

CONTINGENCY PLANNING FOR ABSENCE OF DELIVERY VEHICLE DRIVERS

Michael A. Haughton
Davoud Shirazi

School of Business & Economics, Wilfrid Laurier University
Waterloo, Ontario, Canada N2L 3C5, Telephone: (519) 884-0710
Email: mhaughton@wlu.ca

1. Introduction

Very few studies explicitly treat unexpected worker absenteeism as an important consideration. The present study extends this treatment to the context of vehicle routing and scheduling. A key contribution of this work is to show how features of optimal vehicle routing problem (VRP) yield an alternative operational response to the problem of uncertainty about the day-to-day number of scheduled drivers that fail to show for work. Thus, our central goal is performance comparison of this alternative response with the typical default option of exclusive reliance on a mix of standby drivers. Since the primary comparison criterion is labour cost, a related goal of the paper is to portray the response-specific labour costs caused by driver absenteeism. The next section details the context for the comparison and the literature review section provides further elaboration on how that context represents an extension of existing treatment of the worker absenteeism problem. The analysis of response-specific labour costs includes determination of how costs are affected by: (i) the rate of driver absenteeism, (ii) the variability of customer orders, (iii) the capacity of delivery vehicles, and (iv) the number of customers requiring deliveries.

Key practical motivations for this work are the significant costs of worker absenteeism and the persistent need for emergency operational tactics to avoid undue disruption of business operations when some of the scheduled workers are *unexpectedly* absent. Various reports underscored the significant costs. For example, the annual direct cost to absenteeism to Canadian employees has been estimated at \$12 billion by the Human Resources Professionals Association of Ontario (Buffet and Bachman, 2001/2002). Also, the Ontario Medical

Association (OMA) reports that the additional costs or indirect costs to the economy –e.g., health care services and lost productivity– for absences related solely to mental illnesses such as stress, depression, anxiety, and other mental disorders adds a further \$14 billion to absenteeism costs¹. This suggests a total cost (direct and indirect) of 2 to 3% of Canada's average annual GDP of around \$1.2 trillion since 2000. At the individual employee level, a report by Buffett and Bachman (2001/2002, p. 2) cites an estimate of \$3,500 per year per worker as the employer's directly incurred cost of absenteeism.

Within the occupation of commercial vehicle driving, the ongoing need for emergency operational responses seems to be epitomized by a study of Alabama's Metro Transit Authority (Phillips and Phillips, 1998). This study reported that while the intervention of several human resources policies significantly *reduced* driver absenteeism, the post-intervention absenteeism rate was still substantial: 4.8% (down from the pre-intervention rate of 8.7%) and causing 13.9% of all bus delays (Phillips and Phillips, p. 6). This suggests that, since an absenteeism rate of zero is probably unattainable, the human resources programs to reduce absenteeism must be complemented by operational tactics to counter the disruptive effects of the inevitable failure of all scheduled drivers to show up for work. In Canada, where Statistics Canada reports the employee "inactivity rate" and the number of lost days per employee as metrics for unscheduled absenteeism, the annual averages over the period 2000-2004 for the inactivity rate and the number of lost days for transport equipment operators averages were, respectively, 4.48% and 11.26 days (or 4.5% of a 250-workday year).

Using the multiple vehicle routing context, this study introduces a tactic that exploits a property of optimal solutions to the classic vehicle routing problem (VRP). This tactic counters absence-induced staff shortfalls by assigning extra routes to drivers that are present for work (akin to assigning overtime work). The study measures the labour cost performance of this tactic and compares it with the normally default tactic of exclusive reliance on each day's reserve pool of drivers. That reserve pool comprises part-time drivers and full-time off-duty drivers. The VRP solution structure that gives rise to the proposed tactic creates an opportunity to reduce the size of the required pool of part-time staff. Aside from the potential cost savings

¹"Mental illness and workplace absenteeism", accessed January 8th, 2006 at the OMA website (www.oma.org).

