Representing Employee Requirements in Labor Tour Scheduling

Gary M. Thompson

University of Utah

July 1993

Address for Correspondence:

Professor GM Thompson,

Department of Management,

University of Utah, Salt

Lake City, UT 84112, USA.

Abstract

In this paper, we use the methodology of simulation to evaluate six approaches for handling employee requirements In an LP-based labour tour scheduling heuristic. We model employee requirements both as minimum acceptable staffing levels-where understaffing is unacceptable-and as target staffing levels-where both under- and overstaffing are acceptable. For each representation of employee requirements, we evaluate forms of the heuristic that use problem-specific and problem-independent information on the costs of employee surpluses and, if appropriate, employee shortages. Over an extensive test data set, the target-staffing approach using problem-specific cost Information outperformed all other procedures. Specifically, it generated schedules costing less than 87% of those developed using the approach most commonly found in the literature. Its schedules were also almost 5% cheaper than those of its closest competitor. We discuss the managerial and research implications of the findings and provide suggestions for future research.

Keywords - scheduling, manpower planning, LP, simulation

Representing Employee Requirements in Labor Tour Scheduling

1. Introduction

Services represent a major component of the economy in many countries. Hence, attaining high service productivity is of broad concern. Labour scheduling, which is frequently a significant determinant of a service organization's efficiency, has received a good deal of attention in the literature. Its basic aim is ensuring that enough appropriately-skilled employees are present, at the times needed, to provide the level of service specified by management. Increasing the difficulty of the task is customer demand that exhibits wide temporal variation and management's inability to inventory customer- staff interaction activities.

The labour scheduling literature has used two methods of matching the number of employees working to the number of employees needed. Of these, the far more common method treats employee requirements as a lower bound, by prohibiting understaffing. The less common method treats the employee requirements as a target, and in doing so allows both employee shortages and surpluses. We choose to call these types of employee requirement restrictions 'at-least' and 'target' staffing requirements, respectively.

In this paper, we use the methodology of simulation to evaluate six approaches for handling employee requirements in an LP-based labour tour scheduling heuristic. These approaches arise from both at-least and target staffing requirements combined with problem-specific and problem-independent costs associated with having too many or too few employees on hand. Problem-independent costs are typically constant across and within periods across problems, while problem-specific costs typically vary across problems and across and within periods within problems. In conducting the simulation experiment, we consider 576 environmental scenarios. These scenarios arise from two replications of each combination of three across-day and three within-day customer arrival patterns, four mean service times, two

curves linking customer dissatisfaction to customer waiting times, and four relative costs of customer dissatisfaction.

We consider a scheduling environment typical of those found in service organizations employing full-time staff working regular schedules. We assume the facility operates 20 h daily and that planning occurs using hour-long periods. Allowable weekly work schedules, or tours, have (1) 5 working days, (2) consecutive or non-consecutive days-off, (3) shifts of 9 h, with an hour-long break taken during the fifth hour and (4) identical starting times for all shifts. These restrictions yield a total of 252 unique tours.

The remainder of the paper presents relevant background material on labour scheduling, describes the experimental design, defines the modelling approaches, presents the results of the simulation experiments, discusses the managerial implications of the study, and offers suggestions for future research.

2. LABOUR SCHEDULING BACKGROUND

Dantzig [12] provided the first mathematical programming formulation of the labour scheduling problem. His model, which we call Ml, took a typical set-covering approach to the problem by using at-least staffing requirements. Baker [2] proposed modifying M 1 by allowing both under- and overstaffing, but suggested limiting each period's maximum under- and overstaffing. If care is not exercised when setting the bounds on the under- and overstaffing variables in such a model, however, no feasible solution would exist [16]. To avoid this potential difficulty, each period should have both bounded and unbounded under- and overstaffing variables [16]. Keith's [16] formulation of the labour scheduling problem, which we call M2, is the universal tour scheduling model used in this paper. Appendix A presents M2. In applying M2, the unbounded employee shortage and surplus variables have higher costs than the corresponding bounded variables. Thus, a model like M2 can incorporate

more accurate cost information than Ml can [2]. Also, solutions to M2 tend to distribute any shortage or surplus of employees within the limits specified by the bounds on under- and overstaffing.

Table 1 summarizes the type of cost information and employee requirement modelling approaches used in the labour scheduling literature. Most of the research has considered the employee requirements as predetermined, and in doing so has used problem-independent costs associated with employee surpluses and, if appropriate, employee shortages. Unfortunately, these problem-independent costs may bear little relationship to the true costs of employee shortages and surpluses. One goal of this paper is to compare the effectiveness of problem-independent and problem-specific costs in tour scheduling. Table l also shows the predominant use of the at-least approach to modelling employee requirements. Our other goal is to compare the effectiveness of at-least and target staffing approaches.

3. EXPERIMENTAL DESIGN

Within this section we list the assumptions of the study, identify and describe the experimental factors and provide details on the simulation.

Assumptions

To clarify scope of this study, it may be helpful to list our assumptions explicitly: (l) the labour pool is unlimited (labour staffing is not addressed); (2) the employees are homogeneously- skilled; (3) the employees work as scheduled (there is no absenteeism, for example); (4) no breaks, other than the hour-long meal periods, are either taken or scheduled; (5) the mean service rate is constant (employees do not work faster or slower when the facility is busy, for example); (6) the mean weekly number of customer arrivals is constant over the duration of the simulation period; (7) a customer is either satisfied or unsatisfied with the service based on the time he/she spends waiting for service; (8) the likelihood of a customer

being dissatisfied with any given waiting time is constant for the duration of the simulation; (9) the organization loses future sales for each customer who is dissatisfied with the service, but this decrease in sales occurs beyond the horizon of the simulation, and (10) there is no balking from the single, first-in-first-out queue.

Experimental factors

To evaluate the modelling approaches' performance, we varied five environmental factors potentially influencing their relative performance: (1) the customer arrival pattern (two factors); (2) the mean customer service time (one factor); (3) the relative cost of customer dissatisfaction (one factor); and (4) the functional relationship between the probability of customer dissatisfaction and customer waiting time (one factor). Table 2 summarizes the environmental factors, which the following subsections describe in detail.

Customer arrival patterns. We generated nine distinct customer arrival patterns from the combinations of three within-day (IND) and three across-day (ACR) patterns. Within-day customer arrival variation occurs as the underlying (true) customer arrival rate changes over the operating day. The three in-day arrival patterns had one, two and three daily arrival peaks. In contrast to within-day variation, across-day variation occurs as the mean daily customer arrival rate changes across the operating week. There were three levels of ACR: the low-variation ACR pattern had equal mean customer arrival rates on all days; while the medium-variation (high-variation) ACR pattern had relative mean daily customer arrival rates of l, 3/4, 3/4, 1, 1, 5/4 and 5/4 (1, 1/2, 1/2, 1, 1, 3/2 and 3/2) for Sunday-Saturday, respectively. All combinations of ACR and IND had the same mean weekly average of 48 customer arrivals per hour.

Customer inter-arrival times followed an exponential distribution (the models describing the arrival processes are available from the author). Figure 1 illustrates a typical example of each weekly pattern. Although the illustrated arrival patterns are not smooth and

although all the days do not exhibit their characteristic shapes, the patterns would become more clearly defined should customer arrivals be averaged over multiple weeks.

Service times. We altered the mean service time (TIM) to measure the effect of increasing the required staff size, while using the same sequence of random numbers in the simulation. Actual service times followed an exponential distribution because this distribution reflects the natural variability in customer service and inter-arrival times often seen in organisations. The four levels of TIM had mean service times of 2.5, 5, IO and 20 minutes per customer.

