

1-1-2011

A two-stage stochastic program for scheduling and allocating cross-trained workers

Gerard Campbell

Fairfield University, gcampbell@fairfield.edu

This is a post peer-review, pre-copyedit version of an article published in *Journal of the Operational Research Society*. The definitive publisher-authenticated version is available online at <http://www.palgrave-journals.com/jors/journal/v62/n6/abs/jors201016a.html>

Peer Reviewed

Repository Citation

Campbell, Gerard, "A two-stage stochastic program for scheduling and allocating cross-trained workers" (2011). *Business Faculty Publications*. Paper 13.
<http://digitalcommons.fairfield.edu/business-facultypubs/13>

Published Citation

Campbell, Gerard. A two-stage stochastic program for scheduling and allocating cross-trained workers. *Journal of the Operational Research Society* (2011) 62, 1038–1047. doi:10.1057/jors.2010.16

This Article is brought to you for free and open access by the Charles F. Dolan School of Business at DigitalCommons@Fairfield. It has been accepted for inclusion in Business Faculty Publications by an authorized administrator of DigitalCommons@Fairfield. For more information, please contact digitalcommons@fairfield.edu.

A two-stage stochastic program for scheduling and allocating cross-trained workers

Gerard M. Campbell

Information Systems and Operations Management, Charles F. Dolan School of Business,
Fairfield University, North Benson Road, Fairfield, CT 06824, USA.

Abstract

A two-stage stochastic program is developed for scheduling and allocating cross-trained workers in a multi-department service environment with random demands. The first stage corresponds to scheduling days-off over a time horizon such as a week or month. The second stage is the recourse action that deals with allocating available workers at the beginning of a day to accommodate realized demands. After the general two-stage model is formulated, a special case is introduced for computational testing. The testing helps quantify the value of cross-training as a function of problem characteristics. Results show that cross-training can be more valuable than perfect information, especially when demand uncertainty is high.

Keywords: manpower planning; labour scheduling; stochastic programming; cross-training

Introduction

In recent years, there has been a growing interest in the study of workforce cross-training. Hopp and Van Oyen (2004) describe, from a strategic perspective, how cross-training offers advantages along competitive dimensions of cost, time, quality, and variety. They also discuss how cross-training relates to the more general concept of workforce agility, observing that cross-training is “broadly applicable, powerful, and also highly complex.” Iravani, Van Oyen and Sims (2005) also emphasize the strategic value of cross-training. Jack and Powers (2004) discuss the importance of “volume flexible” strategies, including cross-training, in health services. Wallace and Whitt (2005) describe the value of cross-training in call center environments. Nembhard (2007) addresses a variety of service and manufacturing environments in a recently-published book on workforce cross-training.

Figure 1 displays a three-level framework for manpower planning and scheduling decisions. This framework was presented in Abernathy et al. (1973), and has been frequently cited elsewhere in the literature. Each of the levels in Figure 1 is further described below as relevant literature is reviewed. As indicated in the figure, this paper develops an integrated model that encompasses the scheduling and allocation levels of the framework. The first stage of the model, which corresponds to the middle level of the framework shown in Figure

1, schedules the days-off of workers over a multi-period planning horizon, such as a week or month. At that stage of the problem, requirements are random, and the objective is to maximize expected value. The second stage of the model corresponds to the lowest level of the framework shown in Figure 1. It allocates workers at the beginning of a day, considering the capabilities of available workers and the actual requirements that have been realized.

***** Insert Figure 1 about here *****

An example of an application environment for this paper's model is hospital nurse scheduling, which is the application area that motivated the Figure 1 framework developed by Abernathy et al. (1973). Warner et al. (1990) discuss automated nurse scheduling, covering all three levels shown in Figure 1. Campbell (2009) highlights healthcare employee scheduling as one of four key workforce scheduling application areas. For nurse scheduling, Ernst et al. (2004) and Burke et al. (2004) discuss how cost objectives must be balanced against various constraints, such as the variable requirements and heterogeneous employee capabilities included in this paper's model.

In previous research, cross-training has been modeled at each of the levels shown in Figure 1. At the top level, Nembhard, Nembhard, and Qu (2005) use a real options approach to model workforce planning in a manufacturing environment. Li and King (1998) present a planning-level model aimed at determining the number of cross-trained employees in a health service clinic. Agnihothri et al. (2003) use queuing models and simulation to find the best mix of dedicated and cross-trained workers in field service systems with two types of jobs. Kao and Queyranne (1985) develop a series of models focused on workforce planning for nurses at a hospital. One of their models is a two-stage stochastic program, where the first stage deals with sizing the regular-time workforce, and the second stage relates to the use of overtime and/or workers from outside agencies. Zhu and Sherali (2007) develop a two-stage stochastic program where the first stage relates to personnel hiring at multiple locations and the second stage allows for reassignments of work amongst the locations. Wright, Bretthauer and Cote (2006) present a bicriteria integer programming model for nurse scheduling that spans the top two levels of the framework shown in Figure 1.

The middle level of Figure 1 has been the subject of much attention in the literature. Ernst et al. (2004) present a survey of scheduling and rostering methods, which includes categories such as days-off scheduling and shift scheduling. Brusco and Johns (1998) develop a workforce scheduling model that enables the evaluation of alternative cross-training structures. Bard (2004) models staff scheduling with the option of moving workers with higher skill levels to areas requiring less skill. Wan and Bard (2007) look at shift

scheduling and task assignment when restrictions exist on the movement of flexible workers between workstation groups. The scheduling research presented by Morton and Popova (2004) does not include cross-training, but it does use a two-stage stochastic program that has monthly and daily stages, enabling adjustments to schedules over a rolling horizon. Bard and Purnomo (2005) use a rolling horizon approach in conjunction with an integer program for reactive scheduling of nurses in a hospital. Easton and Rossin (1996) present a stochastic program for workforce scheduling and demonstrate its advantages relative to deterministic models.

Moving from scheduling to the allocation level in Figure 1, the research of Campbell (1999) and Campbell and Diaby (2002) is most directly relevant. Brusco (2008) developed an exact algorithm for the model presented in Campbell (1999). The two-stage formulation developed in the current study is an extension of the model presented in Campbell (1999).