(merely having a part-time pool has a cost independent of the utilization level of the part-timers) well documented concerns about driver shortages suggest that the proposed tactic is a necessity in some situations. The proposed tactic, along with the default/baseline tactic, is described in the next section, which also outlines the context for the study. Following that section and the literature review section, the paper presents in sequence, the relevant models, the experimental methodology, the findings, and the conclusions.

2. Alternative Tactics and Comparison Context

This study's setting is a square geographic zone consisting of N customers, each with delivery order sizes (demands) that vary each day according to some known probability distribution. A vehicle of capacity Q units serves each route comprising a separate subset of the N customers. As part of their medium-term staff planning activities to determine parameters such as the mix of full-time/part-time drivers, specified personnel at the depot require actual (preferably) or forecasted customer orders for several months into the future. This study considers two sources of staff planning uncertainty: (i) each day's actual number of absent drivers (e.g., a reserve cadre of drivers will be planned to provide a specified probability that all orders will still be delivered) and (ii) the gap between the day's forecasted and actual orders (again, surplus drivers will be needed to avoid excessive staff shortfalls due to under-forecasting). In cases where each day's operational response to realized absenteeism is exclusive reliance on the reserve pool (comprising part-timers, and off-duty full-timers), an objective is to determine the full-time/part-time mix that minimize the tactic's cost. The definitions and formulations for that tactic (labeled Tactic 1 from this point on) follow.

2a. Tactic 1: The Baseline/Default Option

Define:

l as the day's required number of drivers

FT as the number of full-time drivers on staff.

$s_l(l)$ as the day's planned number of surplus drivers to substitute for absent co-workers if the recovery from absenteeism is to call exclusively on a mix of part-time and off-duty full-time drivers.

$PT(l)$ as the day's planned number of part-time drivers

PT as the overall number of part-time staff to maintain

p as the absenteeism rate and a as a random variable for the day's number of absent drivers; a is binomially distributed with (l, p) .

C_s and C_{PT} as the daily per driver cost of respectively, having surplus drivers "on call" as a contingency for absenteeism and maintaining a pool of part-time drivers.

r as the maximum acceptable probability of not having enough drivers to deliver all customer orders for the day.

The formulation to determine the cost-minimizing set of values for $s_l(l) \forall l$ can be stated as:

$$\text{Minimize } C_s \sum_{l=0}^{FT} s(l) + C_{PT} [\max\{PT(l), l = 0, 1, \dots, FT\}] \quad (1)$$

$$\text{Subject to: } \sum_{l=0}^{FT} P[L = l] \times P[A > s(l)] \leq r \quad (2)$$

The above problem would have to be solved at the start of the medium-term planning horizon. In the above formulations, constraint (2) ensures attainment of the $(1 - r)$ service level: the probability of delivering all orders, irrespective of how many scheduled drivers are unexpectedly absent or the size of the forecast error. This study assumes a 99.99% service level ($r = 0.0001$). A key underlying assumption of the above formulations concerns the issue of standby or on call wages. The generally accepted principle of standby wages is to compensate workers for the inconvenience of being obligated to immediately drop whatever non-work/personal activity they may have planned for the day and get to work at a moment's notice. Still, as noted in Hirschman (1999), there are instances in which standby pay is not an automatic entitlement for on call workers; e.g., the on call worker is not asked to sacrifice his/her personal time by having to be report for work (but accessible through, say, a beeper). In such instances $C_s = 0$ so the first term in (1) and (3) would be zero. Another way for the term to be zero is when a vehicle routing and dispatch manager need not put potentially needed surplus drivers on standby but can wait until the very last minute (e.g., the pre-dispatch morning period when actual absences are known) to summon the exact number of drivers needed to cover the staff shortfall. This would mean that the aforementioned inconvenience lasts for only a part of the morning. This study's analysis will acknowledge the possible existence of such instances, with the caveat noted by Hirschman that they might be rare. That is because, *inter alia*, shortage of capable workers (a situation that applies to commercial vehicle drivers) might make employers wary of displeasing and thus losing those workers through the tight-fisted policy of not paying standby wages.

