2004

A Comparison of Heuristics for Assigning Individual Employees to Labor Tour Schedules

John C. Goodale
University of Oregon

Gary Thompson
Cornell University School of Hotel Administration, gmt1@cornell.edu

A Comparison of Heuristics for Assigning Individual Employees to Labor Tour Schedules

Abstract

The labor tour scheduling literature has focused on the development of schedules, and with a few exceptions, employees were assumed to have identical cost and productivity. Even the few exceptions in the literature that solved tour problems considered employees within a work group to have identical cost and productivity. In this paper we evaluated heuristics for assigning individual employees – who differed in cost and productivity – to labor tour schedules. Our results showed that considering productivity levels when assigning individuals to tours increased profitability. We found that a simple managerial heuristic of assigning individuals in descending order of their productivity to cost ratio was both fast and effective over a broad range of service environmental scenarios.

Keywords

individual productivity, labor scheduling, simulated annealing, tour scheduling

Disciplines

Business Administration, Management, and Operations | Performance Management

Comments

Required Publisher Statement:

A Comparison of Heuristics for Assigning Individual Employees to Labor Tour Schedules

John C. Goodale

University of Oregon, Eugene

Gary M. Thompson

Cornell University

Abstract

The labor tour scheduling literature has focused on the development of schedules, and with a few exceptions, employees were assumed to have identical cost and productivity. Even the few exceptions in the literature that solved tour problems considered employees within a work group to have identical cost and productivity. In this paper we evaluated heuristics for assigning individual employees – who differed in cost and productivity – to labor tour schedules. Our results showed that considering productivity levels when assigning individuals to tours increased profitability. We found that a simple managerial heuristic of assigning individuals in descending order of their productivity to cost ratio was both fast and effective over a broad range of service environmental scenarios.

Keywords: individual productivity, labor scheduling, simulated annealing, tour scheduling

1. Introduction

Managers of labor-intensive service operations have the important responsibility of generating tour schedules, that is, schedules that specify each employee's daily starting time and days-off during the week. Usually focused on maximizing profitability, managers are constrained in this task by the need to have sufficient service capacity in each planning period to meet the expected level of customer demand. A variety of other factors influence and constrain employee scheduling decisions, including the number of employees available to work, employee costs, shift length, the level of variation in customer demand, and customers' displeasure with waiting in line.

Employee tour scheduling, which is typically treated as a general set-covering problem, has received a substantial amount of attention in operations research. The amount of work conducted in this area reflects the complex and rich environment in which staffing decisions are made. For example, recently published research has focused on scheduling full- and part-time employees (Easton and Rossin, 1991a, 1991b; Li, Robinson, and Mabert, 1991), scheduling with a limited number of starting times for employees (Brusco and Jacobs, 1998b), employees with dynamic productivity (Goodale and Tunc, 1997), employees with travel times and costs (Beaumont, 1997), and heuristic algorithms for solving the set covering problem due to its NP-hardness (Brusco and Jacobs, 1993, 1998a; Brusco, Jacobs, and Thompson, 1999; Caprara, Toth, and Fischetti, 2000; Easton and Mansour, 1999). In addition, other literature has focused on sub-problems of tour scheduling, like common start times and only specifying days of the week employees work (days-off scheduling; see, for example, Dowling et al., 1997; Emmons and Fuh, 1997), or considered only starting times and ignored the days of the week that employees work (shift scheduling; see, for example, Brusco and Johns, 1998; Goodale, Verma, and Pullman, 2003).

In this paper, we examine rules for assigning individual, heterogeneous employees to tour schedules. Previously, models in the labor tour scheduling literature considered employees as homogeneous (see, for example, Baker, 1976; Bechtold, Brusco, and Showalter, 1991; Brusco and Jacobs, 1993; Dantzig, 1954; Keith, 1979; Thompson, 1995b) or homogenous within groups when generating tour schedules (see, for example, Easton and Rossin, 1991a, 1991b; Li, Robinson, and Mabert, 1991). Related research showed that significant benefits were achieved when different productivity levels of cross-trained employees were considered and scheduled (Brusco and Johns, 1998), and when heterogeneous, cross-trained employees were allocated without regard to schedules (Campbell, 1999). Thus, we hypothesized that more profitable tour schedules are possible by considering individual productivity and hourly wage when workers are assigned to tour schedules. In this paper, we evaluated four assignment heuristics. Three heuristics were simple rules that managers could easily implement, and one was a more complex heuristic based on the simulated annealing heuristic.

The paper is organized as follows. The next section contains a review of the relevant work in labor tour scheduling and discusses how we solved the tour scheduling problem assuming generic employees. The individual worker assignment (IWA) heuristic methods are described in section 3. Section 4 provides the design of the full-factorial simulation experiment we conducted to evaluate the IWA heuristics. Section 5 presents results of the experiment. Finally, section 6 discusses the results and provides conclusions.

2. Labor scheduling background

In this section we review relevant literature and present the model we employ.

2.1. Relevant literature

Dantzig (1954) provided the original formulation of labor scheduling, for a problem of determining the number of toll collectors required to staff various shifts to meet demand. His work spawned a wealth of relevant, applied research in operations management and operations research. Labor scheduling is defined as having four component tasks (Thompson, 1993, 1995a), which we refer to as T1–T4. T1 forecasts customer demand for the service (Thompson, 1998a). T2 converts the forecasts of demand into requirements for service workers/servers (Thompson, 1998b). T3 generates the least costly or most profitable workforce schedule (Thompson, 1999a). This schedule specifies the number of employees that work the different work patterns (typically identified by days worked, start time, hours worked, and break times), and identifies the number of employees that will be working during each planning period. Last, T4 controls the real-time delivery of the service (Thompson, 1999b). To this point in time, most of the labor scheduling literature has focused on T3 using daily (the shift scheduling problem) or weekly (the tour scheduling problem) planning horizons.