Model development

Formulation of the general two-stage model

The first stage of the two-stage model corresponds to scheduling days-off for a fixed set of cross-trained workers over a finite planning horizon (e.g., a week or a month). Each worker is assigned to a tour that represents their schedule for the entire planning horizon. At this stage, labour requirements for each day in each department are represented by random variables. The second stage models the allocation of cross-trained workers to departments once requirements have been realized. Thus, the second stage is actually a series of single-day allocation problems with deterministic labour requirements. In practice, one of these problems would be solved at the beginning of each day, based on actual requirements and the set of workers scheduled for that day.

The following summarizes the notation used for the two-stage model:

$x_{ij} = 1$ if worker i is assigned to tour j , 0 otherwise;

X = the set of established tours;

c_{ij} = the cost when worker i is scheduled for tour j ;

(Note that cross-trained workers will typically cost more than non-flexible workers.)

$a_{tj} = 1$ if period t is a work period in tour j ; 0 otherwise;

$y_{idt} = 1$ if worker i is allocated to department d in period t , 0 otherwise;

w_{dt} = labour used in department d in period $t = \sum_{i=1}^I p_{id}y_{idt}$;

p_{id} = the productivity of worker i in department d ($0 \leq p_{id} \leq 1$);

r_{dt} = the random amount of labour required in department d in period t;

R = the set of random labour requirements;

r_{dt} = the realized deterministic amount of labour required in department d in period t;

R = the set of realized deterministic labour requirements;

$u_{dt}(w_{dt})$ = the random concave utility function for department d in period t.

This function depends on the random labour requirements (r_{dt}). For example, based on Campbell (1999), an objective function that reflects quadratic shortage costs can be defined using the following:

$$u_{dt}(w_{dt}) = k_d[r_{dt}^2 - (r_{dt} - w_{dt})^2] \text{ for } w_{dt} < r_{dt}, k_d r_{dt}^2 \text{ otherwise,}$$

where k_d is a constant representing a weighting factor for department d. With this type of objective function, departments that have higher priorities would have higher k_d values.

U = the set of random departmental utility functions;

$u_{dt}(w_{dt})$ = the realized concave utility function for department d in period t.

This function depends on the realized labour requirements (r_{dt}). For example:

$$u_{dt}(w_{dt}) = k_d[r_{dt}^2 - (r_{dt} - w_{dt})^2] \text{ for } w_{dt} < r_{dt}, k_d r_{dt}^2 \text{ otherwise.}$$

U = the set of realized departmental utility functions;

T is the number of periods in the planning horizon;

D is the number of departments;

I is the number of workers;

N_i is the set of possible tours for worker i;

S_i is the set of departments for which $p_{id} > 0$;

E_R denotes mathematical expectation with respect to R .

The two-stage stochastic integer programming (TSSIP) formulation can now be presented as follows:

Problem TSSIP:

$$\text{Stage 1:} \quad \begin{aligned} & \text{maximize } Z = \sum_{i=1}^I \sum_{j \in N_i} -c_{ij}x_{ij} + E_R[f(X, U)] \end{aligned} \quad (1)$$

$$\text{subject to} \quad x_{ij} = 0 \text{ or } 1 \quad \text{for all } i, \text{ and } j \in N_i. \quad (2)$$

Stage 2:

$$f(X, U) = \text{maximize} \sum_{d=1}^D \sum_{t=1}^T u_{dt}(w_{dt}) \quad (3)$$

subject to

$$w_{dt} = \sum_{i=1}^I p_{id} y_{idt} \quad \text{for all } d, t, \quad (4)$$

$$\sum_{d \in S_i} y_{idt} = a_j x_{ij} \quad \text{for all } i, t, \text{ and } j \in N_i, \quad (5)$$

$$y_{idt} = 0 \text{ or } 1 \quad \text{for all } i, t, \text{ and } d \in S_i. \quad (6)$$

The formulation given by (1) – (6) represents a two-stage stochastic integer program with recourse. In solving stage 1, the x_{ij} values are established before departmental requirements are known with certainty. Note that the number of workers, I , is the model's only limit on capacity. Also, infeasibility is not an issue because there are no constraints stating that departmental requirements must be met.

Stage 2 corresponds to allocating cross-trained workers at the beginning of each day, based on the workers available that day (given by X) and the departmental utility functions for the day (which depend on R). Note that (3) – (6) is in the same form as the allocation problem studied in Campbell (1999) and Campbell and Diaby (2002). The current formulation extends the prior work to include multiple time periods. The multi-period problem separates into a series of single-period problems that, in practice, would be solved day-by-day as actual departmental requirements were realized. These stage 2 problems can be solved using the heuristic developed in Campbell and Diaby (2002), and problems where all p_{id} values are 0/1 can be solved optimally using a simple assignment algorithm.

Other formulations from previous research correspond to special cases of TSSIP. For example, when there is only one department, then there is no cross-training and no relevant second stage – i.e., no allocation problem. In that case, TSSIP reduces to a problem similar to that studied by Easton and Rossin (1996). What happens to TSSIP if requirements are deterministic rather than stochastic? Again, there is no need for a second stage. If all departmental requirements are known with certainty when tours are being assigned, then allocations can also be fixed at that time. In this case, the problem is similar to that studied by Brusco and Johns (1998), who modeled the scheduling of cross-trained workers in a multi-department environment with deterministic requirements. The special case of TSSIP with one department and deterministic requirements corresponds to the classical single-stage tour scheduling problem, as described in Ernst et al. (2004).

Having presented the general formulation and discussed its relationship to prior research, the current study now focuses on a special case of the two-stage problem.

Network representation of a special case

For two-stage stochastic programming problems, Birge and Louveaux (1999) define the following:

RP = the value of the optimal solution to the stochastic programming problem;

WS = the value of an optimal “wait-and-see” solution – i.e., the value of the optimal solution if future requirements were known with certainty; and

EVPI = the expected value of perfect information.

EVPI is calculated as follows for a maximization problem: $EVPI = WS - RP$.