The expressions in (1) and (2) address costing at the staff planning level. Costing must also be done at the daily operational level. That is, invariably there will be days when the actual number of absentees is such that some standbys are unutilized (Phillips and Phillips, 1998). Thus, by denoting C_w as the daily wages paid for utilized standbys (C_s is still the amount paid to idle standbys), the calculations of overall per day costs of Tactic 1 is as in (3). The third term in the expression ensures that standby pay is not double-counted when a standby worker is actually utilized; i.e., all standby workers are guaranteed C_s but those utilized earn a full day's wage of C_w ; i.e., an additional $(C_w - C_s)$.

$$C_s \sum_{l=0}^{FT} P[L=l] \times s_1(l) + C_{PT} [\max\{PT(l), l = 0, 1, \dots, FT\}] + (C_w - C_s) \sum_{l=0}^{FT} \sum_{a=0}^l P[A=a, L=l] \times \min\{a, s_1(l)\} \quad (3)$$

2b. Tactic 2: A Proposed Alternative Option

The key difference between the baseline and proposed options is that the latter exhausts overload work opportunities presented by optimal (distance minimizing) VRP solutions before calling on the standby drivers. How the opportunities arise is illustrated in Table 1, which shows the optimum solution to a VRP for 200 delivery addresses with gamma distributed orders ($\alpha = 100$, $\beta = 1$) and served by vehicles with a capacity of 2000 units each. The service zone is a square of side 100 kilometres in southwestern Ontario, Canada, and the 200 addresses are randomly selected from a larger set of 1000 delivery addresses (spatially positioned according to a uniform probability distribution for both the latitude and longitude of actual address locations in the area). The centrally located depot is the start and end point of all routes. In this illustration, the only addition to the standard VRP formulation is the constraint that no driver's workload should exceed a travel distance of 200 kilometres. The optimal solution is shown in Table 1, with the optimal routes listed in ascending order of travel distance. Four routes comprising two pairings (route 1 paired with route 6 and route 2 paired with route 5) satisfy the distance constraint.

The significance of this observation is that if no more than two of the day's l scheduled drivers are unexpectedly absent then the shortfall can, without detriment to the workload limit of 200 kilometres per driver, be fully made up by the $(l - a)$ drivers in attendance. That is, Tactic 2 would involve reallocating the workload so that a of those l

- *a*) drivers will each be responsible for two routes. As the example in Table 1 shows, if John and Kevin are absent then routes can be reassigned so that the 9 drivers who showed up for work will handle the workload (with two of them each handling two routes). Drivers from the surplus/reserve pool will only be called on if more than two drivers are absent (i.e., $a - 2$ of them will be called on). Thus, the benefit of this tactic is that by reducing the need to call on reserves, it makes it possible to operate with a smaller pool of part-timers. The corollary of achieving this benefit is to maximize the number of route pairings that satisfy the stated constraint. That maximization problem can be formulated as an assignment problem as per (4). Constraint 5a and 5b ensure that each route can be paired with only one other route; constraint 5c ensures that the pairings satisfy the workload constraint.

$$\text{Maximize } \sum_{i=1}^l \sum_{j=1}^l X_{ij}^{(l)} \quad (4)$$

$$\sum_{i=1}^l X_{ij}^{(l)} \leq 1; \forall j \quad (5a) \quad \sum_{j=1}^l X_{ij}^{(l)} \leq 1; \forall i \quad (5b) \quad c_{ij}^{(l)} X_{ij}^{(l)} \leq 200 \forall i, j \quad (5c)$$

where:

$$X_{ij}^{(l)} = 1 \quad \text{if routes } i \text{ and } j \text{ are reallocated to one driver and } l \text{ drivers are scheduled to work}$$

$$c_{ij}^{(l)} = D_i^{(l)} + D_j^{(l)} \text{ and } D_i^{(l)} \text{ is the distance (length) of route } i.$$

TABLE 1: Solved sample VRP and contingency for absences

The Sample Solution			Adjusting for absences	
Driver	Rout	Route Length	Rout	Workload
Andy	1	71	1 & 6	71+113=184
Bobby	2	79	2 & 5	79+109=188
Charlie	3	103	3	103
Dave	4	104	4	104
Eric	5	109	7	167
Fred	6	113	8	173
George	7	167	9	188
Harry	8	173	10	198
Ian	9	188	11	200
John	10	198		
Kevin	11	200		
	Mean	136.8	Mean	167.2

