



2013

Workforce management in periodic delivery operations

Karen Smilowitz
Northwestern University

Maciek Nowak
mnowak4@luc.edu

Tingting Jiang
Northwestern University

Recommended Citation

Smilowitz, Karen; Nowak, Maciek; and Jiang, Tingting, "Workforce management in periodic delivery operations" (2013). *Information Systems and Operations Management: Faculty Publications & Other Works*. Paper 8.
http://ecommons.luc.edu/isom_facpubs/8

This Article is brought to you for free and open access by the Faculty Publications at Loyola eCommons. It has been accepted for inclusion in Information Systems and Operations Management: Faculty Publications & Other Works by an authorized administrator of Loyola eCommons. For more information, please contact ecommons@luc.edu.



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 License](https://creativecommons.org/licenses/by-nc-nd/3.0/).

Workforce management in periodic delivery operations

Karen Smilowitz

Industrial Engineering and Management Sciences, McCormick School of Engineering, Northwestern University, Evanston, IL, 60208, ksmilowitz@northwestern.edu,

Maciek Nowak

Information Systems and Operations Management, Graduate School of Business, Loyola University, Chicago, IL, 60611, mnowak4@luc.edu,

Tingting Jiang

Industrial Engineering and Management Sciences, McCormick School of Engineering, Northwestern University, Evanston, IL, 60208, tingting-jiang@northwestern.edu

This work was partially supported by grant DMI-0348622 from the National Science Foundation.

Service quality and driver efficiency in the delivery industry may be enhanced by increasing the regularity with which a driver visits the same set of customers. However, effectively managing a workforce of drivers may increase travel distance, a traditional metric of the vehicle routing problem (VRP). This paper evaluates the effect that workforce management has on routing costs, providing insight for managerial decision making. The analysis is presented in the context of the period vehicle routing problem (PVRP), an extension of the VRP with vehicle routes constructed to service customers according to preset visit frequencies over an established period of time. We develop models to apply workforce management principles. We show that multi-objective PVRP models can achieve a balance between workforce management and travel distance goals, through a computational study with standard PVRP test cases and real-world delivery data. With the proper parameters in place, workforce management principles may be successfully applied without sacrificing other operational objectives.

Key words: periodic vehicle routing problems; workforce management; multi-objective optimization

1. Introduction

With intense external pressure from competitors and internal pressure from shareholders mandating increases in both customer service and profit margins, carriers in the delivery industry must utilize every advantage possible. This has raised attention on the efficient use of drivers and the importance of service quality provided to customers. Both service quality and driver efficiency may be enhanced by increasing the regularity with which a driver visits the same set of customers. UPS, the biggest operator in the domestic package express industry, relies on their drivers and the system for managing drivers to maintain a high level of customer satisfaction:

“...Many UPS drivers work the same route for 20 or 25 years. In contrast, a major competitor

reserves the right to reconfigure some drivers' routes with five days' notice, meaning their customers, service area and earnings power can change quickly. UPS drivers form a real bond with customers...A formal program that gathers sales leads from drivers generates volume of more than 60 million packages a year, largely because drivers take tremendous ownership of their customers and routes." [UPS Corp. (2006)]

This example shows the potential for revenue generation from better workforce management in delivery services. Competitive advantage through workforce management has been recognized and extensively studied [Pfeffer (1998) and Bartlett and Ghoshal (2002)]. However, an increase in workforce effectiveness may come with additional costs elsewhere. While consistency is generally the goal with delivery workforce management, such that drivers are assigned to the same customers and regions as often as possible, routes are most commonly constructed to minimize the travel distance for drivers with little consideration for consistency. These two objectives may conflict with each other. In a recent review of commercial vehicle routing software, one vendor notes, "If we tell a customer that it costs them \$100 a day to keep the same drivers servicing the same customers, most fleets will decide it isn't worth the cost." [Partyka and Hall (2010)]

In this paper, we study the relationship between the objectives of workforce management and travel costs, and examine how incorporating multiple objectives impacts managerial decision making. Designing vehicle routes to minimize distance may result in multiple drivers visiting a customer over the planning horizon. Designing vehicle routes to enhance customer-driver relationships may result in the same driver visiting a customer each time, yet travel cost could increase. Estimating the magnitude of this cost increase is difficult. Often routing solutions that minimize distance have the characteristic that customers in close proximity are visited by the same vehicle route, which is the aim of workforce management. However, given customers with unique visit frequency requests, and with other operational constraints such as fleet size and vehicle capacity, the dynamics can be quite intricate. In order to balance workforce management and travel costs one must first quantify workforce management in route construction. We consider two metrics:

1. *Customer familiarity*: As the driver becomes familiar with a set of customers, she may more efficiently serve her customer base. By visiting the same customers repeatedly, service quality is elevated and the customers become more likely to use the carrier in the future. The customer familiarity metric is modeled as an objective to *maximize the number of times that a unique driver visits a customer*.