In this paper we focused on an aspect of T3 that is important to managers wishing to maximize profitability, but has been overlooked by researchers. Specifically, we developed and evaluated methods of deciding which employees, given their unique productivities and costs, should be assigned to specific work schedules. Essentially, we relaxed the commonly used assumption of employee homogeneity. In the academic literature, some attempts were made to address worker heterogeneity by including differences in productivity levels. However, previous modeling and solution-generating efforts were limited to the categories of full- and part-time employees (Easton and Rossin, 1991a, 1991b; Li, Robinson, and Mabert, 1991), where employees were assumed to be homogenous within these categories (in other words, each full-time employee was assumed to have the same productivity and cost characteristics as all the other full-time employees), or methods for generating solutions were not generalized to individual employees. Since individuals vary in their cost and productivity, in this study we investigated the effect of using this information in order to match service capacity to customer demand more profitability.

2.2. Tour scheduling problem

We used the model introduced by Thompson (1995a) to represent the workforce tour scheduling problem, which is referred to hereafter as the Tour Scheduling Model, or TSM. Compared to earlier representations of the problem (based on the work of Dantzig (1954) and Keith (1979)), TSM finds better workforce schedule for two reasons. First, it relaxes the assumption of inter-period independence that affects earlier models. Second, it considers the non-linear incremental value of work schedules across their whole duration. TSM is:

$$\text{maximize } Z = \sum_{i \in I} d_{ik}\tau_{ik} - \sum_{j \in T} c_j x_j \tag{1}$$

subject to the following constraints:

$$\sum_{j \in T} a_{ij}x_j - \sum_{k=1}^{q_i} \tau_{ik} \geq m_i \text{ for } i \in I \tag{2}$$

$$\tau_{ik} = 0, 1 \text{ for } i \in I \text{ and } k = 1 \ldots, q_i, \tag{3}$$

$$x_j \geq 0 \text{ and integer for } j \in T \tag{4}$$

where relevant terminology is defined in table 1.

INSERT TABLE 1 HERE

TSM's objective is to maximize the long-run benefit (net present value profit) of providing good service. Equation (1) represents the trade-off between the cost of labor and the effect of staffing levels on profits, by formulating the objective function using separate terms for the cost of tours and incremental changes in net present value profit. For each period, equation (2) ensures a minimum reasonable level of staffing and identifies the number of workers scheduled over the minimum reasonable level. As envisioned by Thompson (1995a), the minimum reasonable staff size in a period is the smallest number of staff with an aggregate service rate that exceeds the expected customer arrival rate. The binary variables

τ_{ik}(3) are used in the objective function to determine the corresponding change in net present value profit. Last, equation (4) specifies the integrality of the variables that represent assigned workers.

3. The IWA heuristic methods

This section describes the IWA heuristic methods for post hoc assignment of individuals to tours. In all cases we began with a profit maximizing solution to TSM, where employees were assumed homogeneous. Specific individuals were selected from a fixed pool of workers and assigned to each tour in the schedule using the IWA heuristics. Associated experimental factor levels determined the size of the worker pool. For a base-case application of TSM, we defined H-RAN as a random assignment of individuals to the schedule. This means, then, that H-RAN is equivalent to the majority of the labor scheduling literature, in which no consideration was given to the individuals assigned to tours.

In addition to H-RAN, we examined three IWA heuristic methods that do not require a computer and one simulated annealing heuristic method that requires the use of a computer. We describe them in the following subsections:

3.1. Assigning individuals in Descending Order of Productivity (H-DOP)

H-DOP sorts individuals in descending order of estimated productivity and assigns the workers in order, first to the longer tours and then to the shorter tours. H-DOP's rationale is the common assumption that more productive workers are more beneficial.

3.2. Assigning individuals in Ascending Order of Cost (H-AOC)

H-AOC sorts individuals in ascending order of estimated relative cost and assigns the workers in order, first to the longer tours and then to the remaining tours. H-AOC is based on a low-cost strategy, albeit pursed to the extreme.

3.3. Assigning individuals in Descending Order of Productivity/Cost Ratio (H-DOR)

H-DOR combines H-DOP and H-AOC by considering each individual's estimated productivity per unit cost. This incorporates both common assumptions of high estimated productivity and low cost being beneficial in scheduling/assigning labor. H-DOR sorts the individuals in descending order of their estimated productivity/cost ratio and assigns the workers in order, first to the longer tours and then to the remaining tours.

3.4. Maximizing the net present value profit of the schedule (H-NPV)

H-NPV, the fifth heuristic, randomly takes individuals from the pool and assigns the workers to tours in the solution to TSM. A simulated annealing heuristic (according to the method of Brusco and Jacobs, 1993) is used to find the combination of individuals that maximizes the objective of profit for the entire schedule. The parameters of the simulated annealing algorithm are specified in section 4.7. An explicit description of H-NPV follows:

1. Randomly assign individual workers to the scheduled tours for which no individuals are assigned.
2. Calculate the net present value profit of the schedule. The schedule and the profit in the incumbent solution are retained as the current solution according to the simulated annealing algorithm.
3. Drop workers randomly to perturb the schedule. The number of workers dropped from tour assignments varied randomly from approximately 10–40%.
4. Repeat steps one through three a specified number of times.
5. Repeat steps one through four a specified number of replicates, each time resetting the annealing parameters and clearing all assignments.

