4–6	8–12

REFERENCES

[1] P. Cappelli, "A supply chain approach to workforce planning," *Org. Dyn.*, vol. 38, no. 1, pp. 8–15, 2009.
[2] T. H. Davenport, J. Harris, and J. Shapiro, "Competing on talent analytics," *Harvard Bus. Rev.*, vol. 88, no. 10, pp. 52–58, Oct. 2010.
[3] B. Dietrich, "Resource planning for business services," *Commun. ACM*, vol. 49, no. 7, pp. 62–64, 2006.
[4] W. J. Abernathy, N. Baloff, J. C. Hershey, and S. Wandel, "A three-stage manpower planning and scheduling model—A service-sector example," *Oper. Res.*, vol. 21, no. 3, pp. 693–711, 1973.
[5] P. Drucker, "Knowledge-worker productivity: The biggest challenge," *California Manage. Rev.*, vol. 41, no. 2, pp. 79–107, 1999.
[6] R. Greenwood, T. Morris, S. Fairclough, and M. Boussebaa, "The organizational design of transnational professional service firms," *Org. Dyn.*, vol. 39, no. 2, pp. 173–183, 2010.
[7] C. Gray and E. Larson, *Project Management - The Managerial Process*. New York, NY, USA: McGraw Hill, 2008.
[8] M. Engwall and A. Jerbrant, "The resource allocation syndrome: The prime challenge of multi-project management?," *Int. J. Project Manage.*, vol. 21, no. 6, pp. 403–409, 2003.
[9] B. V. Dean, D. R. Denzler, and J. J. Watkins, "Multiproject staff scheduling with variable resource constraints," *IEEE Trans. Eng. Manage.*, vol. 39, no. 1, pp. 59–72, Feb. 1992.
[10] E. E. Lawler, "Make human capital a source of competitive advantage," *Org. Dyn.*, vol. 38, no. 1, pp. 1–7, 2009.
[11] D. L. Gresh, D. P. Connors, J. P. Fasano, and R. J. Wittrock, "Applying supply chain optimization techniques to workforce planning problems," *IBM J. Res. Develop.*, vol. 51, no. 3–4, pp. 251–261, 2007.
[12] H.-C. Huang, L.-H. Lee, H. Song, and B. T. Eck, "SimMan—A simulation model for workforce capacity planning," *Comput. Oper. Res.*, vol. 36, no. 8, pp. 2490–2497, 2009.
[13] H. Cao, J. Hu, C. Jiang, T. Kumar, T.-H. Li, Y. Liu, Y. Lu, S. Mahatma, A. Mojsilovic, M. Sharma, M. S. Squillante, and Y. Yu, "OnTheMark (OTM): Integrated stochastic resource planning of human capital supply chains," *Interfaces*, vol. 41, no. 5, pp. 414–435, 2011.
[14] C. Santos, T. Gonzalez, H. Li, K.-Y. Chen, D. Beyer, S. Biligi, Q. Feng, R. Kumar, S. Jain, R. Ramanujam, and A. Zhang, "HP enterprise services uses optimization for resource planning," *Interfaces*, vol. 43, no. 2, pp. 152–169, 2013.
[15] S. E. Bechtold and M. J. Showalter, "Methodology for labor scheduling in a service operating system," *Decision Sci.*, vol. 18, no. 1, pp. 89–108, 1987.
[16] A. Ernst, H. Jiang, M. Krishnamoorthy, and D. Sier, "Staff scheduling and rostering: A review of applications, methods and models," *Eur. J. Oper. Res.*, vol. 153, no. 1, pp. 3–27, 2004.
[17] A. J. Schaefer, E. L. Johnson, A. J. Kleywegt, and G. L. Nemhauser, "Airline crew scheduling under uncertainty," *Transport. Sci.*, vol. 39, no. 3, pp. 340–348, 2005.
[18] N. Gans, G. Koole, and A. Mandelbaum, "Telephone call centers: Tutorial, review and research prospects," *Manuf. Service Oper. Manage.*, vol. 5, no. 2, pp. 79–142, 2003.
[19] J. Van Den Bergh, J. Belin, P. De Bruecker, E. Demeulemeester, and L. De Boeck, "Personnel scheduling: A literature review," *Eur. J. Oper. Res.*, vol. 226, no. 3, pp. 367–385, 2013.
[20] P. De Bruecker, J. Van den Bergh, J. Belin, and Demeulemeester E. (2015). Workforce planning incorporating skills: State of the art [Online]. Available: http://ssrn.com/abstract=2437372
[21] B. Cheang, H. Li, A. Lim, and B. Rodrigues, "Nurse rostering problems - A bibliographic survey," *Eur. J. Oper. Res.*, vol. 151, no. 3, pp. 447–460, 2003.
[22] L. S. Franz and J. L. Miller, "Scheduling medical residents to rotations - Solving the large-scale multiperiod staff assignment problem," *Oper. Res.*, vol. 41, no. 2, pp. 269–279, 1993.