It is important to note that, in order to verify that route pairing for the example in Table 1 was not a hastily chosen recourse (vis-à-vis the recourse of re-solving the VRP under the constraint that there are only

$l - a$ drivers), an attempt was made at re-solving the VRP with the new constraint. The lack of a feasible solution to the new VRP helped to confirm the necessity of the route pairing recourse. This confirmation, which was sought and obtained throughout the simulation experiments, is not surprising since minimizing the number of delivery routes is a part of the process of solving VRPs. That is, having found l as the minimum number of drivers (delivery routes) to satisfy a given day's outcome of demands across all customers, finding a feasible solution with fewer than l drivers is highly improbable (not impossible given that VRP solutions for realistic problems cannot be guaranteed as unequivocally optimal). For $l - a$ drivers, the reason that the route pairing tactic works yet re-solving is infeasible is that the demand outcomes invariably lead to violation of vehicle capacity constraints if one attempts to merge two paired routes for delivery by a single vehicle (i.e., the type of route merger that is tested in the search procedure involved in re-solving the VRP).

Confirmation of the route pairing tactic's necessity aside, two important assumptions are required for the tactic to be feasible. First is that the customers' time window constraints would have to be flexible enough for a driver to complete one route then proceed to his second route. Second is that assuring workload equity across drivers would have to be through alternatives to a requirement that all of any given day's routes must have roughly the same length. That is, taking the illustration in Table 1 as typical of the depot's delivery operations, an alternative might be to ensure that over time, the mean scheduled workload of each driver is approximately equal to the per route mean of 136.7 kilometres. This route rotation scheme would require that no driver be consistently scheduled to do, say, the long route of nearly 200 kilometres (or, for that matter, a short or an "easy" route like the 71-kilometre route). Instead, like his colleagues, he would get to ply other routes within the range 71 to 200 kilometres.

Each day on which at least one absence occurs, one or more drivers will experience a larger than scheduled workload; i.e., using the illustration in Table 1 as an example, a driver's actual workload will, on average, exceed his mean scheduled workload of 136.7 kilometres. In practice, that extra ("overtime") workload attracts wage premiums of at least 50% (i.e., time and a half). Thus, Tactic 2 is only beneficial if the overtime cost is exceeded by the savings from carrying a smaller pool of part-time drivers. Determining the size of the part-time pool for Tactic 2 uses analogous formulations to those for Tactic 1. The important difference is that those for Tactic 2 account for how

the required number of surplus drivers is affected by the maximum number of feasible route pairings from the formulation in (4) and (5). Specifically, defining $X^{(l)}$ as the maximum number based on a (perfect) demand forecast that l drivers should be scheduled, the required surpluses under Tactic 2 would be as specified in (6).

$$s_2(l) = \max\{0, s_1(l) - X^{(l)}\} \quad (6)$$

As with the planning procedure, the mean daily cost calculations for Tactic 2 are analogous to those for Tactic 1 but with the adjustment for the solution to (4). The computational formula in (7) shows the average daily operational cost. In these formulae, C_{OT} is the daily wage for performing overload/overtime driving duties. Derivation of the amount of overload work in (7) is fairly straightforward but page length restrictions prevent inclusion of the procedure here.

$$\begin{aligned} & C_s \sum_{l=0}^{FT} P[L=l] \times s_2(l) + C_{PT} \times [\max\{PT(l), l = 0, 1, \dots, FT\}] \\ & (C_w - C_s) \sum_{l=0}^{FT} \sum_{a=0}^l P[A=a, L=l] \times \min\{s_2(l), (a - \min\{a, X^{(l)}\})\} + \\ & C_{OT} \sum_{l=0}^{FT} \sum_{a=0}^l P[A=a, L=l] \times \left[\frac{(l-a) \times \min\{a, X^{(l)}\}}{(l-a) + \min\{s_2(l), (a - \min\{a, X^{(l)}\})\}} \right] \end{aligned} \quad (7)$$

In summary, the core research task involves comparing (7) with (3) to assess the cost effectiveness Tactic 2 vis-à-vis Tactic 1. The tactic-specific costs and the tactic-to-tactic comparison will cover the impact of four factors: the rate of absenteeism plus three vehicle routing parameters –the number of customers, their order size variability, and the capacity of delivery vehicles.