The net present value profit estimates are determined by the method of Thompson (1995a) as follows: (1) an estimate of the probability distribution of customer waiting times is generated for each period, and (2) these probability distributions are used as inputs for determining whether customers will be very

INSERT FIGURE 1 HERE

satisfied, satisfied, or dissatisfied for the various waiting times – customer satisfaction is specified by an estimate of the functional relationship between waiting and customer satisfaction. Logistic functions describe these relationships, which is consistent with Davis (1991). Furthermore, expected waiting time in each period is determined using M/M/n queueing formulae with the arrival rate in the period, average service rate of workers scheduled for the period, and the number of workers scheduled in the period. For a full specification of these relationships, see Thompson.

4. Research design

This section presents the full-factorial experiment we used to analyze the effects of the IWA heuristic methods across a broad range of environmental conditions. Figure 1 is the graphical depiction of the experimental process. A set of 2,304 tour scheduling problems with 10 factors was used in the full-factorial design of the experiment. The four IWA heuristics and H-RAN were used to generate employee assignments for each problem, yielding a total of 11,520 observations. The experimental factors and factor levels are given in table 2. The experimental factors are separated into four categories and reviewed in sections 4.2–4.5. Section 4.1 specifies assumptions of our service environment. Sections 4.2

INSERT TABLE 2 HERE

and 4.3 contain explanations of the factors associated with customer arrivals and service providers, respectively. We describe the factors associated with the service transaction and calculating net present value profit in section 4.4. Section 4.5 contains an explanation of the factors associated with accuracy in parameter estimates. We describe how simulated annealing parameters were specified in sections 4.6 and 4.7. Lastly, section 4.8 contains the description of the simulation component of the experiment.

4.1. Assumptions

The operating environment for the study reflected a high contact service operating seven days per week with each day consisting of 18 hour-long planning periods, in which the mean customer arrival rate was assumed to be one customer per minute. Each day begins with no customers in the system and continues until all customers arriving in the 18 hours are served. The workforce was comprised of full- and part-time employees. Full-time employees worked eight hours each day and they received a one-hour break between the fourth and fifth work hours in their shift. Part-time employees worked four-hour shifts with no breaks. All employees worked five consecutive days out of the seven-day week.

We made assumptions regarding individual employee productivity and cost. We assumed that individual employees were from a single population, and thus their productivity levels were normally distributed about a relative mean of 1.0 with a standard deviation of 0.2. This range of productivity variation is consistent with that observed in published studies (Schmidt and Hunter, 1983; Hunter, Schmidt, and Judiesch, 1990). For this experimental design, productivity levels were distributed ±2 standard deviations about the mean. Considering the purpose of this study, we felt that a range of relative productivity from 0.6 to 1.4 was reasonable and conservative.

We assumed the costs of individuals to be approximately proportional to their individual productivity levels. For example, an individual who has a service rate that is 50% higher than the average worker will receive a wage that is approximately 50% more than average. Indirectly this happens when workers have developed a performance history and their value to the organization is reflected in periodic adjustments in their compensation. Labor costs are correlated directly with productivity when workers' compensation is commission based.

We assumed a mean service duration of five minutes per customer. This represented a service transaction like a bank teller, telemarketer or customer service representative. We also adjusted each period's arrival rate to account for the multi-period impact of service, following the method introduced by Thompson (1993).

Customer waiting time was estimated with standard M/M/n queuing formulae by using the average service rate in the period. To justify using average service rates in the queuing calculations, we created a simulation for sets of various individual service rates used in the study and then compared the queue time results to the calculated customer queue time. This sensitivity analysis confirmed that, for this experiment's range and distribution of individual productivity levels, average service rates were a good estimator of waiting time.

4.2. Factors associated with customer arrivals

The experimental design included three levels for the factor of within-day arrival patterns (sine functions). The unimodal pattern (such as might be seen in a grocery store on a weekend), the bimodal pattern (such as might be seen in any number of commute-based services, such as toll booths), and the trimodal pattern (such as might be seen at a restaurant open for breakfast, lunch and dinner) parallel the customer arrivals at various types of organizations. The two levels of daily arrival-rate variation (coefficients of variation of 0.25 and 0.50 in the within-day demand) enabled us to judge the effect of greater and lesser within-day demand variation. We note that the greater the within-day variation in demand, the greater the interdependency of periods, due to the increased magnitude of the service spillover effect (Thompson, 1993).

We selected two factor levels of the across-day arrival-rate patterns. One level had identical relative mean arrival-rates, while the other had relative mean arrival-rates of 1.0, 0.7, 0.7, 1.0, 1.0, 1.3 and 1.3. The latter case is comparable to that seen in service organizations that may be busier on weekends (such as restaurants).

4.3. Factors associated with service providers

Three factors were included for modeling the service provider's environment. The first factor was the size of the worker pool from which the manager assigned individuals to TSM's solution. The two levels for the size of the worker pool were 1.2W and 2.0W, where W equals the number of tours in the best solution to TSM found within the first 50 iterations of the simulated annealing heuristic. To control extraneous variance in the experiment, we used the same worker pool for assigning workers for all IWA heuristic methods under the same factor levels.

The second factor was the relative cost of the individual workers. Based on our assumption of a worker's labor cost being related to his/her true individual productivity level, we created two levels for the degree of the relationship. A worker's relative cost was generated according to a normal distribution of error values about the individual's productivity level, with a coefficient of variation determined by the experimental factor level. The coefficients of variation were 0.05 and 0.20, which respectively represent labor costs that are tightly and loosely correlated with productivity.

Third, we had two levels of shift-length flexibility when solving TSM. The lower flexibility level used only tours made up of eight-hour shifts (i.e., only full-time employees). In contrast, the higher flexibility level used tours comprised of eight- and four-hour shifts (i.e., full- and part-time employees).