Costs of poor service. In this study we assumed that there were only two relevant costs: the cost of service delivery labour and the cost of poor service. For all investigations, the labour cost per employee per hour was normalized to unity, an action consistent with earlier research [3, 6, 7, 10, 12, 14, 15, 18, 23, 26, 28]. To develop a cost of poor service, we assumed that a customer's waiting time affected his satisfaction with the service. We also assumed that the organization would lose future contribution if a customer became dissatisfied with the service.

Two customer dissatisfaction curves (CDCs) depicted in Fig. 2 specified the functional relationship between customer waiting time and customer satisfaction. Both curves have 50% of customers dissatisfied with a 2 min wait for service. Appendix B provides the mathematical forms of these curves.

There were four, widely-varying levels of opportunity costs (CST) associated with customer dissatisfaction: net-present-values of 0.3125, 1.25, 5 and 20, measured in labour-hour equivalents (LHEs). Obviously, an organization that considers the long-term applications of poor service will provide better customer service than one that has only a short-term focus. Given labour scheduling's short-term decision scope, then, it is particularly important to use net-present- value costs of customer dissatisfaction. The respective levels of CST justify a

manager scheduling an additional employee to a 40- working-hour tour if, on average, the employee could prevent at least 128, 32, 8 and 2 customers from becoming dissatisfied.

Simulation details

To keep the computational requirements within reason, the 9 weekly customer arrival patterns were each replicated only twice. For each replication, the data on customer arrival times and the random numbers used for calculating customer service times and for determining customer satisfaction were stored in a data file. To control variance, schedule development and evaluation for every combination of TIM, CDC, CST and modelling approach (MDL) used this stored data. Thus, each arrival pattern replication yielded 192 data points, for a total of 3456 (= 192 x 18) observations.

We generated 30 weeks of customer data for each arrival pattern replication. Ten 'historical' weeks provided an average number of customer arrivals in each of the 140 planning periods. These data enabled us to first set the employee requirements, as described in the next section, and then develop a labour schedule.

Twenty 'future' weeks of data allowed a simulation of the service facility with the labour schedule in effect. Each customer's 'actual' waiting time yielded the probability of his being dissatisfied. The appropriate stored random number determined whether the customer would be satisfied, given their probability of dissatisfaction. Twenty-week averages of the labour and customer dissatisfaction costs together gave the total schedule cost. The twenty-week evaluation period is used only to give a good measure of the true average weekly cost of the schedule; it does not imply that a service organization should use the same schedule for 20 weeks.

In trial runs, how well the models performed relative to each other was insensitive to the random number seeds used in generating the customer characteristics (inter-arrival and service times and probability of dissatisfaction). Because of this, we judged the 20 weeks of simulated future operation to be of acceptable duration. The simulation model itself was coded

in FORTRAN. Completing the experiment required over 400 h on an 486DX33-based personal computer.

4. TOUR SCHEDULING MODELLING APPROACHES

We investigate six tour scheduling modelling approaches, designated MDLI-MDL6. As Table 3 shows, MDLI and MDL2 use at-least staffing, while MDL3-MDL6 use target staffing. MDLI, MDL3, MDL4 and MDL5 follow most of the literature in using problem-independent costs of employee shortages and surpluses. In contrast, MDL2 and MDL6 use problem-specific cost information, both in the tour scheduling LP model and in subsequent heuristic schedule improvement.

The rest of this section addresses the setting of employee requirements, provides details on the tour scheduling model forms, and describes the heuristic schedule improvement activities.

Setting the employee requirements

Each of the 140 weekly planning periods had its employee requirement set independently. For each period, we identified the staff size yielding the lowest total cost for the period (consisting of labour costs and customer dissatisfaction costs, with the latter described in Appendix B). As Table 4 shows, the total expected cost for a period varies non-linearly with respect to the number of employees available during the period. Based on this cost information, the optimal staffing levels for periods one through three are 9, 15 and 13 employees, respectively. We define the minimum acceptable number of employees for a period as the smallest staff size for which the aggregate service rate exceeds the period's average historical customer arrival rate. For example, the minimum acceptable staff sizes for periods one through three in Table 4 are 5, 10 and 8 employees, respectively.

Tour scheduling LP model form

The following subsections identify the variables included in each MDL and describe the procedures used in setting the bounds on and the costs of these variables.

Variables. Table 3 summarizes the employee shortage and surplus variables used in each modelling approach. All MDLs allow unbounded surplus staffing variables. MDL2 and MDL6 include bounded surplus variables since these models use problem-specific information on the costs of employee shortages and surpluses.

MDLI and MDL2 include neither type of shortage variable, since both take an at-least approach to modelling employee requirements. MDL3-MDL6 include bounded shortage variables since these models take a target approach to modelling employee requirements.

Bounds and costs of the variables. In all modelling approaches, the costs of the unbounded surplus variables per employee-hour were equal to one LHE (from the normalized labour cost per period). As Table 3 shows, the problem-independent costs of employee shortages in MDL3-MDL5 were 2.5, 5 and 10 LHEs, respectively. The bounds on understaffing in MDL3-MDL5 were such that no fewer than the minimum acceptable number of employees would be present in each period.

In four of the six MDLs, the model structure is the same for both the LP and heuristic phases. MDL2 and MDL6, which use problem-specific cost information, are the exceptions. The reason for the difference in model structures is that although a heuristic can readily incorporate non-linear under- and overstaffing costs, an LP model cannot. Thus, the LP versions of MDL2 and MDL6 attempt to approximate the true, non-linear, problem-specific cost information using bounded under- and overstaffing variables. Clearly, the rationale for using bounded variables is that small changes from the desired number of staff (those within the bounds) are likely to be less costly than dramatic deviations (those outside the bounds) [16]. Table 4's cost data bear this out. As the staff size increases beyond the size yielding the minimum cost for the period, waiting times decrease and the lowered cost of customer

dissatisfaction partly offsets the increased labour cost. However, as the number of employees continues to grow, the savings resulting from the lower customer dissatisfaction costs become insignificant. Decreasing the staff size has the opposite effect. Initially the savings from having fewer staff partially offset the increased customer dissatisfaction costs. One reaches a point, though, when further decreases in the number of staff dramatically increase customer waiting times and customer dissatisfaction costs.

The point where adding an additional employee resulted in more than a 0.95 unit increase in the total relevant cost (in LHEs) provided the bound on surplus staffing in the LP versions of MDL2 and MDL6. This is where we judged the marginal reduction in customer dissatisfaction resulting from an additional employee to be essentially insignificant. Consider period one from Table 4, for example. Having 9 employees on-hand yields the period's lowest total cost. As the number of employees incrementally increases from 9 to 14, the period's total relevant costs increase by 0.51, 0.85, 0.94, 0.98 and 1.00 LHEs, respectively. Thus, the bound on overstaffing would be set at three employees-the third incremental employee satisfies the criterion of no more than a 0.95 unit increase in the period's relevant costs, but the fourth employee does not.

For understaffing, the point where employee shortages became too costly imposed the bound on employee shortages in the LP version of MDL6. We decided that this would be where a further reduction in the number of staff resulted in more than a five-unit increase in the period's total relevant cost, measured in LHEs. (Obviously, the bound would always have to be small enough to ensure that the minimum acceptable number of employees would be present in the period.) Consider, for example, period three from Table 4. Here, the lowest cost arises with 13 employees, while the minimum acceptable staff size is 8 employees. As the staff size incrementally decreases from 13 to 8 employees, the period's relevant costs increase by 1.40, 4.64, 11.71, 26.31 and 55.33 LHEs, respectively. Thus, period three's bound on understaffing

is set at two employees-any further staffing decreases raise relevant costs by more than the acceptable level.