3. Literature Overview and Extensions

Largely founded on the works of Dantzig (1954) and Keith (1979), the scientific literature dealing with the problems of worker scheduling is quite extensive (see, for example, the review by Baker, 1976). Predominantly, the staffing and scheduling models in this literature assume perfect job attendance records by workers. Notable exceptions are in two recent studies: Easton and Goodale (2005) and Pinker and Larson (2003), as well as in Wild and Schneeweiss (1993). Like Wild and Schneeweiss (1993), Easton and Goodale (2005) also addressed the interplay between the short-term and longer-term levels in the hierarchy of staff planning/scheduling decisions but considered

forecasting a matter for follow-up work. The present study contributes to the literature by simultaneously treating that interplay.

This study's treatment of unplanned absenteeism in the context of vehicle routing also contributes to the body of work on vehicle routing problems (VRPs). First, the literature on uncertainty VRPs is dominated by customer/demand uncertainty (see, for example, reviews by Gendreau, Laporte and Séguin, 1996 and Baita, Ukovich, Pesenti, and Favaretto, 1998). Thus, in studying cases of uncertainty in each day's number of absentees, we expand the set of sources of uncertainty covered by that literature. The other contribution to the VRP literature is showing how optimal VRP solutions depicting unbalanced travel distances across routes can be exploited to add another operational tactic to correct for unplanned absenteeism (Tactic 2).

4. Research Methodology

Evaluating the expressions in (3) and (7) in order to determine the cost for each tactic required two inputs that had to be obtained through probabilistic simulation: the baseline staffing levels (l) and the route lengths $D_i^{(l)}$, both of which were required to solve (4) as a basis for obtaining $X^{(l)}$ and $s_2(l)$. Obtaining these inputs involved generating and solving a large number of vehicle routing problems using a full factorial experiment covering different combinations of three vehicle routing factors: variability in customer order sizes (two levels), vehicle capacity (three levels), and number of customers to be served (two levels). Table 2 summarizes the parameter settings for each of the three vehicle routing factors as well as those for the absenteeism factor (absenteeism rates of $\pi = 0.05$ and 0.10). Table 2 also shows the run length (number of simulated days) for each of the twelve (12) combinations of routing factors. Using the stated gamma distribution of demand, each simulated day featured a different randomly generated set of demand/order outcomes across all customers. The rationale for different run lengths was that the combinations exhibited differences in volatility of the values required to evaluate each tactic's cost; i.e., the staffing levels and the solutions to (7).

In particular, having initially decided on five replicates for each of the twelve routing factor combinations, the run length was then determined by visually inspecting plots of the replicate-specific per-

day mean of $X^{(l)}$ versus run length to determine the run length at which the mean appeared firmly settled into steady state. The procedure, based on recommendations in Law and Kelton (1991) was sufficient to yield an acceptably small margin of error in estimating $E[X^{(l)}]$: a maximum of 2.56% of the mean across all twelve combinations. Each replicate involved a separate randomly selected subset of N ($= 200$ or 500) customer addresses from a total of 1000 actual commercial and residential addresses within the region used as the experimental setting (a 100x100 kilometre square region in southwestern Ontario). Using Roadshow®, each of the $390 \times 5 = 1950$ vehicle routing problems (see Table 2B) was solved under the constraint that the maximum allowable workload of any single driver is 200 kilometres. This constraint also applied to the corresponding 1950 formulations in (4) and (5); i.e., the 200 kilometre workload limit was also imposed for drivers performing overload duty by taking on an extra route to fill in for absent colleagues.