4.4. Factors associated with the service transaction and calculating net present value profit

There were three factors that affected the service transactions, which converted the scheduler's decisions into customer satisfaction and net present value profit/loss for the organization. The first factor was the functional relationship between customer waiting time and customer satisfaction. We used logistic functions that were a function of customer waiting time to predict customer satisfaction. We used two levels, gradual and rapid, which describe the rate at which customer dissatisfaction increases along with customer waiting time. These relationships were identical to those used by Thompson (1995a).

A customer's satisfaction or dissatisfaction has a benefit or cost, respectively, to the organization. The second experimental factor was the magnitude of the increase in net present value profit (one or five labor-hour equivalents) as a result of a satisfied customer, while the third experimental factor was the net present value loss (one or five labor-hour equivalents) due to a dissatisfied customer. The level of these costs in real firms vary based on a variety of factors, including the nature of the business and the competitive environment.

4.5. Factors associated with the accuracy in parameter estimates

Managers must make staffing decisions based on historical data and projections about what will happen in the future. To test the robustness of the IWA heuristics, we included a factor with three levels of predictable variation about the expected values of key parameter estimates. We considered three levels of information accuracy (uncertainty): perfect information, high accuracy and low accuracy. Table 3 lists the effects on the key parameter estimates under the three levels of information accuracy.

4.6. Simulated annealing parameters for solving TSM

Prior to conducting our experiment, we performed tests to set the parameters for the simulated annealing heuristic method we used to solve TSM. We found that 1,200 iterations was a reasonable run length since we observed that the objective function hit plateaus by that point. With longer runs we observed only very small incremental increases in the objective and relatively long time intervals between finding new best solutions. We also ran the heuristic five times for each problem, selecting the best solution we found across all five runs. Again, in our initial testing we observed only a very slight improvement in the objective as we considered more than five replicates. With the five runs of 1,200 iterations, then, our best schedule to TSM was based on a total of 6,000 iterations of the simulated annealing heuristic.

The temperature parameter and reduction constant in the simulated annealing algorithm were chosen based on the first fifty iterations of the heuristic for each problem. We set the initial probability of accepting an incumbent solution to 0.50, for the mean difference between the profitability of incumbent and current solutions. We adjusted the temperature and cooling rate so that this same profitability differential would have a 0.05 probability of being accepted at the end of the 1,200 iterations.

4.7. Simulated annealing parameters for assigning individuals to tours

We determined the simulated annealing parameters for assigning individual employees to tours in a manner similar to finding the parameters for solving TSM described in section 4.6. Again, we found that 1,200 iterations with five replicates, or 6,000 iterations was successful in finding good solutions before additional iterations made only relatively small improvements to the "best" profit value. The temperature parameter and reduction constant in the simulated annealing algorithm were chosen using the same cooling conditions described in section 4.6.

INSERT TABLE 3 HERE

4.8. The simulation component

The simulation component tested the schedules generated by the heuristic procedure by implementing the schedules in a simulated service operation (with Markov distribution assumptions for service and inter-arrival times in each planning period) for 50 simulated weeks. Using the net present value of customer satisfaction, this stochastic process yielded a value for net present value profit that is comparable across each IWA heuristic method. Across the worker assignment methods, the customer arrivals to the service operation were generated using a stream of random numbers, which continued from the random number stream (and corresponding seed) used in determining the mean arrival rates for the tour scheduling component. In addition, across the worker assignment methods, employee productivity and cost parameters were created in advance and fixed for the duration of the experimental run of each combination of factor level. Our simulation of 50 weeks of operation did not mean to imply that a service organization would use the same schedule for 50 weeks; rather it provided us with a reasonable estimate of the true profitability of implementing the schedule for one week.

5. Results

Duncan's multiple-range test was used to analyze the results of the experiment. Table 4 displays a summary of the results and the percentage improvement of each heuristic when compared to H-RAN based on the mean weekly net present value profit. H-NPV consistently outperformed all of the other heuristics. Furthermore, at the factor level of low profit, the H-NPV solutions were significantly different from the next best heuristic (H-DOR) at the 0.05 level. For the factor level of high profit, the H-NPV solutions were not significantly different from the H-DOR solutions; however, the H-NPV and H-DOR

INSERT TABLE 4 HERE

---

heuristic solutions were significantly different from the next best heuristic at the 0.05 level.

When comparing the H-NPV and H-DOR heuristics to the baseline H-RAN, the differences for all four Profit/Loss combinations were statistically significant, and relatively substantial for the factor levels indicating low profit values for very satisfied customers (perhaps a very competitive environment with small margins). The percentage improvement of H-NPV over H-RAN was 3.23% and 2.71% for low profit/high loss and low profit/low loss, respectively.

The difference for the factor level of high profit was a relatively small percentage, but the value in absolute terms may be substantial. For example, for one operation in a service organization, a 0.53% (high profit/high loss) improvement on profit values of 33,082.48 man-hour equivalents is 176.82 manhours. If labor rates are estimated at $10.00/hour, then the increase in profit is $1768.20 per week or $91,946.40 per year. This increase in profits is for one operation. Many large retailers, call centers, and telecommunications companies have regional or multiple sites of customer service representatives. So even small, statistically significant differences could yield substantial increases in profits.