Because the costs of increasing or decreasing the staff size from the desired level are nonlinear, we calculated an average linear cost for the bounded shortage and surplus variables used in the LP versions of MDL2 and MDL6. Based on the cost and staffing data in Table 4, Table 5 provides examples of the limits on and costs of bounded shortage and surplus staffing variables. To see how to determine these costs, consider our earlier example of overstaffing in period one. With three extra employees, the total relevant cost is 2.32 LHEs higher than the lowest possible value, which linearly is 0. 77 LHEs per employee.

Heuristic schedule improvement

In heuristically modifying solutions to the relaxed LP tour scheduling models, we followed the general procedure of [16], supplemented with additional schedule-improvement actions based on a vector-exchange heuristic [14]. We selected Keith's [16] heuristic for two reasons. First, this heuristic performed best in a recent evaluation of various tour scheduling heuristics [7]. Second, one can readily adapt the heuristic for both at-least and target staffing approaches.

Schedule modification begins by rounding fractional variables in the LP tour model solution to the nearest integer. The heuristic next adds employees to or drops employees from the schedule, if beneficial, and then undertakes a second improvement phase. We describe each action more fully below.

Adding employees to tours. To investigate the benefit of adding employees to the schedule, the heuristic examines all alternate tours. MDLI and MDL2 (MDL3-MDL6) iteratively add an employee to the tour making the greatest reduction in total understaffing until satisfying the at-least (minimum acceptable) staffing requirements. All the modelling approaches break ties by choosing to add an employee to the tour yielding the minimum sum of squares of overstaffing. Based on Table 4's data, for example, MDLI and MDL2 (MDL3-

MDL6) would add employees to tours until at least 9 (5), 15 (IO) and 13 (8) employees were present in periods one-three, respectively.

Having satisfied the minimum acceptable employee requirements for all periods, MDL3- MDL5 each iteratively adds employees to tours if this lowers the schedule cost. As described earlier, employee shortages below and surpluses above the target staffing level have relative costs of 2.5:1 in MDL3, 5:1 in MDL4 and I0:1 in MDL5. Thus, beneficial tours in MDL3, MDL4 and MDL5 reduce the total shortage by at least 12, 7 and 4 employee-periods, respectively.

After satisfying the minimum acceptable employee requirements for all periods, MDL6 uses the complete, problem-specific cost information to find the net benefit of adding more employees to tours. To illustrate how MDL6 identifies the value of adding employees to tours, consider again Table 4's data. Now assume that there are respectively 7, 15 and 15 employees scheduled to work in periods one-three. If one adds an employee to a tour covering these periods, periods one-three will have 8, 16 and 16 employees working, respectively. The relevant costs will consequently be lower in period one (from 13.23 to 10.17) but higher in periods two (from 16.22 to 16.47) and three (from 15.20 to 16.07). Since the complete effect will be to reduce relevant costs by l.94 LHEs, it helps to add an employee to the tour.

Improving the schedule, part I-dropping employees from tours. The first attempt at improving the schedule seeks to drop employees from non-beneficial tours. All MDLs use the same criteria for selecting a tour from which to remove an employee: drop an employee from the tour yielding the greatest improvement in schedule cost and break ties by dropping an employee from the tour yielding the greatest improvement in the sum of the squared overstaffing. In MDLI and MDL2 (MDL3-MDL6), tours are only candidates for having employees removed if the number of employees currently scheduled exceeds the at-least (minimum acceptable) staffing levels in all of the tour's working periods.

To evaluate the schedule cost change resulting from dropping an employee from a tour, MDL2 and MDL6 use the exact, problem-specific cost information. In contrast, MDLI and MDL3- MDL5 use the problem-independent cost information to make the evaluation. As such, the value of primary criterion will be identical for all tours from which one can remove an employee in MDLI, while MDL3, MDL4 and MDL5 will drop an employee from any tour that does not increase the total understaffing by more than 11, 6 and 3 employee-periods (LHEs), respectively. To see how MDL2 and MDL6 evaluate dropping employees from tours, again consider the information in Table 4, but now assume that periods one-three respectively have 11, 16 and 15 employees scheduled to work. Also assume that one is considering dropping an employee working a tour covering periods one and two, or dropping an employee working a tour covering periods two and three. One can remove either employee, since every period in both tours has more than the target number of employees scheduled to work. However, we drop an employee working the first tour because doing so lowers total relevant costs the most (I.IO LHEs vs 0.88 LHEs for dropping an employee working the second tour).

Improving the schedule, part II-drop two employees and add a third. The second attempt at improving the schedule, which is based on the improvement actions of a vector-exchange heuristic [14), works by dropping employees from two tours and adding an employee to a third tour. When dropping employees from tours, the heuristic uses the criteria identified in the previous section, but now the heuristic can tentatively drop an employee from any tour, even if doing so means that some period(s) will have fewer than the at-least (MDLI and MDL2) or minimum acceptable (MDL3- MDL6) number of staff on-hand. After tentatively dropping two employees from the schedule, the heuristic adds an employee to a third tour using the same criteria as used in the initial adding of employees to tours. If changing the schedule lowered its cost, and if each period has no fewer than the at-least (MDLI and MDL2) or the minimum acceptable (MDL3-MDL6) number of staff on-hand, the tentative changes to the schedule become permanent and the process repeats.

5. RESULTS

Table 6 summarizes the experimental results. From high to low total cost, the modelling approaches were MDLI, MDL2, MDL3, MDL5, MDU and MDL6. In percentage terms, MDL6's schedules were over 13% cheaper than those of MDLI, and just under 5% cheaper than those of its closest competitor, MDU, differences significant at the $\alpha = 0.0005$ level.

Table 6 also reports each modelling approach's mean time required to generate tour schedules. The values are the seconds required to perform all actions associated with developing the schedule. In order of least to greatest total time, the modelling approaches were MDLI, MDL5, MDL4, MDL3, MDL2 and MDL6. Since the reported times are for a 486DX33-based personal computer and since the longest time is just over 6 min, none of the time requirements should be viewed as burdensome.

Figures 3 and 4 graphically present comparative schedule cost results. The vertical axis in these figures is the schedule cost savings offered by modelling approaches MDL2-MDL6. These figures express the savings as a percentage of the cost of MDLI's schedules. Figure 3 presents the results by levels of ACR, CDC and IND, while Fig. 4 presents the results by levels of TIM and CST. Figure 3 shows the performance ordering, with MDL6 clearly superior to the next best model (MDL4), holds across the levels of ACR, IND and CDC. MDL6's percentage advantage over MDLI declines as the duration of customer service increases, as shown in Fig. 4. Also to be noted in Fig. 4 is that MDL3 and MDL4 did very well compared to MDL6 with a relative customer dissatisfaction cost (CST) below 20 LHEs.

Figure 5 illustrates the minimum percentage cost advantages of MDL6's schedules for the levels of TIM and CDC (ACR and IND), by level of CST. The vertical axis in this figure is relative savings (or costs) of MDL6's schedules. The figure expresses these savings (or costs) as a percentage of the best other (non-MDL6) scheduling approach's schedule costs. An important point is that the best other scheduling approach is context specific. For example, the

best other scheduling approach when TIM = 2.5 min and CST = 20 LHEs, is not necessarily the same as the best other scheduling approach when TIM= 20 min and CST= 0.3125 LHEs. Figure 5 clearly shows (1) that MDL6's advantage was greatest with the highest customer dissatisfaction cost and (2) that MDL6 was slightly outperformed when the cost of a dissatisfied customer equaled 1.25 LHEs (results were similar for the levels of ACR and IND by level of CST). MDL6 was broadly superior to the other models over all other two-way combinations of the experimental factors' levels.