TABLE 2A: Experimental Parameters and Factor Levels/Values

Variability of customer demand (β); Each customer's daily demand (<i>iid</i>) is gamma distributed with $\alpha\beta = 100$ so variance $(\alpha\beta^2) = 100\beta$; Two (2) levels of β (α) tested: 1(100) and 100(1).
Capacity of delivery vehicles in number of units of the product (Q); Three (3) levels tested: $Q = 1000, 1500, 2000$
Number of customers (N); Two (2) levels tested: $N = 200, 500$
Rate of absenteeism (π); Two (2) levels tested; $\pi = 5\%, 10\%$

TABLE 2B: Simulation Run Lengths for Each Factor Combination

Factor Set	Demand Variability (β)	Vehicle Capacity (Q)	Number of Customers (N)	Simulation Run Length
1	1	10	200	20 days
2	1	10	500	20 days
3	100	10	200	50 days
4	100	10	500	50 days
5	1	15	200	20 days
6	1	15	500	20 days
7	100	15	200	40 days
8	100	15	500	50 days
9	1	20	200	20 days
10	1	20	500	20 days
11	100	20	200	30 days

12	100	20	500	50 days
----	-----	----	-----	---------

Regarding the labour cost coefficients (C_w , C_s , C_{PT} , and C_{OT}), existing reports were heavily relied on to determine their values used in practice. The daily wage cost (C_w) was taken as \$154. This was based on the most recent published edition in the series of quadrennial studies by Transport Canada to determine the cost of commercial trucking operations: "Operating Costs of Trucks In Canada". The daily on-call or standby wage (C_s) was $0.25C_w = \$38.50$. The 25% is based on two actual examples of on-call remuneration from Hirschman (1999). Based on Phillips and Phillips (1998), the cost to maintain a part-time driver pool (C_{PT}) was also set at $0.25C_w$. The cost of taking on overtime/overload work (C_{OT}) was set at the seemingly conventional rate of time-and-half; i.e., $C_{OT} = 1.5C_w$. Having determined the cost for each tactic –the expressions in (3) and (7)– for each factor combination, they were then converted to more standardized form by expressing them as a percentage of what the payroll would be for that factor combination if no absenteeism occurred. That is, costs are viewed in terms of the percentage increase in payroll costs necessitated by the need for contingency action to mitigate the effects of absenteeism. The data were then subjected ANOVA procedures as a key step towards identifying (i) the factor combinations that might make Tactic 2 a more cost-effective counter-measure for driver absenteeism than Tactic 1 and (ii) the magnitude of the cost differences between the two tactics.

5. Research Findings

Table 3 show is an extract of the key results from the ANOVA for the cost difference between Tactic 1 and Tactic 2. For extract shows only the statistically significant (0.01 level) main and interactive effects. The reason that many of the effects are not statistically significant for the cost difference is that they impact each tactic's cost in very similar ways. More important than the ANOVA's depiction of what is statistically significant and what is not, are the particulars of the effects. These are discussed with reference to Figure 1.

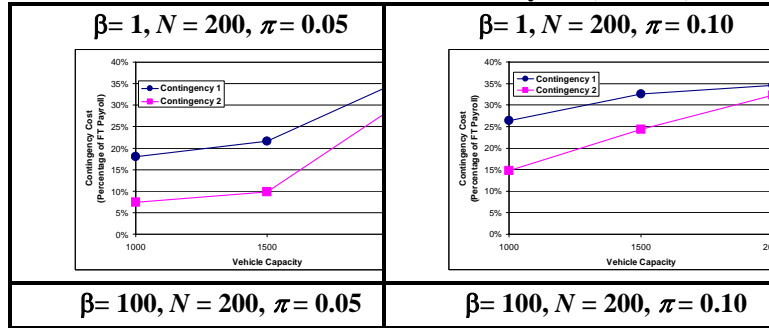
TABLE 3: Statistically Significant Impacts on the Cost Difference Between Tactic 1 and Tactic 2


Main or Interactive Factor Effects	P-value
Demand Variability	0.000
Vehicle Capacity	0.000
(Vehicle Capacity)*(Number of customers)	0.000