Table 5 presents a summary of the heuristics' operating performance. H-NPV (simulated annealing) yielded the best solution on 82.3% of the problems. H-DOR (managerial heuristic) yielded the best solution on 10.7% of the problems. So the H-NPV and H-DOR heuristics accounted for 2,141, or 93.0%, of the best solutions, and they were consistently first and second, respectively, in yielding the highest profit. So in summary, the difference between the performance of the H-NPV and H-DOR

INSERT TABLE 5 HERE

---

heuristics was relatively small according to table 4; however, as seen in table 5, the difference in computation time is very large. To yield its modest profitability advantage over H-DOR, H-NPV requires large numbers of iterations. Since H-DOR heuristic is almost as good as H-NPV and since it could be easily implemented without a computer by managers, it could prove a useful practical heuristic for managers.

5.1. Factors associated with customer arrivals

For factors associated with customer arrivals, the results clearly indicated that the performance of H-NPV and H-DOR were most comparable when higher variability existed. Specifically for variation in daily customer arrival-rate pattern (amplitude of sine function), statistically insignificant and smaller differences were seen for the smaller amplitude in the customer arrival-rate pattern, that is, the factor level of the coefficient of variation equal to 0.25. Under these conditions, the improvement that H-DOR provided over the baseline H-RAN ranged from 0.46% under the factor level of high profit/high loss to 2.7% under the factor level of low profit/high loss.

In contrast, a higher amplitude (coefficient of variation equal to 0.50) of the customer arrival-rate pattern led to a distinct difference between the H-NPV and H-DOR heuristics. Under these conditions, H-NPV was superior. The reason that H-NPV performed significantly better than H-DOR may have been due to the period to period interdependency – and H-NPV's better ability to capture its effects – caused by the large differences in period to period changes in customer arrivals. Recall that we accounted for the spill-over effect by adjusting the customer arrivals (Thompson, 1993). Improvements that H-NPV provided over the baseline H-RAN ranged from 0.51% under the factor level of high profit/low loss to 3.3% under the factor level of low profit/high loss.

5.2. Factors associated with service providers

For a small worker pool, improvements that H-NPV provided over the baseline H-RAN ranged from 0.34% under the factor level of high profit/low loss to 2.1% under the factor level of low profit/high loss. For the large worker pool, improvements that H-NPV provided over the baseline H-RAN ranged from 0.69% under the factor level of high profit/low loss to 4.4% under the factor level of low profit/high loss. With a large worker pool size, all differences between H-NPV and H-DOR were statistically insignificant. This indicated that H-DOR may work comparably to H-NPV.

The differences between solutions generated by both H-NPV and H-DOR over H-RAN across variation in the shift length were statistically significant, but relatively small in size. The flexibility gained by using four- and eight-hour shifts compared to only eight-hour shifts was reflected in these significant differences.

5.3. Factors associated with the accuracy in parameter estimates

Across various levels of environmental uncertainty, again, H-NPV and H-DOR performed comparably. As expected, the distinction between schedules generated with H-NPV and H-DOR decreased as uncertainty increased. With perfect information, the differences between H-NPV and H-DOR were statistically significant. Under conditions of low accuracy, the difference between H-NPV and H-DOR became statistically insignificant. Table 6 presents a matrix of the percentage of solutions where each heuristic yielded the higher profit when compared to every other heuristic. As can be seen, the dominance of H-NPV and H-DOR dissipated as environmental uncertainty increased.

6. Discussion and conclusions

We hypothesized that IWA heuristic methods for assigning workers with individual productivity levels to schedules would differentially affect profitability. By considering individuals, a manager can better match service capacity with customer demand. We identified efficient heuristics for assigning individual workers to labor tour schedules. Secondly, we determined how much managers can expect to increase the profitability of the tour schedule by using the best-performing IWA heuristic methods.

We examined four heuristics – one based on simulated annealing and three simple heuristics that would not require the use of a computer. In addition, we established baseline performance using post hoc random assignment of individuals to tours. The results indicated that two heuristics dominated. H-NPV and H-DOR were always the first- and second-best performing heuristics, respectively, when the results were broken out for each factor separately. Generally H-DOR distinguished itself from the next-best performing heuristic by statistical significance and a substantial increase in profit level. Upon inspection of the results, H-NPV and H-DOR performed comparably in magnitude of profits (table 4) and often were indistinguishable based on their performance when good information was utilized (table 6).

INSERT TABLE 6 HERE

However, H-DOR took a fraction of the computation time expended by H-NPV (table 5). In addition, the ease of implementing H-DOR makes it a useful practical heuristic for managers.

Comparing the profitability of H-NPV and H-DOR to baseline H-RAN provided information about how they would work in practical settings. As table 4 shows, under conditions of low profit levels (e.g., competitive environments with small profit margins), H-NPV and H-DOR increased profits by an average 3.0% and 2.4%, respectively, over the baseline H-RAN.

Further benefits in profit levels were generated for particular conditions. For low profit levels and a large worker pool, H-NPV and H-DOR provided an average 4.0% and 3.5% increase in profits, respectively, over the baseline H-RAN. Similarly, under low profit levels and high variation in the relative cost of individual workers, H-NPV and H-DOR provided an average 4.0% and 3.7% increase in profits, respectively, over the baseline H-RAN. Of course, managers have many considerations when assigning individuals, however, within the freedom at their disposal for assigning employees to tours, managers should use H-NPV and H-DOR in lieu of inferior assignment methods.

Our findings clearly show that considering productivity and costs when assigning employees to tours does matter and it matters to both research and to practice. First, it matters to practice in that we have found an easy and low cost method of improving profitability. Second, it matters to research since this issue has not been examined, but it is one that can improve the profitability of real organizations.

Extensions of this work include incorporating additional real-world constraints into the problem of assigning individuals; for example, considering different productivity distributions. Also, improving the methodology to consider scheduling tours and individuals concurrently instead of sequentially is a logical extension of this work.