Due to the complexity of TSSIP, optimal solutions are difficult to obtain for problems of realistic size. However, for the special case outlined below, WS can be established using a network programming formulation. Also, an heuristic (available from the author) provides feasible solutions, which can be used to establish lower bounds on RP. Thus, upper bounds on EVPI can be obtained for the special case of TSSIP developed below.

The special case studied in this paper is one where scheduling is done for a seven-day planning horizon, and the set of feasible tours for each employee, N_i , includes all five-day schedules (i.e., there are no additional requirements, such as two consecutive days off). The workforce size is fixed, and the c_{ij} values do not vary, so departmental utilities control the objective function. Furthermore, the special case includes productivity values (p_{id}) that are 0/1 for all employees in all departments. This enables the network model shown in Figure 2 to be used for obtaining optimal wait-and-see solutions.

***** Insert Figure 2 about here *****

The supply of five shown for each worker in Figure 2 represents the worker’s capacity of five work days. Although not shown in the figure, all arcs have an upper bound on flow equal to one. Utility values are only shown for the final set of arcs, but if worker costs differed by day of week, that could be reflected in the first set of arcs in the network. The second set of arcs shown in Figure 2 reflects the capabilities of the workers – i.e., worker 1 is trained for departments 1 and 3, worker 2 is trained for 2 and 4, and worker I is trained for department 3 only. The u_{dt} values shown along the final set of arcs actually represent a series of values, as shown for one of the department/day combinations in the figure. Utility contributions along the arcs are defined as follows: $\delta_{dtn} = u_{dt}(n) - u_{dt}(n-1)$, where δ_{dtn} is the marginal contribution of the nth worker in department d in period t. In the figure, these values are shown for department 4/day 2 for $n = 1, \dots, Q$, where Q is the maximum number of

workers that can be assigned to a department in one day. Network problems of the form shown in Figure 2 can be quickly solved using available software, so optimal solutions to the special case when requirements are deterministic are easily obtained. The computational testing described in the next section includes an evaluation of an heuristic (available from the author) that uses an enhanced network formulation to establish solutions to problems with random requirements.

Computational testing

The purposes of the computational testing are: 1) to evaluate the value of cross-training; and 2) to gain insight into the value of perfect information.

Solution values were collected using the measures outlined below. Note that mean values are based on 10 sets of realized requirements per problem.

M_{fixed} = mean value of Z for a fixed assignment solution. A fixed assignment solution fixes departmental allocations at stage 1 and does not allow subsequent changes.

M_{cross} = mean value of Z for a solution with cross-utilization, where X is obtained using the heuristic.

M_{pi} = mean value of Z for wait-and-see solutions, which are based on perfect information.

M_{upper} = mean value of Z with perfect information, with the added constraint of equal numbers of workers in all periods. For the test problems used in this study, the validity of M_{upper} as an upper bound on the optimal solution to TSSIP is based on symmetry and concavity. Each test problem has requirements distributions that are the same for all days of the planning horizon, and the average number of workers per department per day is always a whole number. Because objective functions are all concave, any deviation from an equal number of workers per day in the stage 1 solution to TSSIP would result in a decrease in expected utility. Therefore, no optimal solution to TSSIP can provide higher utility for a realized set of requirements than a wait-and-see solution that includes an equal number of workers in all periods.

The following relationships apply for the solution measures:

$$M_{pi} \geq M_{upper} \geq M_{cross} \geq M_{fixed}.$$

Using these solution measures, the following performance measures are defined:

$$GAP = (M_{upper} - M_{cross}) / M_{upper},$$

$$V_{cross} = (M_{cross} - M_{fixed}) / M_{fixed},$$

$$V_{pi} = (M_{pi} - M_{cross}) / M_{pi},$$

where GAP is the gap between heuristic solution and upper bound, V_{cross} is the value of cross-utilization, and V_{pi} is the value of perfect information.

Experimental factors and the test environment

Test problems were generated to investigate the effects of five problem characteristics using a $2^4 \times 4$ experimental design. The experimental factors and their treatment levels are summarized in the first column of Table 1. These factors and the test environment were established so that stage 2 problems in the current study are comparable to problems solved in Campbell (1999) and Campbell and Diaby ((2002).

***** Insert Table 1 about here *****

One of the fixed characteristics of the test environment is the functional form of departmental utility functions, which is defined as follows:

$$u_{dt}(w_{dt}) = k_d[r_{dt}^2 - (r_{dt} - w_{dt})^2] \text{ for } w_{dt} < r_{dt}, k_d r_{dt}^2 \text{ otherwise.}$$

This form is consistent with that used in Campbell (1999). For the current study, departments are homogeneous for all test problems (i.e., $k_d = 1$ for all $d = 1, \dots, D$). Note that departmental differences did not significantly affect the value of cross-utilization in Campbell (1999). Requirements distributions are also identical across departments and time periods, as described below under the forecast error factor.

Another fixed characteristic of the test environment is the planning horizon, which is seven days. The set of feasible tours for all workers includes all five-day schedules – i.e., there are no additional restrictions such as a requirement for two consecutive days off. This enables the network formulation shown in Figure 2 to be used for solving wait-and-see versions of the problems. The costs of all five-day tours are assumed to be equal, so all c_{ij} values could be set to zero. Figure 3 shows an example of problem data used in the experiment. For each test problem, \mathbf{R} , the set of random labour requirements, is defined by treatment levels of the shortage level and forecast error factors. Each of the five experimental factors shown in Table 1 is further described below.

***** Insert Figure 3 about here *****

Level of cross-training. This factor is explored at the following four treatment levels: 1.5, 2.0, 2.5 and 3.0. A cross-training level of 2.0 means that all workers are trained for two departments, and 1.5 means that half the workers are trained for two departments and half are only trained for one. In generating the test problems, each worker is assumed to have a “primary” department, and these primary capabilities are evenly distributed across departments. Mfixed is based on each worker being assigned to their primary department. Secondary and tertiary departments were selected randomly to achieve the desired level of cross-training for each test problem. Note that all p_{id} values are 0/1 for all test problems.