(Vehicle Capacity)*(Absentee rate)	0.002
(Number of customers)*(Absentee rate)	0.000
(Vehicle Capacity)*(Number of customers)*(Absentee rate)	0.000

Figure 1 plots each tactic's cost for the four combinations of absentee rate and vehicle capacity for $N = 200$. A similar plot mirroring Figure 1 was done for $N = 500$ but in light of the lack of statistical significance of N and space limitations, it is not shown here. The impact of absentee rate is intuitive: its increase from 0.05 to 0.10 raised contingency costs from an average of 13.94% to an average of 21.98% of payroll across all other factor combinations. Aside from these two impacts, the plots depict three conspicuous cost effects. First, Tactic 1 never outperforms Tactic 2. The fact that the latter tactic never had a higher cost suggests that if the previously discussed two supporting conditions for its feasibility are present (flexible delivery time windows and route rotation across drivers to assure workload equity) then it is a more financially sound contingency to deal with driver absenteeism. In fact, across all factor combinations covered in the simulation, the contingency cost for Tactic 1 raised payroll costs by 21.68% while the contingency cost for Tactic 2 raised payroll cost by 14.23%, roughly 7.50 percentage points less. The second noteworthy pattern is that contingency costs generally increase with increases in vehicle capacity but decreases with increases in the number of customers.

FIGURE 1: Each Tactic's Cost as a % of Payroll ($N = 200$)





The reason is that both factors have a very direct influence on staff level (larger vehicles mean fewer drivers since more customer deliveries can be handled on a single route and larger customer numbers increase the required number of drivers). Now, as a proportion of total payroll, the cost impact of having a contingency to assure a 99.99% service level will be larger for a smaller number of drivers. As such larger vehicle capacities, which translate to fewer drivers, will result in higher relative contingency cost, and larger customer numbers, which translate to more drivers, will yield lower relative contingency costs (visible in the left-to-right comparison of the plots in each figure). Using Tactic 1 for illustration, factor combination 3, which typically required a baseline staff of 24 drivers also required a contingency reserve crew of 6 drivers (25% addition of driving personnel) while in the case of combination 4, the corresponding numbers were 57 and 8 (a smaller relative addition of 14.05%). The third pattern of note is that demand variability actually lowers the cost of each tactic (top-to-bottom comparison of the plots in each figure). The explanation is that increased customer demand variability, like increased customer numbers, also raises the required baseline staff (though less dramatically than the impact cause by increased customer numbers), and thus lowers the relative cost.

The cost reducing impact of increased demand variability, though similar to the impact of increased customer numbers, is somewhat more intricate, particularly because it is more significant for Tactic 1 than for Tactic 2. This stems from the greater reliance of Tactic 1 on reserve crews (Tactic 1 relies on those present taking on extra work to handle the shortfall created by their absent colleagues). Now, a larger baseline staff means an increase in the available number of off-duty full-time workers that can be called on to fill in for their absent colleagues; this reduces the needed size of the crew of part-time drivers; i.e., reduces the term by which C_{PT} is multiplied in the expressions in (3). The final pattern of consequence is that Tactic 1 comes close to being on par with Tactic 2 only for the following

combination of factors: large vehicle capacity, few customers, large demand variability, and high absenteeism rate. The joint effect of large capacity and few customers creates cost parity (see plots to the left in each figure) by leading to fewer drivers (fewer pre-planned routes). With fewer routes, the travel distances tend to be more balanced so it was much harder for Tactic 2 to exploit its source of superiority over Tactic 1; i.e., to find a large number of feasible route pairings on which to base overload assignment. To illustrate the difficulty, for factor combination 4 which had an average of 52 routes per day, typically almost 25 feasible pairings of them were possible (i.e., 96% of them) but a smaller percentage of 81% for factor combination 3 which had an average of 21 routes per day. As regards the reason that high demand variability is also an element in the factor combination that creates parity, the explanation lies in the previously noted impact of demand variability on the required staff of reserve drivers; i.e., since Tactic 2 benefits less than Tactic 1 from this impact, its margin of superiority over Tactic 1 will be narrowed. The margin is further narrowed when all this is combined with a high absentee rate; i.e., given the cited factor combination's relatively large number of possible absentees that must be planned for vis-à-vis the limited overload assignment options, the reserve crew for Tactic 2 will tend to be just marginally smaller than that of Tactic 1.