References

Baker, K.R. (1976). "Workforce Allocation in Cyclical Scheduling Problems: A Survey." *Operational Research Quarterly* 27(1), 155–167.

Beaumont, N. (1997). "Scheduling Staff Using Mixed Integer Programming." *European Journal of Operational Research* 98(3), 473–484.

Bechtold, S.E., M.J. Brusco, and M.J. Showalter. (1991). "A Comparative Evaluation of Labor Tour Scheduling Methods." *Decision Sciences* 22(4), 683–699.

Bechtold, S.E. and G.M. Thompson. (1993). "Optimal Scheduling of a Flexible-Duration Rest Period for `a Work Group." *Operations Research* 41(6), 1046–1054.

Brusco, M.J. and L.W. Jacobs. (1993). "A Simulated Annealing Approach to the Cyclic Staff-Scheduling Problem." *Naval Research Logistics* 40, 69–84.

Brusco, M.J. and L.W. Jacobs. (1998a). "Eliminating Redundant Columns in Continuous Tour Scheduling Problems." *European Journal of Operational Research* 111(3), 518–525.

Brusco, M.J. and L.W. Jacobs. (1998b). "Personnel Tour Scheduling when Starting-Time Restrictions Are Present." *Management Science* 44(4), 434–447.

Brusco, M.J., L.W. Jacobs, and G.M. Thompson. (1999). "A Morphing Procedure to Supplement a Simulated Annealing Heuristic for Cost- and Coverage-Correlated Set-Covering Problems." *Annals of Operations Research* 86, 611–627.

Brusco,M.J. and T.R. Johns (1998). "Staffing a Multi-skilled Workforce with Varying Levels of Productivity: An Analysis of Cross-Training Policies." *Decision Sciences* 29(2), 499–516.

Campbell, G.M. (1999). "Cross-Utilization of Workers whose Capabilities Differ." *Management Science 45*(5), 722–732.

Caprara, A., P. Toth, and M. Fischetti (2000). "Algorithms for the Set Covering Problem." *Annals of Operations Research* 98, 353–371.

Dantzig, G.B. (1954). "A Comment on Edie's "Traffic Delays at Toll Booths"." *Operations Research* 2(3), 339–341.

Davis, M. (1991). "How Long Should a Customer Wait for Service?" *Decision Sciences* 22(2), 421–434.

Dowling, D., M. Krishnamoorthy, H. Mackenzie, and D. Sier. (1997). "Staff Rostering at a Large International Airport." *Annals of Operations Research* 72, 125–147.

Easton, F. and N. Mansour. (1999). "A Distributed Genetic Algorithm for Deterministic and Stochastic Labor Scheduling Problems." *European Journal of Operational Research* 118(3), 505–523.

Easton, F.F. and D.F. Rossin. (1991a). "Equivalent Alternate Solutions for the Tour Scheduling Problem." *Decision Sciences* 22(5), 985–1001.

Easton, F.F. and D.F. Rossin. (1991b). "Sufficient Working Subsets for the Tour Scheduling Problem." *Management Science* 37(11), 1441–1451.

Emmons, H. and D.S. Fuh. (1997). "Sizing and Scheduling a Full-Time and Part-Time Workforce with Off-Day and Off-Weekend Constraints." *Annals of Operations Research* 70, 473–492.

Goodale, J.C. and E. Tunc. (1997). "Tour Scheduling with Dynamic Service Rates." *International Journal of Service Industry Management* 9(3), 226–247.

Goodale, J.C., R.V. Verma, and M.E. Pullman. (2003). "A Market Utility-Based Model for Capacity Scheduling in Mass Services." *Production and Operations Management*, to appear.

Hunter, J.E., F.L. Schmidt, and M.K. Judiesch. (1990). "Individual Differences in Output Variability as a Function of Job Complexity." *Journal of Applied Psychology* 75, 28–42.

Keith, E.G. (1979). "Operator Scheduling." *AIIE Transactions* 11(1), 37–41.

Li, C., E.P. Robinson, Jr., and V.A. Mabert. (1991). "An Evaluation of Tour Scheduling Heuristics with Differences in Employee Productivity and Cost." *Decision Sciences* 22(4), 700–718.

Schmidt, F.L. and J.E. Hunter. (1983). "Individual Differences in Productivity: An Empirical Test of Estimates Derived from Studies of Selection Procedure Utility." *Journal of Applied Psychology* 68(3), 407–414.

Thompson, G.M. (1992). "Improving the Utilization of Front-Line Service Delivery System Personnel." *Decision Sciences* 23(5), 1072–1098.

Thompson, G.M. (1993). "Accounting for the Multi-Period Impact of Service when Determining Employee Requirements for Labor Scheduling." *Journal of Operations Management* 11(3), 269–287.

Thompson, G.M. (1995a). "Labor Scheduling Using NPV Estimates of the Marginal Benefit of Additional Labor Capacity." *Journal of Operations Management* 13(1), 67–86.

Thompson, G.M. (1995b). "Improved Implicit Optimal Modeling of the Labor Shift Scheduling Problem." *Management Science* 41(4), 595–607.

Thompson, G.M. (1998a). "Labor Scheduling, Part 1: Forecasting Demand." *Cornell Hotel and Restaurant Administration Quarterly* 39(5), 22–31.

Thompson, G.M. (1998b). "Labor Scheduling, Part 2: Knowing How Many On-Duty Employees to Schedule." *Cornell Hotel and Restaurant Administration Quarterly* 39(6), 26–37.

Thompson, G.M. (1999a). "Labor Scheduling, Part 3: Developing a Workforce Schedule." *Cornell Hotel and Restaurant Administration Quarterly* 40(1), 86–96.