This enables the network formulation shown in Figure 2 to provide optimal solutions to the wait-and-see problems used to establish M_{pi} and M_{upper} .

Number of departments. $D = 4$, or $D = 8$.

Workers per department (wpd). $wpd = 7$, or $wpd = 14$. With each worker getting two days off out of seven, these treatment levels correspond to the wpd levels of five and ten used in the single-stage environment studied in Campbell (1999) and Campbell and Diaby (2002).

Shortage level. This factor relates to the departmental requirements, r_{dt} . The treatment levels for this factor are 10 % and 20 %. Considering the number of workers per department (wpd), each working five days out of seven, the shortage level corresponds to the following:

$$\text{shortage level} = (\mu - \text{wpd}(5/7)) / \mu,$$

where μ is the mean r_{dt} value. For example, for the problem shown in Figure 3, $\mu = 6.25$.

The manner in which individual r_{dt} values were established is described under the next factor.

Forecast error. For each test problem, the r_{dt} values were generated from a single normal distribution. The mean of the distribution (μ) was established to provide the desired treatment level of the shortage factor, as described above. The standard deviation of the distribution was either 0.3μ or 0.6μ , providing the factor treatment levels of 0.3 and 0.6. For realized requirements, randomly-generated values were adjusted proportionally to ensure that the specified shortage level was achieved for each test problem.

The $2^4 \times 4$ full factorial design results in 64 factor combinations, or cells. With four replications per cell, there are 256 test problems. For each test problem, ten sets of realized requirements values (r_{dt}) were generated. Performance measures represent averages over the ten sets of realized requirements per problem. To enable direct comparisons, the same sets of requirements were used across all treatment levels of the cross-training factor.

Results

The last three columns in Table 1 summarize results based on the GAP, V_{cross} , and V_{pi} performance measures. Figs. 4, 5 and 6 provide further details.

The overall mean GAP measure of 0.9 % shown in Table 1 demonstrates that the heuristic provides solutions that are close to optimal. Looking at GAP as a function of problem characteristics, it is clear that level of cross-training has the most significant effect, with lower levels of cross-training having larger GAP values. This makes sense because the solution to the stage 1 problem represents more of a constraint when there is less flexibility to move workers around at stage 2. The forecast error factor is significant because when there are larger variations in requirements, the cost of not knowing requirements with certainty at

stage 1 increases. The effect of the workers per department factor on the GAP performance measure occurs because of the reduced amount of flexibility available when there are fewer cross-trained workers.

Figures 4a and 4b show interaction effects between level of cross-training and forecast error for both the GAP and Vpi performance measures. In terms of the GAP measure, a low degree of cross-training and high forecast errors result in the worst performance, with GAP averaging 4.7 % for those types of problems. A large GAP value can be the result of a large difference between the heuristic solution and the (unknown) optimal solution, and/or a large difference between the optimal solution and the upper bound. As described later when Figure 6 is discussed, evidence suggests that it is the difference between the optimal solution and the upper bound that is increasing most dramatically as the level of cross-training decreases.

***** Insert Figure 4 about here *****

With respect to Vcross, the value added by cross-utilization, the overall average shown in Table 1 is 9.1 %. Vcross is most seriously affected by the forecast error factor, with a mean value of 13.3 % when forecast errors are high, compared to 4.9 % when they are low. Obviously, the flexibility of moving workers around to accommodate random departmental requirements is more valuable when variability is higher. The level of cross-training also has a significant effect, but changes in Vcross are relatively small above a cross-training level of 2.0. This is consistent with results from earlier studies, which demonstrated that high levels of cross-training are not required to achieve the benefits of cross-utilization. Figure 5 displays Vcross as a function of forecast error and level of cross-training.

***** Insert Figure 5 about here *****

Looking at Vpi, Table 1 shows that the value of perfect information averaged 2.9 % for the problems solved in this study. Forecast error obviously has the strongest effect on this measure, with Vpi values averaging 1.6 % and 4.2 % for low and high forecast errors, respectively. Level of cross-training also had an effect, especially at the 1.5 treatment level. Figures 4a and 4b show Vpi as a function of level of cross-training and forecast error. Notice the sharp drop in the value of perfect information as level of cross-training increases from 1.5 to 2.0. Figs. 6a and 6b provide further insight into the behavior of Vpi.

***** Insert Figure 6 about here *****

Figures 6a and 6b show average solution values as a percentage of the highest Mpi values, which were obtained when the level of cross-training was at 3.0. The curves show how Mcross, Mupper, and Mpi relate to each other and how they increase as the level of

cross-training increases. The figures include values for fixed assignment solutions, which correspond to a level of cross-training of 1.0. At the 1.0 level of cross-training, M_{cross} values represent optimal solutions, while they represent heuristic solutions at higher levels of cross-training.

Figures 6a and 6b show that cross-training can be more valuable than perfect information, especially when forecast errors are high. For example, in Figure 6b, M_{cross} at the 2.0 level of cross-training is higher than M_{pi} at the 1.0 level of cross-training. For those problems, it was more valuable to have flexible workers than to have perfect information regarding daily requirements. Even if you knew when high requirements were going to occur in certain departments, without cross-training you would not be able to schedule enough workers for those days to avoid high shortage costs. With cross-training, you can accommodate higher requirements in some departments as stage 2 allocation problems are solved. This analysis of the value of cross-training relative to the value of perfect information has important practical implications for workforce managers faced with uncertain requirements.

Another observation with respect to Figures 6a and 6b relates to the performance of the heuristic and the GAP performance measure. The GAP measure is based on the difference between M_{upper} and M_{cross} , and it was noted previously that this measure was largest at the 1.5 level of cross-training. Figures 6a and 6b provide evidence that the larger values of GAP at the 1.5 level of cross-training are due primarily to large differences between optimal solutions and upper bounds, rather than differences between heuristic and optimal solutions. This can be inferred because when the level of cross-training is 1.0, M_{cross} represents optimal solutions, yet the difference between M_{cross} and M_{upper} is still large. Inferences are not required at cross-training levels of 2.0 and above because then the small differences between M_{cross} and M_{upper} ensure that heuristic solutions are close to optimal. Therefore, conclusions regarding the value of cross-training could be drawn, even though the complexity of the two-stage stochastic program made it impractical to try to obtain optimal solutions.