6. DISCUSSION

Within this section we address the managerial and research implications of the study and suggest research extensions.

Managerial and research implications of the findings

This section discusses the superiority of the target staffing approach, the superiority of problem-specific costs of under- and overstaffing, the impact of heuristic solutions on the results, the occurrence of multiple optimal schedules, and the use of the preferred modelling approach when insufficient information is available.

Superiority of the target staffing approach. Our first major finding is the general superiority of the target staffing approach, as reflected by all four target-based modelling approaches (MDL3-MDL6) performing better, on average, than either of the at-least-based modelling approaches (MDLI and MDL2). The general superiority of the target staffing approach relates to the nature of employee requirements. Recall that the desired staff size in a period is the number of employees yielding the minimum total relevant costs for the period. Clearly, then, there would be no difference between the at-least and target staffing approaches when the desired number of employees were present in all periods. This is commonly not possible because of management policy and employee desires regarding acceptable work

schedules. When the desired staffing levels cannot be exactly satisfied in all periods, the target approach to staffing offers the flexibility of having fewer than the optimal (for a single period) number of staff on hand in a period. The target approach to staffing also enables a clearer evaluation of the tradeoff between increased labour costs and lowered customer dissatisfaction costs.

The difference in performance between the target and at-least modelling approaches illustrated in Figs 3 and 4 is greater than an astute (and patient) manager using the at-least modelling approach might expect. An astute user of the at-least staffing approach recognizes that having extensive periods of overstaffing (but no understaffing) is intuitively of greater cost than having fewer than the ideal number of staff in some periods and greatly reduced overstaffing. Hence, the astute user of the at-least staffing approach will use a trial-and-error method: solve the scheduling problem, look at the distribution of surplus/shortage staffing, and adjust the employee requirements in the at-least model. This process repeats until the balance of under-and overstaffing satisfies the manager, or until the manager tires of the process. What is this manager, in effect, doing? Merely implementing a target approach to staffing, albeit in a cumbersome way. Clearly, then, another advantage of using a target staffing modelling approach is its directness.

As evidenced by the literature summary provided in Table 1, researchers have used the at-least approach to staffing restrictions much more commonly than they have used the target staffing approach. Given the superiority of the target staffing approach, researchers should ensure that any optimal or heuristic procedures they develop can use it. In addition, comparisons of technique performance should be undertaken primarily using a target staffing approach. The implication for managers from this discovery is clear: better schedules can be more readily obtained using a target staffing approach.

An important observation is the failure of the target-based modelling approaches to totally dominate the at-least approaches. MDLI and MDL2 both performed better than MDL3

under the highest customer dissatisfaction cost (CST = 20). This result is not surprising when one recollects the relationship between MDL3's problem-independent costs of under- and overstaffing (a 2.5: l ratio). Based on the results shown in Figs 3 and 4, a manager choosing to use the target-staffing approach and who wishes to use problem-independent costs of under- and overstaffing is probably best advised to select a higher relative cost of understaffing than has generally been used in the literature (see Table 1 for historical values). A danger in using problem-independent costs is that the target-staffing modelling approach will always be outperformed by the at-least approach, if the cost of customer dissatisfaction is high enough. Such a concern does not apply for a target staffing approach that uses problem-specific costs. The implication of this is the heading of the next section.

Superiority of problem-specific costs of under and overstaffing. Our second major finding is that the use of accurate cost data in the target staffing approach (MDL6) yielded markedly lower cost schedules than any of the other modelling approaches. Overall, MDL6 generated schedules costing 86.6% of those generated by the most common approach in the literature, MDLI. MDL6 also notably outperformed the next-best target staffing approach, MDL4. MDL6's schedules were almost 5% less costly than those of MDL4, on average. The likely reason for MDL6 outperforming the other target staffing approaches is that MDL6 accurately identifies the tradeoffs between labour and customer dissatisfaction costs, while MDL3-MDL5 only approximate it. This point becomes particularly relevant when one examines the comparative results presented in Fig. 4. With the highest customer dissatisfaction cost (CST= 20 LHEs), MDL5 did better than either MDL3 or MDL4. When the cost of a dissatisfied customer was intermediate (CST = 1.25 or 5 LHEs), MDL4 did better than either MDL3 or MDL5. Finally, with a low customer-dissatisfaction cost (CST= 0.3125 LHEs), MDL3 did better than either MDL4 or MDL5. These results are fully consistent with the relationships between the problem-independent costs of employee shortage and surpluses used in MDL3 (2.5:1), MDL4 (5:1) and MDL5 (10:1). Obviously, using problem-independent costs

of employee shortages and surpluses limits the ability of a model to work well over a broad range of environmental conditions. If one counters with the argument that a manager could choose between MDL3, MDL4 and MDL5 based on the cost of poor service, then one is really advocating the use of problem-specific costs of employee shortages and surpluses. This is exactly the reason for using MDL6. The broad superiority of MDL6 (and the performance ordering of the other target-based modelling approaches) convincingly shows that the labour scheduling process should use the information of the relative costs of different staffing levels generated as a byproduct of the process of setting employee requirements. Since the vast majority of published scheduling research cannot, without modification, use such cost information (see Table 1), there is much room for improvement.

How heuristic solutions affect the results. It is useful to consider the potential impact of the heuristic nature of the modelling approaches. Although it is possible that optimal procedures would yield different results, there are several reasons why this is not a pressing concern. First, all the approaches we evaluate follow the general heuristic procedure of [16], which a large comparative study of tour scheduling heuristic performance has shown to work very well consistently [7]. Second, if the number of test cases is small, or if the test cases are peculiar in some regard, a heuristic that is effective over a broad range of problems may be outperformed by one that is not. Because we simulated 576 different environments, representing what we feel is a broad range of realistic environmental characteristics, we are confident in the validity of our results. Third, managers and researchers often use heuristic solution procedures because of the difficulty in solving tour problems to optimality. Thus, when evaluating modelling approaches, it makes sense to report the results for heuristic, rather than optimal procedures. Finally, it is likely that MDL6 would also prove to be the best optimal modelling approach. This is because MDL6's superior performance comes from its ability to represent the true tradeoffs between labour costs and the costs of poor service-something that is duplicated in none of the other approaches.

Occurrence of multiple optimal schedules. Researchers have commonly observed that labour scheduling problems frequently exhibit multiple optimal solutions. Multiple optima are desirable from management's perspective, since there are likely to be additional considerations in choosing a schedule beyond those incorporated in a scheduling model or procedure. Multiple optima are much more likely to occur when using costs of under- and overstaffing that are problem-independent and constant across and within periods. With realistic cost data, under- and overstaffing costs are non-linear and vary across periods. For example, a surplus employee would typically vary in value across the weekly planning horizon, as Table 4 shows. Without multiple optima (when using problem-specific costs), a manager can better judge the true impact of accommodating other, non-quantitative objectives.

Using MDL6 with insufficient information. A practical concern for managers (similar to one raised by [2]) is, "How should employee staffing levels be set without knowledge of the functional form of customer dissatisfaction, and without knowledge of the average cost of dissatisfying a customer?" There are three options available to managers in such cases. The first is to translate the organization's existing customer service policy into a form that MDL6 can use. Doing this might entail transforming the organization's service policy into an implied customer dissatisfaction curve and then estimating the implicit cost of customer dissatisfaction from the current staffing levels in the organization. The disadvantage of this approach is that it assumes that the current customer service policy and current staffing levels are indeed appropriate.