Thompson, G.M. (1999b). "Labor Scheduling, Part 4: Controlling Workforce Schedules in Real Time." *Cornell Hotel and Restaurant Administration Quarterly* 40(3), 85–96.

Table 1
List of variables and constants.

x_j	= the number of employees assigned to work schedule j;
c_j	= the cost of assigning an employee to work schedule j;
a_{ij}	= 1, if period i is a work period of schedule j; 0, otherwise;
r_i	= the number of workers required in period i (in full-time equivalents);
j	= an index for work schedules;
i	= an index for planning intervals;
I	= the set of planning intervals;
T	= the set of work schedules;
τ_{ik}	= 1, if the number of employees working in period i equals or exceeds $m_i + k$; 0, otherwise;
k	= the number of additional workers in period i beyond the minimum acceptable level;
q_i	= the number of employees in period i in excess of the minimum reasonable staff size, who, ignoring labor costs, contribute to increased NPV profits (i.e., $d_{iq} > 0$);
d_{ik}	= the incremental improvement in NPV profit (ignoring labor costs) that occurs with the addition of the $(m_i + k)$th employee in period i (assuming nonincreasing marginal NPV returns for each period – i.e., $d_{ik} \geqslant d_{i,k+1}$ for $i \in I$ and for all j);
m_i	= the minimum reasonable number of workers for period i.

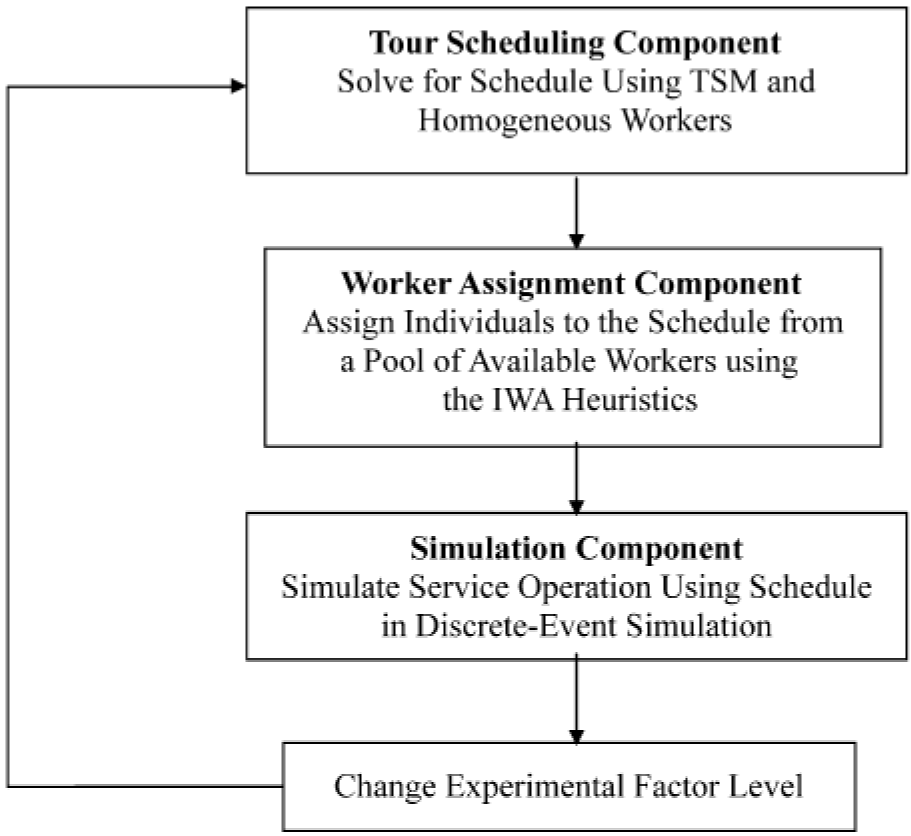

Figure 1. Flow diagram of experimental process.

Table 2
Experimental factors.

Factor (# of levels)	Unit of measure	Levels
Worker assignment methods (5)	Heuristic	H-AOC, H-DOP, H-DOR, H-NPV and H-RAN
Within-day customer arrival-rate pattern (3)	A sinusoidal pattern	Unimodal, bimodal and trimodal
Variation in the weekly customer arrival-rate pattern (2)	Relative mean arrival rates for days in the week	Constant and nonconstant (see section 4.2)
Variation in the daily customer arrival-rate pattern (2)	Coefficient of variation of customer arrival-rate	0.25 and 0.50; sine function with amplitudes of 0.353 and 0.706
Worker pool size (2)	Number of workers available for assignment	1.2 and 2.0 times the number of workers used in the best solution of the first 50 tour schedules
Relative cost (2)	Coefficient of variation of individual's productivity	0.05 and 0.20
Shift length (2)	Hours per shift	(1) All 8-hour shifts or (2) 4 and 8-hour shifts possible
True functional relationship between customer satisfaction and waiting time (2)	Curve relating probability of customer being satisfied with waiting time	Gradual and rapid (see section 4.4)
Relative NPV profit from a very satisfied customer (2)	Labor-hour equivalents	1 and 5
Relative NPV loss from a dissatisfied customer (2)	Labor-hour equivalents	1 and 5
Uncertainty (3)	Accuracy in parameter estimates of unpredictable arrival-rate variation, productivity estimates, profit/loss estimates, and estimate of functional relationship between waiting and satisfaction	Perfect information, high accuracy, and low accuracy
Dependent variable	Net present value profit (in labor-hour equivalents)	

Table 3
Information items and their values under the levels of the information accuracy factor levels.