Conclusion

This paper has made a contribution by extending previous research to develop a two-stage stochastic program for scheduling and allocating cross-trained workers. The general two-stage formulation presented in (1) – (6) is useful for describing the problem faced by workforce managers, but its complexity makes it difficult to solve. Therefore, for computational testing this study used a special case based on 0/1 worker capabilities and five

day-per-week scheduling without additional requirements such as consecutive days off. For this special case, heuristic solutions, upper bounds, and optimal wait-and-see solutions could be obtained. This enabled investigation of the value of cross-utilization and perfect information within the integrated two-stage environment.

Problems for the computational testing were structured to be consistent with those described in Campbell (1999) and Campbell and Diaby (2002), both in terms of problem characteristics and performance measures. The value of cross-utilization, as represented by the V_{cross} performance measure, showed similar behavior to that seen in Campbell (1999). A key contribution beyond previous research is the investigation of the value of perfect information relative to the value of cross-utilization. Results suggest that managers would be better off having workers that are each capable of working in two departments rather than having perfect information and no allocation flexibility. In practice, perfect information is rarely obtainable, but in many environments some degree of cross-training can be implemented. Scheduling and allocation models such as that presented in the current study have the potential to help managers better utilize cross-trained workers. Computational studies based on more realistic versions of the two-stage problem would be necessary to support this paper's conclusions. The development of solution algorithms for more general version of the two-stage problem could be motivated by the results of the experiments based on the special case studied in this paper.

The current study also suggests a number of other areas for future research. For example, the current study has not included alternative methods for handling demand variation, such as the use of overtime seen in Zhang et al. (2009). It should be noted that overtime and outside agencies, which are both widely used in practice for hospital nurse scheduling, could reduce the advantages of cross-training that were seen in the current study.

The planning level shown in Figure 1 may present another attractive area for further investigation. At this level, the costs of cross-training are traded-off against its anticipated value. Future research could also include the development of methods to solve the general version of the problem given by (1) – (6), and solution methods for other special cases. The results of the current study suggest that the development of methods for determining the most appropriate level of cross-training as a function of environmental characteristics is a most important area for further study.

References

- Abernathy, W.J., Baloff, N., Hershey, J.C., Wandel, S. (1973). A three-stage manpower planning and scheduling model – a service sector example. *Opns Res* 21: 693-711.
- Agnihotri, S.R, Mishra, A.K., and Simmons, D.E. (2003). Workforce cross-training decisions in field service systems with two job types. *J Opl Res Soc* 54: 410-418.
- Al-Yakoob, S.M., and Sherali, H.D. (2008). A column generation approach for an employee scheduling problem with multiple shifts and work locations. *J Opl Res Soc* 59: 34-43.
- Bard, J.F. (2004). Staff scheduling in high volume service facilities with downgrading. *IIE Transac* 36: 985-997.
- Bard, J.R., Purnoma, H.W. (2005). Hospital-wide reactive scheduling of nurses with preference considerations. *IIE Transac* 37: 589-608.
- Birge, J.R., Louveaux, F. (1999). *Introduction to Stochastic Programming*. Springer-Verlag New York, Inc.
- Brusco, M.J. (2008). An exact algorithm for a workforce allocation problem with application to an analysis of cross-training policies. *IIE Transac* 40: 495-508.
- Brusco, M.J., Johns, T.R. (1998). Staffing a multiskilled workforce with varying levels of productivity: an analysis of cross-training policies. *Dec Sci* 29: 499-516.
- Burke, E.K., De Causmaecker, P., Vanden Berghe, G., Van Landeghem, H. (2004). The State of the Art of Nurse Rostering. *J Sched* 7: 441-499.
- Campbell, G.M. (1999). Cross-utilization of workers whose capabilities differ. *Manage Sci* 45: 722-732.
- Campbell, G.M. (2009). Overview of Workforce Scheduling Software. *Prod & Inv Manage*, Vol. 45, No. 2, 2009.
- Campbell, G.M., Diaby, M. (2002). Development and evaluation of an assignment heuristic for allocating cross-trained workers. *Eur J Opl Res* 138: 9 – 20.
- Easton, F., Rossin, D. (1996). A stochastic goal program for employee scheduling. *Dec Sci* 27: 541-568.
- Ernst, A.T., Jiang, H., Krishnamoorthy, M., Sier, D. (2004). Staff scheduling and rostering: a review of applications, methods and models. *Euro J Opl Res* 153: 3-27.
- Hopp, W.J., Van Oyen, M.P. (2004). Agile workforce evaluation: a framework for cross-training and coordination. *IIE Transac* 36: 919-940.
- Iravani, S.M., Van Oyen, M.P., Sims, K.T. (2005). Structural flexibility: a new perspective on the design of manufacturing and service operations. *Manage Sci* 51: 151-166.

Jack, E.P., Powers, T.L. (2004). Volume flexible strategies in health services: a research framework. *Prod Oper Manage* 13: 230-244.

Jordan, W., Inman, R., Blumenfeld, D. (2004). Chained cross-training of workers for robust performance. *IIE Transac* 36: 953-967.

Kao, E.P.C., Queyranne, M. (1985). Budgeting costs of nursing in a hospital. *Manage Sci* 31: 608-21.

Li, L.L.X., King, B.E. (1999). A healthcare staff decision model considering effects of staff cross-training. *Health Care Manage Sci* 2: 53-61.

Morton, D.P., Popova, E. (2004). A Bayesian stochastic programming approach to an employee scheduling problem. *IIE Transac* 36 : 155-167.

Nembhard, D.A. (2007). *Workforce cross training*. CRC Press, Boca Raton, Florida.

Nembhard, D.A., Nembhard, H.B., Qin, R. 2005. A real options model for workforce cross-training. *Engin Econ* 50: 95-116.

Wallace, R.B., Whitt, W. (2005). A staffing algorithm for call centers with skill-based routing. *Manuf & Service Oper Manage* 7: 276-294.