2. *Region familiarity*: The set of customers requesting service each day varies; therefore, it may be beneficial to study workforce metrics at a more aggregated level. By visiting a customer in

a region, the driver gains familiarity that is of benefit when any customer in the region is later visited. The more familiar a driver is with a certain region, the faster she can deliver service to all customers in the region. Furthermore, with a good reputation in a neighborhood, the driver is likely to generate more business. Region familiarity is modeled similarly to customer familiarity, as an objective to *maximize the number of times that a driver repeatedly visits a region*.

We study the effect of adding these metrics under the general setting of the Period Vehicle Routing Problem (PVRP). The PVRP, introduced by Beltrami and Bodin (1974) and Russell and Igo (1979), is an extension of the Vehicle Routing Problem (VRP), with vehicle routes constructed to serve customers according to preset visit frequencies over an established period of time. The traditional objective of the PVRP is to create a set of tours for each vehicle on each day in the period to minimize the travel cost, while satisfying operational constraints such as vehicle capacity and customer visit frequency. The PVRP can be found in industries such as package delivery, grocery delivery, waste collection, etc. The multiple-day setting of the PVRP provides a time interval to study the workforce metrics highlighted above. As reviewed in Section 2, prior approaches have enforced workforce management through model constraints.

In this paper, we develop and analyze alternate models of the PVRP to incorporate customer familiarity and region familiarity. We show that multi-objective PVRP models can achieve a balance between workforce management and travel distance goals. Contrary to the belief of some in industry, we find that workforce management principles can be applied with a minimal impact on travel costs, through a computational study with standard PVRP test cases, modified for our problem setting. These results exploit the added flexibility of modeling workforce management in the PVRP as an objective and not a constraint. The computational analysis shows that the solution with the best balance between workforce management and travel costs may result in multiple drivers serving a customer. We demonstrate the usefulness of the new PVRP variations on a data set from a major package delivery carrier. Fundamentally, this paper presents transportation managers with best practice methods for applying workforce management without sacrificing other operational objectives.

The paper is organized as follows. Section 2 presents an overview of literature in workforce management and periodic routing. Section 3 introduces and analyzes PVRP models to capture workforce and travel distance objectives. Section 4 describes the Tabu Search approach to solve the PVRP models and presents the computational study. Section 5 concludes the paper with a discussion of extensions.

2. Literature Review

The interaction between operations and workforce management has been studied from several perspectives, primarily in staffing and control, with applications in call center management [Jennings et al. (1996) and Whitt (2005)]. Hopp et al. (2004) and Iravani et al. (2007) also study effective workforce cross-training structures. In vehicle routing, there is research in creating routes with balanced daily schedules [Levy and Bodin (1988)] and balanced route lengths [Jozefowicz et al. (2007)]. Levy and Bodin (1988) schedule the postal carriers of the United States Postal Service and design Euler cycles for delivering mail. Two procedures for developing a complete set of daily work schedules and Euler cycles are developed. Jozefowicz et al. (2007) present a bi-objective vehicle routing problem with route balancing in terms of the route length for drivers that is solved using a hybrid multi-objective genetic algorithm.

The PVRP is an extension of the classic vehicle routing problem (VRP), where the vehicle routes are constructed over a period, typically in units of days in a week. Most commonly, objective functions for the PVRP focus on minimizing travel costs and/or fleet size; see Francis et al. (2008) for a recent review of the PVRP literature. Early studies of workforce management in the context of the PVRP rely on fixed routes to enforce driver consistency [Christofides (1971), Beasley (1984), and Wong and Beasley (1984)]. As noted in these references, fixed routes can lead to capacity-infeasible solutions or underutilized routes when customers differ over the period. Recent work focuses on driver consistency in the express delivery industry where a company may wish to send the same crew to a customer repeatedly to take advantage of the familiarity the crew establishes with the customer or with the geographic region. Zhong et al. (2007) seek to maintain driver familiarity with their service territories. The authors model a learning/forgetting behavior for drivers and show that dispatching drivers consistently to the same geographic areas results in driver familiarity and improved driver performance.

Francis et al. (2007) considered a range of workforce metrics for the PVRP, including driver coverage, which measures the portion of the total service region visited by a driver over the period, and crew size, which measures the number of different drivers visiting a customer over the period. These metrics are calculated *a posteriori* to evaluate periodic routing solutions obtained with the objectives of minimizing travel time and maximizing visit frequency. Their model does not attempt to optimize these workforce measures. In the work presented here, we incorporate workforce measures in the objective of the PVRP.