A second option would be to adapt MDL6 for use with a specified service level (such as "serve x percent of customers within time y"). There are two disadvantages of this alternative. First, it presumes the customer service policy is appropriate. Second, without knowledge of the cost of customer dissatisfaction, one has two measures of how good a schedule is (customer service and schedule cost) instead of one (total cost) and so MDL6 would require substantial modification.

The third, and preferable choice, is to get the necessary knowledge, since this is the only alternative that does not presume that an organization's current service level and staffing policies are appropriate. Excellent examples of how to acquire the necessary information exist [l, 13]. Given the degree of uncertainty that exists in most organizations regarding the costs of poor customer service, it is perhaps safer to err on the side of conservatism. This is because better service, although more costly in the short-term, may yield increased future revenue [13]. Moreover, since managers must often estimate this information, they should perform extensive sensitivity analyses [13].

Research extensions

Compared to that existing in many service organizations, the tour scheduling environment used in this paper had a low degree of scheduling flexibility (as measured, for example, by the range of times at which shifts could start, the length of shifts, and the flexibility in scheduling breaks). In comparison to published scheduling research, on the other hand, the environment was not too restrictive. An interesting study would examine the relative performance of the modelling approaches under varied levels of scheduling flexibility. Obviously, if the degree of scheduling flexibility was extensive, then it may be possible to match the number of employees scheduled exactly to the number of employees needed. Overstaffing, measured as a percentage of total labour requirements, was shown to decline substantially in a tour scheduling environment as the flexibility in scheduling employees to shorter shifts and fewer days per week increased [19]. However, even in a very flexible scheduling environment, where employees could work 4-h shifts on as few as 3 days a week, surplus staffing was more than 6% of total labour requirements [19]. We expect, therefore, that even in this very flexible environment the target staffing approach would still outperform the at-least staffing approach. Indeed, the target staffing approach offers a form of flexibility that can complement the other forms of flexibility identified in the literature.

A second investigation relates to the assumption of demand stationarity. Relaxing this assumption would require the use of a more complex process for identifying the desired staffing levels. Obviously, any analysis would be more realistic if: (1) good (poor) service raised (lowered) customer demand; and (2) the simulation accounted for employees failing to perform as scheduled.

Finally, the LP implementation of MDL6 may benefit from further refinement. This issue is particularly relevant given MDL6's performance with an intermediate cost of customer dissatisfaction (CST = 1.25 LHEs). Although MDL6 was far superior in this case to the approach most common in the literature (MDL1), MDL6 was slightly inferior to MDL3. We think the reason for this may lie in how the LP versions of MDL3 and MDL6 impose the bounds on understanding. Recall that MDL3 (and MDL4 and MDL5) imposes the understaffing bound so that no fewer than the minimum- acceptable number of employees will be working each period. MDL6 also uses this bound during heuristic schedule improvement. In the LP model, however, MDL6's understaffing bound is set at the point where having one fewer employee working in a period would result in a 5-unit increase in the period's total relevant costs (in LHEs). MDL6's tighter bound on understaffing may undesirably restrict the range of solutions that the LP model can examine (and that the heuristic schedule improvement cannot escape from). Additional research should clarify this matter. Specifically, refinement efforts should address the setting of understaffing bounds and methods for incorporating non-linear costs of under- and overstaffing in MDL6's LP incarnation.

Appendix A

This appendix presents the relevant mathematical formulation of the labour tour scheduling problem, M2. Begin by defining the following:

Variables

t_n = is the number of employees working a tour of type *n*

bu_p = the bounded shortage (underage) of employees in period p

uu_p = the unbounded shortage (underage) of employees in period *p*

bo_p = the bounded surplus (overage) of employees in period *p*

uo_p = the unbounded surplus (overage) of employees in period *p.*

Data and constants

a_{np} = 1 if period p is a working period for tour *n*

= 0 otherwise

c_n = the cost of tour *n*

$c1_p$ = the cost of the unbounded understaffing variable (per employee short) for period *p*

$c2_p$ = the cost of the bounded understaffing variable (per employee short) for period *p*

$c3_p$ = the cost of the bounded overstaffing variable (per surplus employee) for period *p*

$c4_p$ = the cost of the unbounded overstaffing variable (per surplus employee) for period *p*

N = the set of unique tours that can be scheduled

P = the set of planning periods in the week

r_p = the number of employees needed in period *p* to provide the desired level of customer service

ubu_p = the upper limit on the bounded understaffing variable for period *p*

ubo_p = the upper limit on the bounded overstaffing variable for period *p.*

M2, the model of [16], is:

$$\min Z = \sum_{n \in N} c_n t_n + \sum_{p \in P} [c1_p uu_p + c2_p bu_p + c3_p bo_p + c4_p uo_p] \quad \text{(A1)}$$

subject to

$$\sum_{n \in N} a_{np} t_n + uu_p + bu_p - bo_p - uo_p = r_p \text{ for } p \in P \quad \text{(A2)}$$

$$bu_p \leq ubu_p \text{ for } p \in P \quad \text{(A3)}$$

$$bo_p \leq ubo_p \text{ for } p \in P \quad \text{(A4)}$$

$$uu_p, bu_p, bo_p, uo_p \geq 0 \text{ and integer for } p \in P \quad \text{(A5)}$$

M2 has as its objective (A1) the minimization of total tour costs and the cost of deviations from the target staffing levels. Constraint set (A2) measures the deviations from each period's target staffing levels. Constraint sets (A3) and (A4), respectively, impose the limits on the variables measuring the bounded employee shortages and surpluses. Constraint set (A5) imposes the non-negativity and integrality of the employee shortage and surplus staffing variables while constraint set (A6) does the same for the tour variables.

Appendix B

This appendix provides the formulas for determining the probability of a customer's dissatisfaction with the service and each period's expected customer dissatisfaction cost. With the gradual dissatisfaction curve, a waiting time for service of w_q yields a probability of dissatisfaction, k, of:

$$k(W_q) = 1/[1 - 54.5986 * \exp(-2W_q)], \quad \text{(B1)}$$

while the corresponding probability for the rapid dissatisfaction curve is:

$$k(W_q) = 1/[1 - 22026.5 * \exp(-5W_1)]. \quad \text{(B2)}$$

In our study, a period's expected customer dissatisfaction cost in labor hour equivalents (LHEs) is approximately:

$$TDC_{ps} = CST * AA_p * \left[\sum_{j=1}^{50} P(0.1[j-1] \leq W_q \leq 0.1j) * k(0.1[j-0.5])\right] \quad \text{(B3)}$$

where

TDC_{ps} = the total dissatisfaction cost for period p with s employees working in the period

AA_p = the average number of historical customer arrivals in period p

CST = the cost of a dissatisfied customers, in LHEs (0.3125, 1.25, 5 or 20)

$P(x \leq W_q \leq y)$ = the probability of a customer's waiting time falling between *x* and *y* minutes

(from an *M/M/s* queueing model, using the average historical customer arrival rate for the period, and where *s* = the number of employees working in the period) $k(x)$ = the probability of a dissatisfied customer, given a wait of *x* minutes [from equation (B1) or (BE)].

References

1. Andrews BH and Parsons HL (1989) L. L. Bean chooses a telephone agent scheduling system. *Interfaces* 19(6), 1-9,

2. Baker KR (1976) Workforce allocation in cyclical scheduling problems: a survey. *Opl Res. Q.* 27(1), 155-167.

3. Baker KR, Crabill TB and Magazine MJ (1973) An optimal procedure for allocating manpower with cyclic requirements. *AIIE Trans*. 5(2), 119-126.