Wan, L. and Bard, J.F. (2007). Weekly staff scheduling with workstation group restrictions. *J Opl Res Soc* 58: 1030–1046.

Warner, D.M., Keller, B.J., Martel, S.H. (1990) Automated Nurse Scheduling. *J. Soc. Health Systems* 2: 66-80.

Wright, P.D., Bretthauer, K.M., Cote, M.J. (2006). Reexamining the nurse scheduling problem: staffing ratios and nursing shortages. *Dec Sci* 37: 39-70.

Zhang, X., Chakravarthy, A., Gu, Q. (2009). Equipment scheduling problem under disruptions in mail processing and distribution centres. *J Opl Res Soc* 60: 598-610.

Zhu, X. and Sherali, H.D. (2007). Two-stage workforce planning under demand fluctuations and uncertainty. *J Opl Res Soc* 58: 1-10.

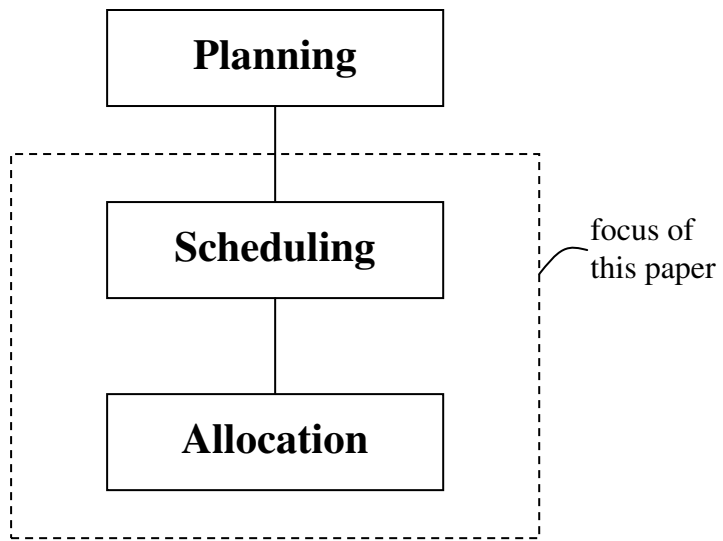


Figure 1. Framework for manpower planning and scheduling decisions.