6. Conclusions

The costs of contingency measures to ensure that unexpected driver absences do not disrupt vehicle routing operations can be significant: averaging between 13.94% of payroll (when the absentee rate of drivers is 5%) and 21.98% (when the rate is 10%). This study shows that a property of vehicle routing solutions provides an opportunity to materially reduce the significant cost of existing contingency measures. This property permits reallocation of workload (routes) among drivers who are present for work as scheduled. In particular, the study found that while adoption of the traditional contingency of relying on a crew of reserve workers can raise payroll costs by approximately 22%, the contingency of workload reallocation raises payroll cost by a smaller margin of 14%. The coverage of the vehicle routing area represents the central component of this study's extension of the extant literature on workforce planning and scheduling under absenteeism.

Though adding an important dimension to existing work on worker absenteeism (vehicle routing context), this modeling work can be viewed as a basis for both empirical follow-up work and future

modeling work to further understand the potential of the proposed contingency of route reallocation. First is the issue of time window flexibility. The paper acknowledges that tight time window constraints can limit the amount of cost savings attainable through route reallocation. It would be interesting to determine actual customers' willingness of customers to provide flexible time windows. Insights on the costs of route reallocation can also be gained by replicating this study under different time window constraints. The study can also be extended to consider forecasting accuracy, an issue which appears to have been addressed in only one previous study (Pinker and Larson, 2003). Finally, studies of contingencies used by managers responsible for actual vehicle routing and dispatch operations would provide enriching complements to this study. The issue of driver absenteeism is important enough for the vehicle routing literature to continue building a rich body of work on the topic.

While the latter study, as does the present one, dealt with the task of developing longer-term staff plans to set the parameters for shorter-term staff scheduling activities, it does not address the issue of possible imperfections in demand forecasts (which lead to imperfect staff projections). Pinker and Larson treat forecasting but their work focused on the short-term (day-to-day) scheduling decisions.

References

- Baita, F.; Ukovich, W.; Pesenti, R.; Favaretto, D. 1998. Dynamic routing-and-inventory problems: A review. *Transportation Research-A*, 32 (8), 585-598.
- Baker, K. R. 1976. Workforce allocation in cyclical scheduling problems: A survey. *Operational Research Quarterly* 27 (1), 155-167.
- Buffett, E., and Bachman, L. Dec. 2001-Jan, 2002. Prevention is the best medicine. *The HR Professional*, 1-4.
- Dantzig, G. B. 1954. A comment on Edie's "Traffic delays at toll booths". *Operations Research* 2 (3), 339-341.
- Easton, F. F. and Goodale, J. C. 2005. Schedule recovery: Unplanned absences in service operations. *Decision Sciences* 36 (3), 459-488.
- Gendreau, M., Laporte, G.; Séguin, R. 1996. Stochastic vehicle routing. *European Journal of Operational Research*, 88 (4), 3-12.
- Hirschmann. 1999. Paying for waiting. *HR Magazine*, 44 (8), 98-105.
- Keith, E. G. 1979. Operator scheduling. *AIIE Transactions* 11 (1), 37-41.

- Law, A., Kelton, W. D. 1991. *Simulation Modeling & Analysis*. (2nd ed.), McGraw-Hill.
- Mitchell, R. W. 2001. Employers taking steps to reduce costs of absenteeism. *National Underwriter* 105 (15), 57.
- Phillips, J. C., and Phillips, P. R. 1998. *Measuring ROI for an absenteeism reduction program*. A case study report to the Alabama Metro Transit Authority. Published by the Jack Phillips Center for Research.
- Pinker, E. J., and Larson, R. C. 2005. Optimizing the use of contingent labor when demand is uncertain. *European Journal of Operational Research* 1446 (1), 39-55.
- Wild, B., and Schneeweiss, C. 1993. Manpower planning – A hierarchical approach. *International Journal of Production Economics* 30/31, 95-106.