4. Bailey J (1985) Integrated days off and shift personnel scheduling. *Computers Ind. Engng* 9(4), 395-404.

5. Bailey J and Field J (1985) Personnel scheduling with flex shift models. *J. Ops Mgrat* 5(3), 327-338.

6. Bartholdi JJ (1981) A guaranteed-accuracy round-off algorithm for cyclic scheduling and set covering. *Ops Res.* 29(3), 501-510.

7. Bechtold SE, Brusco MJ and Showalter MJ (1991) A comparative evaluation of labour tour scheduling methods. *Decn Sci*. 22(4), 683-699.

8. Bechtold SE and Jacobs LW (1990) Implicit optimal modeling of flexible break assignments in labour staffing decisions for service operations. *Mgmt Sci.* 36(11), 1339-1351.

9. Bechtold SE and Jacobs LW (1991) Improvement of labour utilization in shift scheduling for services with implicit optimal modeling. *Int. J. Ops Prodn Mgmt* 11(2), 54-69.

10. Bechtold SE and Showalter MJ (1987) A methodology for labour scheduling in a service operating system. *Decn Sci.* 18(1), 89-107.

11. Buffa ES, Cosgrove MJ and Luce BJ (1976) An integrated work shift scheduling system. *Decn Sci*. 7(4), 620-630.

12. Dantzig GB (1954) A comment on Edie's 'Traffic delays at toll booths.' *Ops Res*. 2(3), 339-341.

13. Davis M (1991) How long should a customer wait for service? *Decn Sci*. 22(2), 421-434.

14. Henderson WB and Berry WL (1976) Heuristic methods for telephone operator shift scheduling: an experimental analysis. *Mgmt Sci.* 22(12), 1372-1380.

15. Henderson WB and Berry WL (1977) Determining optimal shift schedules for telephone traffic exchange operators. *Decn Sci*. 8(2), 239-255.

16. Keith EG (1979) Operator scheduling. *AIIE Trans*. 11(1), 3741.

17. Li C, Robinson EP Jr and Mabert VA (1991) An evaluation of tour scheduling heuristics with differences in employee productivity and cost. *Decn Sci*. 22(4), 700-718.

18. Loucks JS and Jacobs FR (1991) Tour scheduling and task assignment of a heterogeneous work force: a heuristic approach. *Decn Sci.* 22(4), 719-738.

19. Mabert VA and Showalter MJ (1990) Measuring the impact of part-time workers in service organisation. *J. Ops Mgmt* 9(2), 209-229.

20. Mabert VA and Watts CA (1982) A simulation analysis of tour-shift construction procedures. *Mgmt Sci.* 28(5), 520-532.

21. McGinnis LF, Culver WD and Deane RH (1978) One and two-phase heuristics for workforce scheduling. *Computers Ind. Engng* 2(1), 7-15.

22. Moondra SL (1976) An L. P. model for work force scheduling for banks. *J. Bank Res*. 7(4), 299-301.

23. Morris JG and Showalter MJ (1983) Simple approaches to shift, days-off and tour scheduling problems. *Mgmt Sci.* 29(8), 942-950.

24. SAS (1985) *SAS Users Guide: Statistics*, Version 5. SAS Institute Inc., Cary, NC.

25. Segal M (1974) The operator-scheduling problem: a network-flow approach. *Ops Res.* 22(4), 808-823.

26. Showalter MJ and Mabert VA (1989) An evaluation of a full-/part-time tour scheduling methodology. *Int. J. Ops Prodn Mgmt* 8(7), 54-71.

27. Thompson GM (1990) Shift scheduling when employees have limited availability: an L. P. approach. *J. Ops Mgmt* 9(3) 352-370.

28. Thompson GM (1992) Improving the utilization of front-line service delivery system personnel. *Decn Sci.* 25(5) 1072-1098.

Table 1. Classification of relevant labor scheduling literature.

Employee requirements modelling approach	The process by which the employee staffing levels are determined[a]	Objective cost[b]			Solution procedure[d] and references
		Employee shortage[c]	Employee surplus[c]	Shifts or tours	
At-least	Predetermined			No. of tours	CFS[3],H[7,10],LP[6,7,23]
At-least	Predetermined			Cost of tours	LP[17,19,25][e]
At-least	Predetermined			No. of shifts	LP[12,14,15][f]
At-least	Predetermined			Cost of shifts	LP[5,22],IP[8,9,28]
At-least	Predetermined			No. of hours	LP[26]
At-least	Predetermined		1		H[18]
Target	Predetermined	2.5		Cost of tours	LP[4]
Target	Predetermined			Cost of tours	LP[20][g]
Target	Predetermined	300	200	Cost of shifts	LP[16][h]
Target	x% served within time y	1	1		H[11]
Target	Predetermined	Non-linear	Non-linear		H[21][i]
Target	Predetermined	4	1		LP[27]

[a]'Predetermined' indicates that the labour scheduling procedure does not use any of the cost information relating to staffing levels that are generated during the setting of employee requirements.
[b]Missing values indicate the component did not appear in the objective function.
[c]Unless noted otherwise, the costs are linear as a function of the number of employees short or extra.
[d]CFS = close form solution, LP = linear-programming based heuristic, H = heuristic (but not LP-based), IP = integer programming.
[e]Although Segal's [25] basic model used at-least staffing, in a modification he allowed bounded understaffing to occur.
[f]Henderson and Berry [15] developed an LP-based optimal branch and bound procedure.
[g]Although Mabert and Watts' [20] model has the general form of a target staffing approach, it differs from other models in that understaffing can occur because work may be inventoried.
[h]Keith [16] also allowed bounded shortage and surplus variables, each having relative costs of 100.
[i]McGinnis *et al.* [21] developed and evaluated several heuristics, each of which used different, primarily non-linear objectives.

Table 2. Experimental factors.

Factor	Factor abbreviation	Number of levels	Unit of measure	Levels
Within-day variation in the customer arrival rate	IND	3	A pattern	Unimodal (1 peak), bimodal (2 peaks) and trimodal (3 peaks)
Across-day variation in the customer arrival rate	ACR	3	A pattern, based on the relative mean customer arrival rates on Sunday–Saturday	Constant (1,1,1,1,1,1,1), low variation (1,0.75,0.75,1,1,1.25,1.25) and high variation (1,0.5,0.5,1,1,1.5,1.5)
Mean duration of customer service	TIM	4	Minutes	2.5,5,10 and 20
Relative cost of dissatisfying a customer	CST	4	Cost in number of labour-hour equivalents	0.3125,1.25,5,20
Customer dissatisfaction curve	CDC	2	A function relating a customer's probability of dissatisfaction to the time spent waiting for service	Rapid (quick to become dissatisfied) and gradual (slow to become dissatisfied)
Tour scheduling model	MDL	6	A variant of Keith's [16] tour scheduling model (M2)	MDL1,MDL2,MDL3, MDL4,MDL5 and MDL6

Table 3. A summary of the variables and the costs of the variables included in the tour scheduling LP models for six modelling approaches.

Modelling approach identifier	Target (T) or at-least (A) staffing	Variables included in the employee requirement restrictions (cost of the variable in the objective function)			
		Employee shortage		Employee surplus	
		Unbounded	Bounded	Bounded	Unbounded
MDL1	A	N	N	N	Y (1)[a]
MDL2	A	N	N	Y (varies)[b]	Y (1)[a]
MDL3	T	N	Y (2.5)	N	Y (1)[a]
MDL4	T	N	Y (5.0)	N	Y (1)[a]
MDL5	T	N	Y (10)	N	Y (1)[a]
MDL6	T	N	Y (varies)[c]	Y (varies)[b]	Y (1)[a]

[a]Constant across periods.
[b]Varies across periods, approximating the true cost of small surpluses of employees.
[c]Varies across periods, approximating the true cost of small shortages of employees.