[23] S. Topaloglu, "A shift scheduling model for employees with different seniority levels and an application in healthcare," *Eur. J. Oper. Res.*, vol. 198, no. 3, pp. 943–957, 2009.

[24] B. Dodin and A. A. Elimam, "Audit scheduling with overlapping activities and sequence dependent setup costs," *Eur. J. Oper. Res.*, vol. 104, no. 1, pp. 262–264, 1998.

[25] L. L. Brennan, "Operations management for engineering consulting firms: A case study," *J. Manage. Eng.*, vol. 22, no. 3, pp. 98–107, 2006.

[26] K. M. Eisenhardt, "Building theory from case study research," *Academy Manage. Rev.*, vol. 14, no. 4, pp. 532–550, 1989.

[27] R. K. Yin, *Case Study Research: Design and Methods*. Newbury Park, CA, USA: Sage, 2002.

[28] W. K. Roche, "Selecting case studies in business research," in *Business Research Methods: Strategies, Techniques and Sources*. Dublin, Ireland: Oak Tree Press, 1997.

[29] P. M. Swamidass, "Manufacturing strategy: Its assessment and practice," *J. Oper. Manage.*, vol. 6, no. 4, pp. 471–484, 1986.

[30] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Res. Logistics Quarterly*, vol. 2, no. 1–2, pp. 83–97, 1955.

[31] K. R. Baker and S. G. Powell, "Methods for assigning students to groups: A study of alternative objective functions," *J. Oper. Res. Soc.*, vol. 53, no. 4, pp. 397–404, 2002.

[32] D. Hartvigsen, J. C. Wei, and R. Czuchlewski, "The conference paper-reviewer assignment problem," *Decision Sci.*, vol. 30, no. 3, pp. 865–876, 1999.

[33] H. Fei, C. Chu, N. Meskens, and A. Artiba, "Solving surgical cases assignment problem by a branch-and-price approach," *Int. J. Prod. Econ.*, vol. 112, no. 1, pp. 96–108, 2008.

[34] S. Kumar, K. Dutta, and V. Mookerjee, "Maximizing business value by optimal assignment of jobs to resources in grid computing," *Eur. J. Oper. Res.*, vol. 194, no. 3, pp. 856–872, 2009.

[35] A. T. Ernst, H. Jiang, M. Krishnamoorthy, B. Owens, and D. Sier, "An annotated bibliography of personnel scheduling and rostering," *Ann. Oper. Res.*, vol. 127, no. 1, pp. 21–144, 2004.

[36] G. M. Campbell, "Cross-utilization of workers whose capabilities differ," *Manage. Sci.*, vol. 45, no. 5, pp. 722–732, 1999.

[37] S. Kim and D. A. Nembhard, "Cross-trained staffing levels with heterogeneous learning/forgetting," *IEEE Trans. Eng. Manage.*, vol. 57, no. 4, pp. 560–574, Nov. 2010.