Key experimental parameter	Factor level		
	Perfect	High accuracy	Low accuracy
Estimated customer arrival rate, by period (COV compared to actual customer arrival rate, by period)	0.0	0.05	0.20
Estimated individual productivity level (COV compared to actual individual productivity level)	0.0	0.05	0.20
Estimated cost of a dissatisfied customer (COV compared to actual cost of a dissatisfied customer)	0.0	0.05	0.20
Estimated value of satisfied customer (COV compared to actual value of a satisfied customer)	0.0	0.05	0.20
Functional relationship between customer waiting time and customer satisfaction (number of sampled customers)	1,000	100	50

Table 4
Summary of simulation results and percentage improvement compared to H-RAN.*

Heuristic	High profit high loss	High profit low loss	Low profit high loss	Low profit low loss
H-NPV[a]	$33{,}259.30^{A}$ (+0.53%)	$33{,}359.14^{A}$ (+0.51%)	$5{,}159.50^{A}$ (+3.23%)	$5{,}369.99^{A}$ (+2.71%)
H-DOR[b]	$33{,}228.04^{A}$ (+0.44%)	$33{,}332.41^{A}$ (+0.43%)	$5{,}128.88^{B}$ (+2.61%)	$5{,}342.34^{B}$ (+2.18%)
H-DOP[c]	$33{,}118.06^{B}$ (+0.11%)	$33{,}207.01^{B}$ (+0.05%)	$5{,}031.64^{C}$ (+0.67%)	$5{,}232.21^{C}$ (+0.07%)
H-RAN[d]	$33{,}082.48^{B}$	$33{,}188.87^{B}$	$4{,}998.25^{D}$	$5{,}228.51^{C}$
H-AOC[e]	$32{,}905.22^{C}$ (−0.54)	$33{,}080.90^{C}$ (−0.33%)	$4{,}820.00^{E}$ (−3.57%)	$5{,}161.33^{D}$ (−1.28%)

* Statistically significant groupings for Duncan's multiple-stage test indicated by A, B, C, D, and E superscripts, where $n = 576$ and $\alpha = 0.05$.
[a] Maximize net present value profit of schedule (Simulated Annealing).
[b] Descending order of the productivity/cost ratio (Managerial Heuristic).
[c] Descending order of productivity (Managerial Heuristic).
[d] Random (Baseline Scenario).
[e] Ascending order of cost (Managerial Heuristic).

Table 5
Operating performance of heuristics.

Heuristic	Number (%) of problems in which the heuristic yielded the highest profit over each other heuristic ($n = 2304$)					Number (%) highest overall	Heuristic CPU[a] time (sec.)
	H-NPV	H-DOR	H-DOP	H-RAN	H-AOC		
H-NPV	–	1,942 (84.3%)	2,074 (90.0%)	2,215 (96.1%)	2,302 (99.9%)	1,895 (82.3%)	163.16
H-DOR	362 (15.7%)	–	2,035 (88.3%)	2,240 (97.2%)	2,230 (96.8%)	246 (10.7%)	0.07
H-DOP	230 (10.0%)	269 (11.7%)	–	1,201 (52.1%)	1,365 (59.2%)	162 (7.0%)	0.07
H-RAN	89 (3.9%)	64 (2.8%)	1104 (47.9%)	–	1,338 (58.1%)	1 (0.0%)	0.05
H-AOC	2 (0.1%)	74 (3.2%)	939 (40.8%)	967 (41.9%)	–	0 (0.0%)	0.06

[a] Computer: Gateway2000 P5.

Table 6
Performance of heuristics across environmental uncertainty.

Heuristic	Number (%) of problems in which the heuristic yielded the highest profit over each other heuristic ($n = 768$)					Number (%) highest overall
	H-NPV	H-DOR	H-DOP	H-RAN	H-AOC	
Perfect information						
H-NPV	–	713 (92.8%)	764 (99.5%)	767 (99.9%)	768 (100%)	712 (92.7%)
H-DOR	55 (7.2%)	–	742 (96.6%)	753 (98.0%)	732 (95.3%)	53 (6.9%)
H-DOP	4 (0.5%)	26 (3.4%)	–	257 (33.5%)	354 (46.1%)	3 (0.4%)
H-RAN	1 (0.1%)	15 (2.0%)	511 (66.5%)	–	379 (49.2%)	0 (0.0%)
H-AOC	0 (0.0%)	36 (4.7%)	414 (53.9%)	390 (50.8%)	–	0 (0.0%)
High accuracy						
H-NPV	–	725 (94.4%)	757 (98.6%)	768 (100%)	768 (100%)	720 (93.8%)
H-DOR	43 (5.6%)	–	742 (96.6%)	746 (97.1%)	735 (95.7%)	43 (5.6%)
H-DOP	11 (1.4%)	26 (3.4%)	–	279 (36.3%)	382 (49.7%)	2 (0.7%)
H-RAN	0 (0.0%)	22 (2.9%)	490 (63.8%)	–	403 (52.5%)	0 (0.0%)
H-AOC	0 (0.0%)	33 (4.3%)	386 (50.3%)	365 (47.5%)	–	0 (0.0%)
Low accuracy						
H-NPV	–	504 (65.6%)	553 (72.0%)	680 (88.5%)	766 (99.7%)	463 (60.3%)
H-DOR	264 (34.4%)	–	551 (71.7%)	741 (96.5%)	763 (99.3%)	150 (19.5%)
H-DOP	215 (28.0%)	217 (28.3%)	–	665 (86.6%)	629 (81.9%)	154 (20.1%)
H-RAN	88 (11.5%)	27 (3.5%)	103 (13.4%)	–	556 (72.4%)	1 (0.1%)
H-AOC	2 (0.3%)	5 (0.7%)	139 (18.1%)	212 (27.6%)	–	0 (0.0%)