Table 4. Examples of the effect of increasing the number of service personnel on a planning period's total relevant costs

Period											
1	No. of servers[a]	5	6	7	8	9[c]	10	11	12	13	14
	Total cost[b]	46.67	22.67	13.23	10.17	9.70	10.21	11.06	12.02	13.00	14.00
2	No. of servers[a]	10	11	12	13	14	15[c]	16	17	18	19
	Total cost[b]	87.35	47.50	28.59	20.25	17.04	16.22	16.47	17.17	18.06	19.02
3	No. of servers[a]	8	9	10	11	12	13[c]	14	15	16	17
	Total cost[b]	113.93	58.60	32.29	20.58	15.94	14.54	14.57	15.20	16.07	17.02

[a]The left-most value for each period is the smallest possible staffing level that can serve all customers (average arrival rate < aggregate service rate).
[b]Total relevant per-period costs with the given staffing level, in labour-hour equivalents (LHEs).
[c]At-least staffing level for MDL1 and MDL2 (target staffing level for MDL3–MDL6).

Table 5. An example of the limits on and the costs of the bounded under- and overstaffing variables MDL2 and MDL61

Period	Desired staffing level[b]	Bounded shortage[c]		Bounded surplus[d]	
		Limit	Cost	Limit	Cost
1	9	3	4.32	3	0.77
2	15	3	4.12	3	0.61
3	13	2	3.02	3	0.51

[a]Based on the cost data in Table 4.
[b]At-least staffing level in MDL1 and MDL2, target staffing level in MDL3–MDL6.
[c]Used in MDL6.
[d]Used in both MDL2 and MDL6.

Table 6. A summary comparison of the mean schedule coasts and mean schedule generation times for the modelling approaches

Modelling approach	MDL1	MDL2	MDL3	MDL5	MDL4	MDL6
Mean schedule cost	2515.13	2498.85	2443.47	2415.67	2288.88	2177.90
Mean schedule cost as a percentage of the mean schedule cost for MDL1	100.00	99.36	97.15	96.05	91.00	86.59
Grouping[a,b]	—	—	—	—	—	—
Modelling approach	**MDL1**	**MDL5**	**MDL4**	**MDL3**	**MDL2**	**MDL6**
Mean schedule generation time (in s)[c]	211.23	230.76	261.43	280.67	292.31	371.45
Mean schedule generation time as a percentage of the mean schedule generation time for MDL1	100.00	109.25	123.77	132.87	138.38	175.85
Grouping[a,d]	—	—	—	—	—	—

[a]From the results of the Ryan–Einot–Gabriel–Welsch [24] multiple range test. This test controls the maximum experiment-wise type I error rate under any complete or partial null hypothesis.

[b]Means linked by lines are not significantly different at the 0.0005 level.

[c]Average total time required on a 486DX33-based personal computer to generate, solve and interpret the linear programming tour model and to subsequently heuristically modify the schedule. LP models solved using SAS-OR [24].

[d]Means linked by lines are not significantly different at the 0.0001 level. (All modelling approaches are significantly different from each other at the 0.0001 level.)

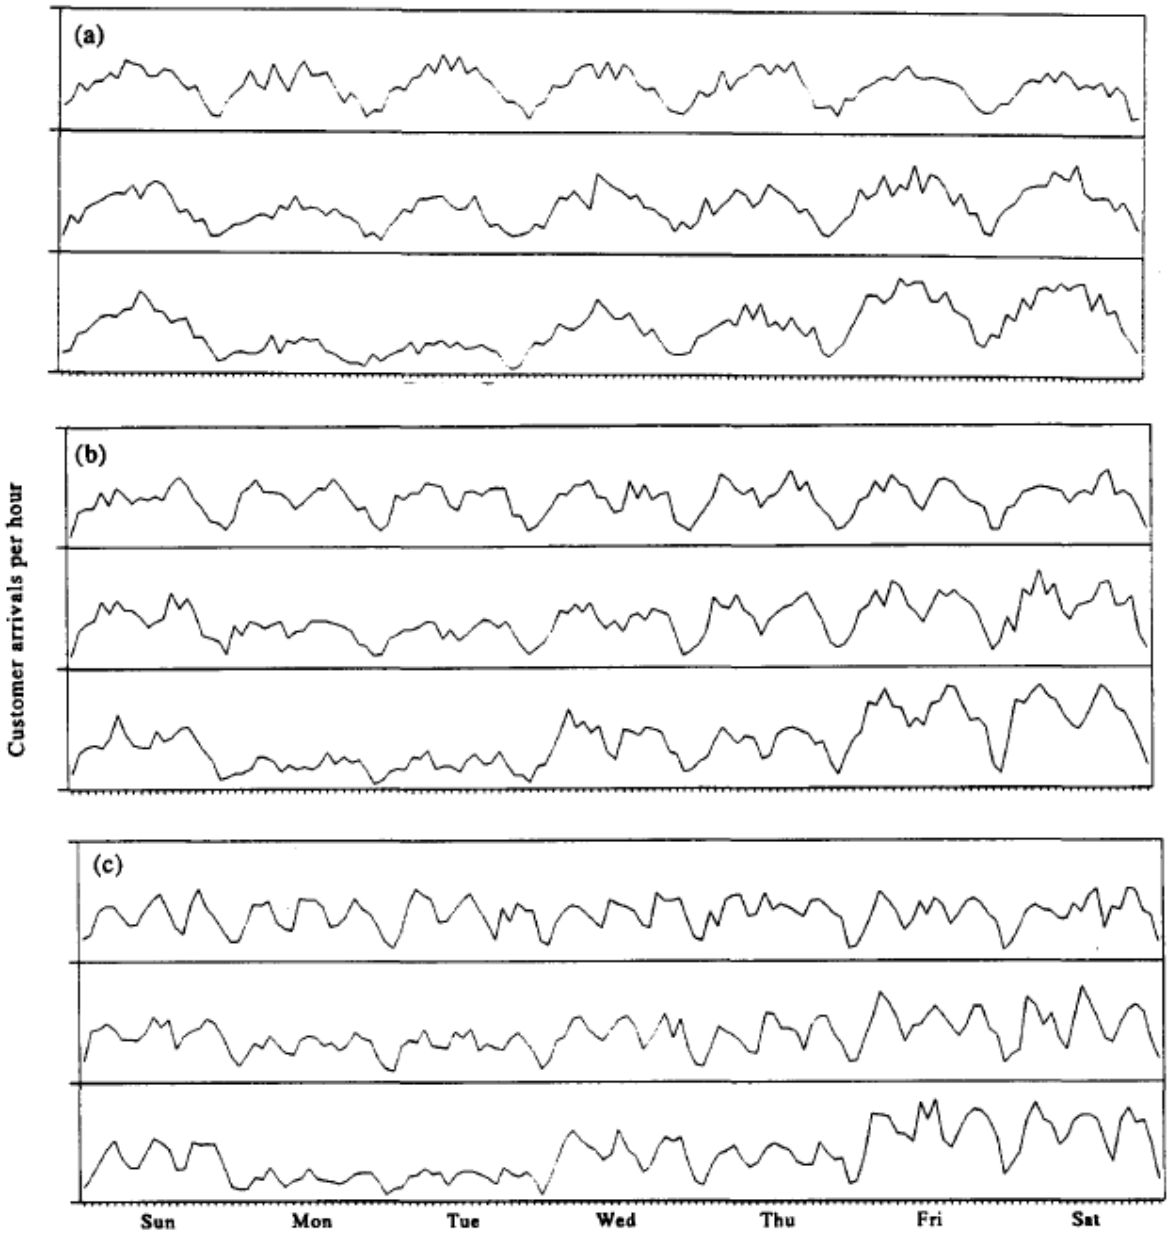

Figure 1. Representative weekly customer arrival patterns. (The scale of all patterns is identical, ranging from 0 to 125 customer arrivals per hour.) (a) Uniform within-day customer arrival patterns. From the top, these patterns are the low-, medium- and high-variation across-day patterns. (b) Bimodal within-day customer arrival patterns. From the top, these patterns are the low-, medium- and high-variation across-day patterns. (c) Trimodal within-day customer arrival patterns. From the top, these patterns are the low-, medium- and high-variation across-day patterns.