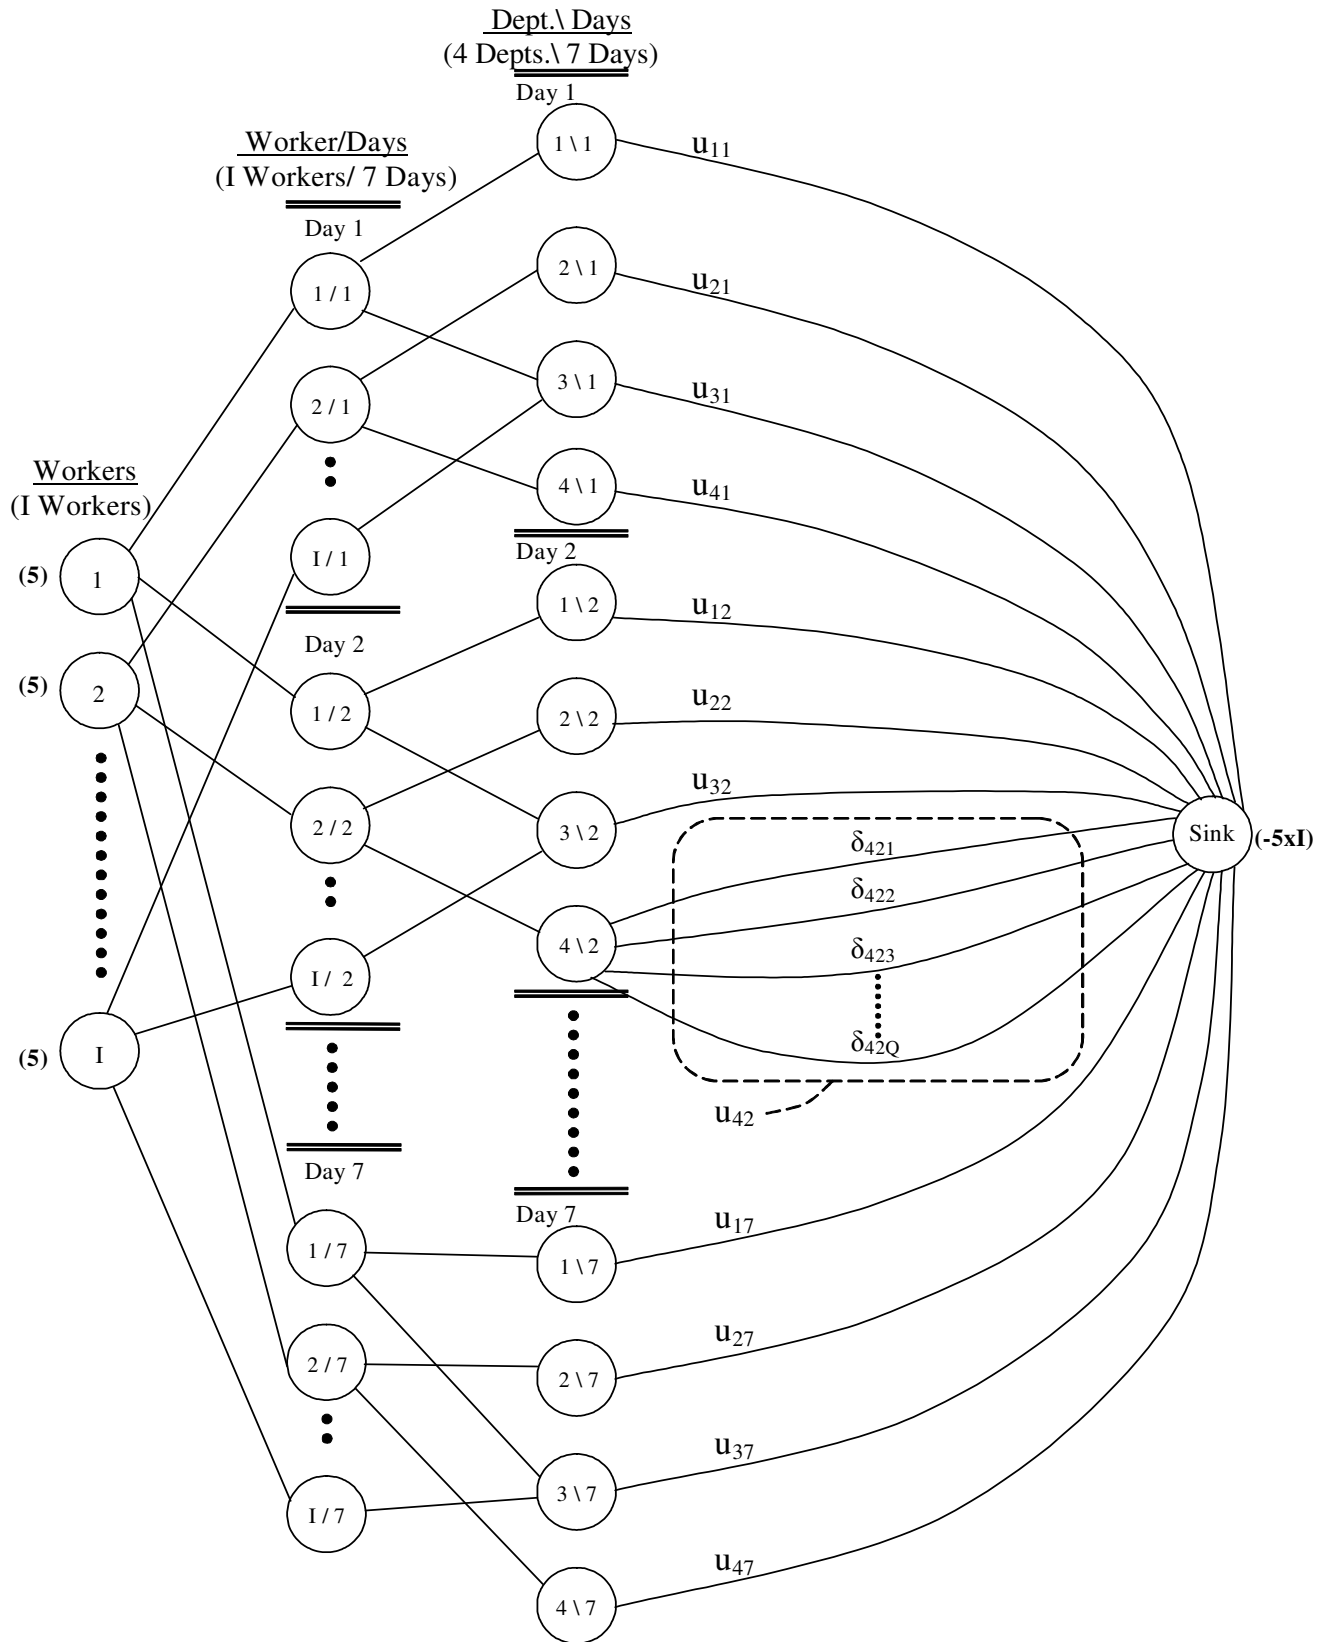


Figure 2. Network representation of a special case with perfect information.

	Dept. 1	Dept. 2	Dept. 3	Dept. 4
Day 1	6.25	9.46	6.44	8.36
Day 2	7.28	5.18	6.41	3.70
Day 3	6.02	8.59	5.58	5.43
Day 4	9.53	5.37	6.59	7.98
Day 5	8.83	5.41	2.74	5.66
Day 6	5.64	6.70	6.41	5.78
Day 7	5.16	4.92	5.53	3.89

Realized
Departmental
Requirements
(r_{dt})

Worker 1	<u>1</u>	1	0	0
Worker 2	0	<u>1</u>	1	0
Worker 3	0	1	<u>1</u>	0
Worker 4	0	1	0	<u>1</u>
Worker 5	<u>1</u>	0	1	0
Worker 6	0	<u>1</u>	0	1
Worker 7	0	0	<u>1</u>	1
Worker 8	0	1	0	<u>1</u>
Worker 9	<u>1</u>	0	1	0
Worker 10	0	<u>1</u>	0	1
Worker 11	0	0	<u>1</u>	1
Worker 12	0	1	0	<u>1</u>
Worker 13	<u>1</u>	0	1	0
Worker 14	1	<u>1</u>	0	0
Worker 15	1	0	<u>1</u>	0
Worker 16	0	0	1	<u>1</u>
Worker 17	<u>1</u>	1	0	0
Worker 18	0	<u>1</u>	1	0
Worker 19	0	0	<u>1</u>	1
Worker 20	0	1	0	<u>1</u>
Worker 21	<u>1</u>	0	1	0
Worker 22	0	<u>1</u>	1	0
Worker 23	1	0	<u>1</u>	0
Worker 24	0	1	0	<u>1</u>
Worker 25	<u>1</u>	0	1	0
Worker 26	1	<u>1</u>	0	0
Worker 27	0	0	<u>1</u>	1
Worker 28	1	0	0	<u>1</u>

Worker
Productivities
(p_{id})

the "primary"
department of
each worker is
underlined

<u>Factor</u>	<u>Treatment Level</u>
Level of Cross-Training	2.0
Number of Departments	4
Workers per Department	7
Shortage Level	20 %
Forecast Error	0.3

Figure 3. Example of problem data used in the experiment.





Figure 4. Value of perfect information (Vpi) and heuristic performance (GAP) as a function of level of cross-training and forecast error.