[38] D. A. Nembhard and M. V. Uzumeri, "An individual-based description of learning within an organization," *IEEE Trans. Eng. Manage.*, vol. 47, no. 3, pp. 370–378, Aug. 2000.

[39] D. A. Nembhard and K. Prichanont, "Cross training in serial production with process characteristics and operational factors," *IEEE Trans. Eng. Manage.*, vol. 54, no. 3, pp. 565–575, Aug. 2007.

[40] B. D. Stewart, D. B. Webster, S. Ahmad, and J. Matson, "Mathematical models for developing a flexible workforce," *Int. J. Prod. Econ.*, vol. 36, no. 3, pp. 243–254, 1994.

[41] J. E. Gomar, C. T. Haas, and D. P. Morton, "Assignment and allocation optimization of partially multiskilled workforce," *J. Construction Eng. Manage.*, vol. 128, no. 2, pp. 103–109, 2002.

[42] W. J. Hopp, E. Tekin, and M. P. Van Oyen, "Benefits of skill chaining in serial production lines with cross-trained workers," *Manage. Sci.*, vol. 50, no. 1, pp. 83–98, 2004.

[43] J. Slomp and E. Molleman, "Cross-training policies and team performance," *Int. J. Prod. Res.*, vol. 40, no. 5, pp. 1193–1219, 2002.

[44] M. Dinger, J. B. Thatcher, L. P. Stepina, and K. Craig, "The grass is always greener on the other side: A test of present and alternative job utility on IT professionals turnover," *IEEE Trans. Eng. Manage.*, vol. 59, no. 3, pp. 364–378, Aug. 2012.

[45] Society for Human Resource Management. (2010). Annual voluntary turnover rate [Online]. Available: http://www.shrm.org/Research/Articles/Articles/Pages/MetricoftheMonthAnnua lVoluntaryTurnoverRate.aspx

[46] S. LaValle, E. Lesser, R. Shockley, M. S. Hopkins, and N. Kruschwitz, "Big data, analytics and the path from insights to value," *MIT Sloan Manage. Rev.*, vol. 52, no. 2, pp. 21–32, 2011.

Vincent Hargaden received the B.E. degree in mechanical engineering from University College Dublin (UCD), Dublin, Ireland, in 1995; the M.Eng.Sc. degree in industrial engineering from the National University of Ireland—Galway, Galway, Ireland, in 1997; the M.B.A. degree in management from UCD, in 2003; the Ph.D. degree in decision sciences and engineering systems from the Industrial and Systems Engineering Department, Rensselaer Polytechnic Institute, Troy, NY, USA, in 2011; and the Professional Diploma in university teaching and learning from UCD, in 2014.

He is a College Lecturer in Engineering Management and the Master of Engineering Management Program Director at UCD. His research interests include supply chain management, specifically design of resilient supply chains, as well as engineering education pedagogy.

Dr. Hargaden is a Member of the Institute of Industrial Engineers, the Institute for Operations Research and Management Science, the Production and Operations Management Society, and the European Operations Management Association.

Jennifer K. Ryan received the B.A. degree in mathematics and social sciences from Dartmouth College, Hanover, NH, USA, in 1990, and the M.S. and Ph.D. degrees from the Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, USA, in 1996 and 1997, respectively.

She is an Associate Professor of supply chain management at the University of Nebraska, Lincoln, NE, USA. Her research has been published in *Management Science*, *Operations Research*, *Production and Operations Management*, *Naval Research Logistics*, *IIE Transactions*, and IEEE TRANSACTIONS ON ENGINEERING MANAGEMENT. She has received several National Science Foundation grants, including a National Science Foundation CAREER grant. Her research expertise is in the area supply chain management, with a particular focus on coping with uncertainty and the role of information in managing supply chains.

Dr. Ryan currently serves as a Department Editor for *IIE Transactions*, as well as an Associate Editor for *Naval Research Logistics* and *OMEGA*.