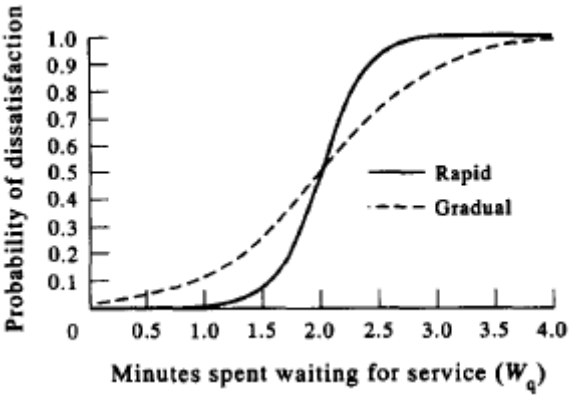

Figure 2. Customer dissatisfaction curves.

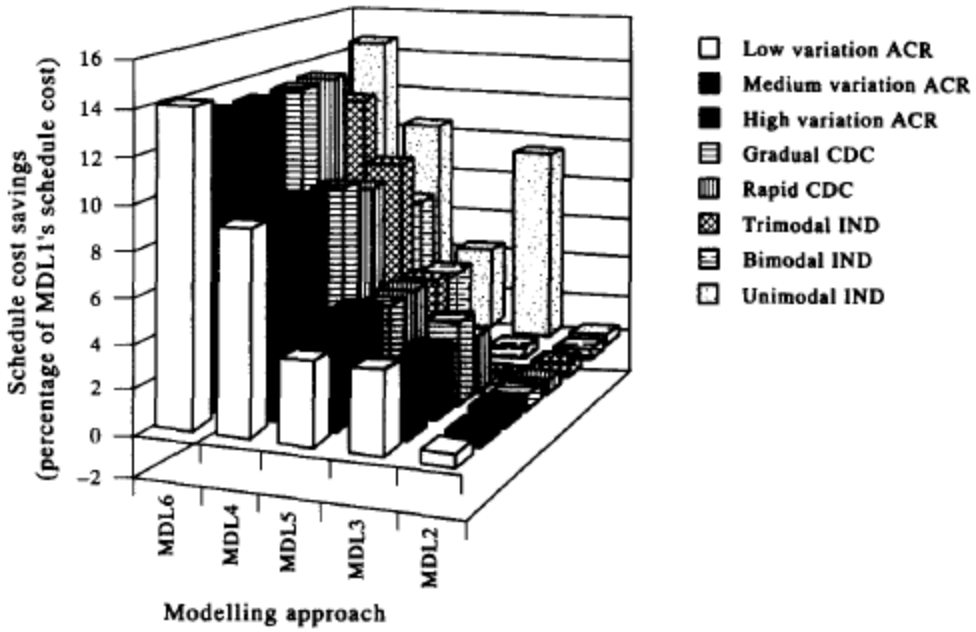

Figure 3. Relative model performance, by level of ACR, CDC and IND.

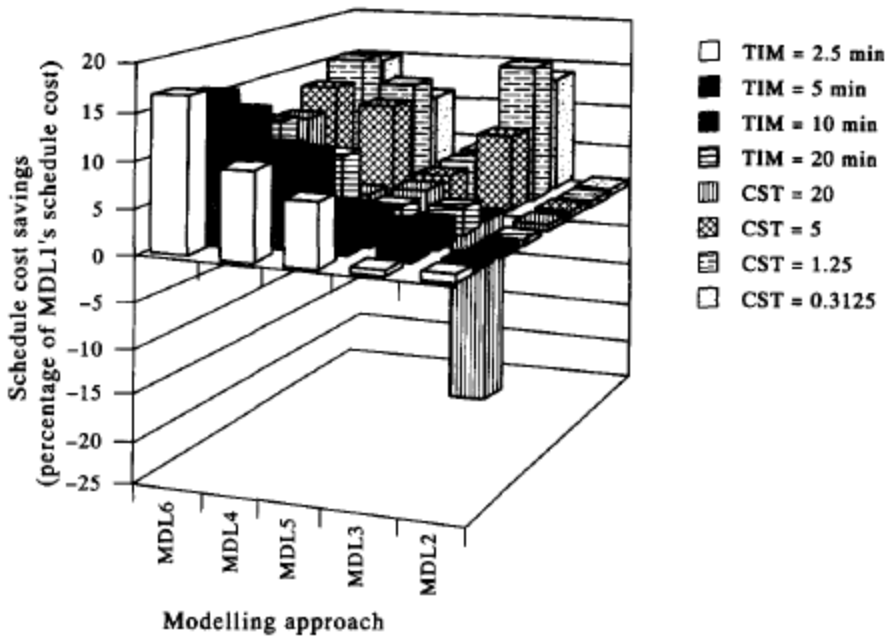

Figure 4. Relative model performance, by level of TIM and CST.

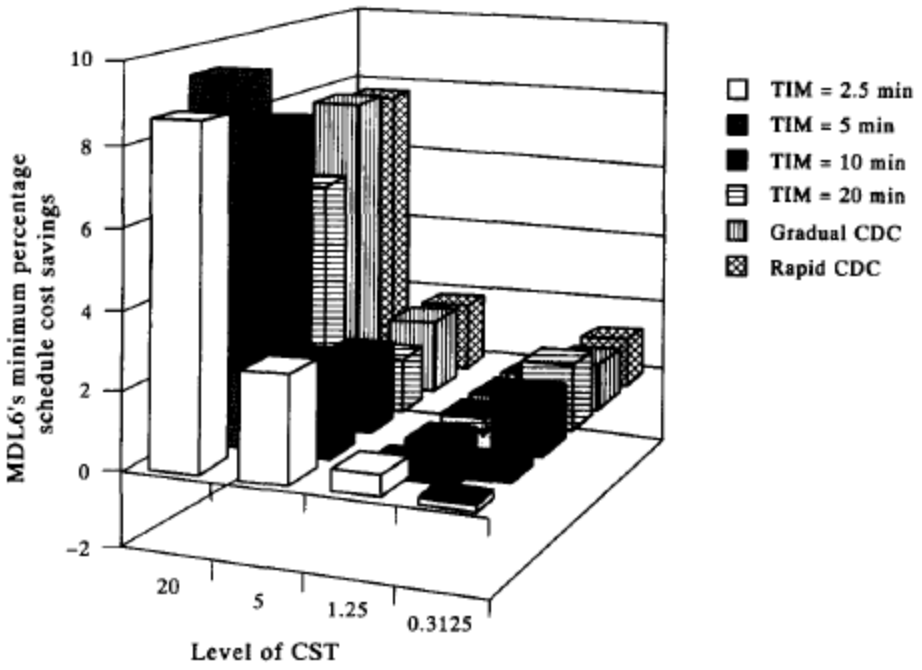

Figure 5. MDL6's minimum percentage schedule cost advantage by level of TIM and CDC, across the levels of CST.