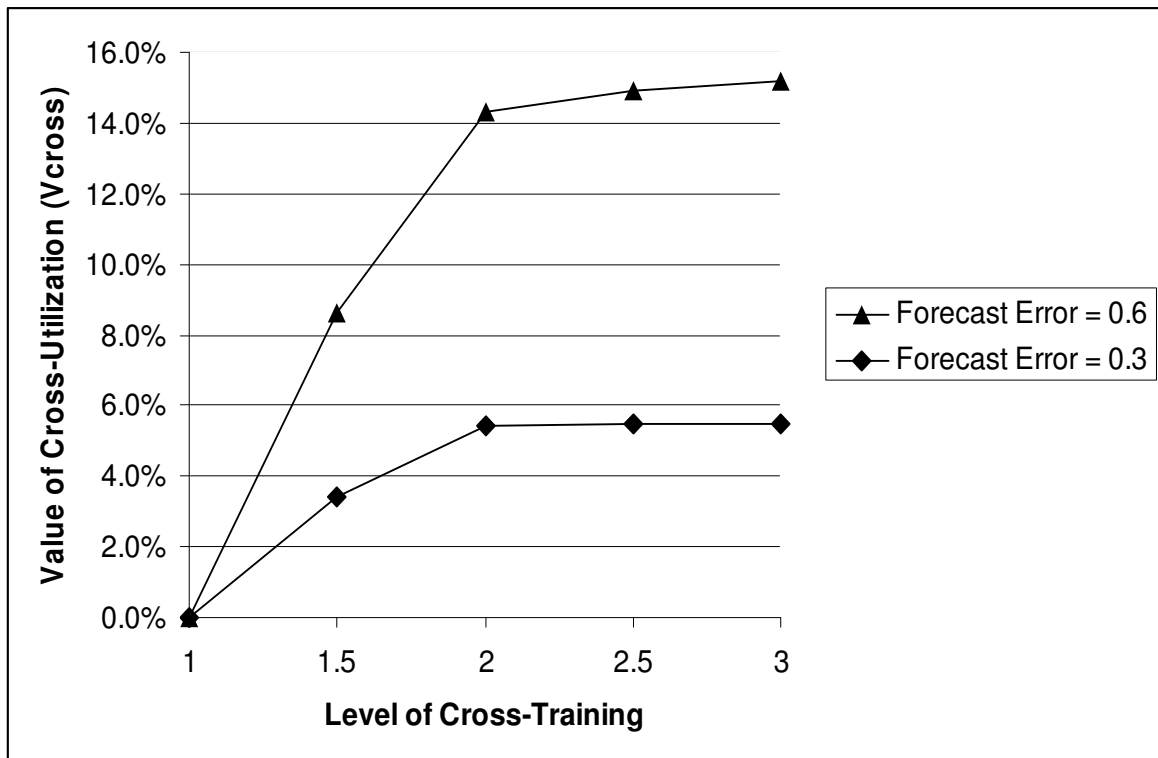


Figure 5. Value of cross-utilization as a function of level of cross-training and forecast error.

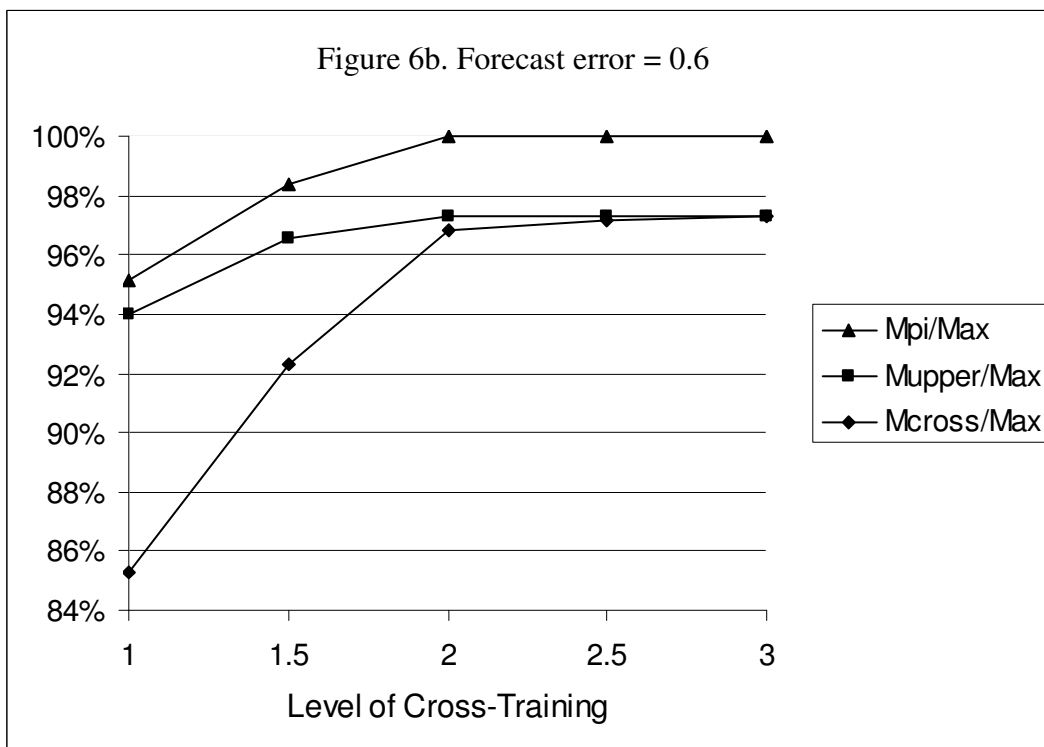



Figure 6. Solution values by level of cross-training, expressed as a percentage of the maximum Mpi (Max) found at each level of forecast error.

Table 1. Summarized results of the computational testing.

Experimental Factor / Treatment Level	GAP	Vcross	Vpi
Level of Cross-training			
1.5	3.2%	6.0%	4.7%
2.0	0.4%	9.9%	2.5%
2.5	0.1%	10.2%	2.3%
3.0	0.0%	10.3%	2.2%
Number of Departments			
4	0.9%	8.7%	3.5%
8	1.0%	9.5%	2.4%
Workers per Department			
7	1.2%	9.9%	3.2%
14	0.7%	8.3%	2.6%
Shortage Level			
10%	0.8%	8.4%	2.6%
20%	1.0%	9.8%	3.2%
Forecast Error			
0.3	0.4%	4.9%	1.6%
0.6	1.4%	13.3%	4.2%
Overall Means	0.9%	9.1%	2.9%