Similar to this paper, Groer et al. (2009) consider a multi-period routing problem where companies want their drivers to develop relationships with customers on a route. They introduce the

Consistent VRP (ConVRP) as a variant of the VRP, requiring that a customer be visited by only one driver over the period at approximately the same time on each day, in addition to the traditional constraints on vehicle capacity and route length. The work presented here differs from the ConVRP in that driver management is modeled as an objective rather than as a constraint. We evaluate the extent to which the added flexibility from modeling customer familiarity as an objective rather than a constraint can improve operations.

Importantly, most multi-period vehicle routing problems are solved with heuristics given the complexity of the problem. Francis et al. (2006) introduce an exact method for solving the PVRP; however, their approach is limited to instances in which each node must be visited by the same driver throughout the time period. With such an assumption, one cannot explore the workforce metrics for more flexible systems. In Francis et al. (2007), the authors show that using Tabu Search, one can relax the restrictive assumptions of Francis et al. (2006) and obtain better solutions, although no longer provably optimal. Further, in a comparative analysis of PVRP solution methods performed by Francis et al. (2008), the Tabu Search heuristic by Cordeau et al. (1997) is found to provide high-quality solutions to a known set of PVRP test cases. Notably, these test cases have not been solved optimally.

3. Modeling

In this section, the general PVRP setting is introduced, which is employed in models based on travel distance and two workforce management metrics. The relationship between these different models is then discussed.

3.1. PVRP problem setting

In the standard PVRP, a homogeneous fleet of vehicles operate from a single depot to serve a set N of customers over a period of time. Let T be the set of days in the time period, indexed by $t \in T$. Associated with each vehicle is a route $k \in K$ with a fixed capacity C . A vehicle completes at most one route per day; an idle vehicle is modeled as a route of no cost. For each day within the time period, a set of routes is developed to meet a given objective (most commonly minimizing total travel cost) while satisfying capacity constraints and demand requirements. An example with three customers requesting service from two vehicles over a two-day period is shown in Figure 1. The PVRP is defined here on a complete network denoted by the underlying graph $G = (N^0, A)$, with node set N^0 representing the customers in set N plus the depot (denoted $i = 0$) and arc set A connecting nodes in N^0 with non-negative travel costs c_{ij} for all $(i, j) \in A$. The arc costs are shown on the arcs in Figure 1.

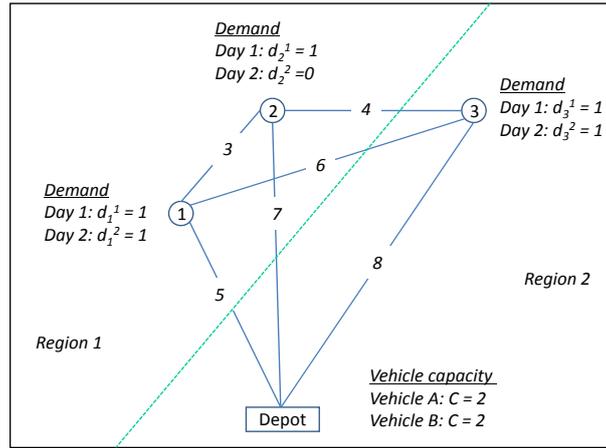


Figure 1 Period vehicle routing problem example

In the example, each customer has a specific demand request per day. More general demand and service assumptions have appeared in the literature; often a customer requests a number of visits over the time period with a constant demand per visit, determined by either the average or maximum accumulation between visits. The PVRP then determines the visitation schedules for customers from a set of feasible options; i.e., Monday-Wednesday or Tuesday-Thursday. In this paper, the visitation schedule is assumed to be set *a priori*; for a given day $t \in T$, the set of customers to visit and their respective demands are given. Let d_i^t denote the demand of customer $i \in N$ on day $t \in T$. When travel distance is the only performance metric, the PVRP decomposes by day given these assumptions. However, workforce metrics link the problem across days.

In the following subsections, mixed integer models of the PVRP are developed. Given a chosen objective, described in Section 3.2, the PVRP models design vehicle routes (set of customers visited and sequence of visits) and assign drivers to routes such that operational constraints are observed. Split delivery is not allowed within one day of service: a customer is visited by at most one route per day. However, over the period T , a customer may be visited by more than one driver.

3.2. Performance metrics for the PVRP

The three performance metrics studied in this paper are described below.

- *Travel Distance: PVRP(TD)*. The objective function seeks to minimize total travel distance across all routes over the time period.
- *Customer Familiarity: PVRP(CF)*. The objective function seeks to maximize the number of visits by a driver to a customer with a cost term that reduces per customer-visit for a driver as the frequency of visits to that customer increases, recognizing that routing solutions can be more efficient and effective if drivers repeatedly visit the same set of customers.

- *Region Familiarity: PVRP(RF)*. The objective function reduces the cost per visit to a geographic region for a driver as the visit frequency to that region increases for that driver. Since the customer set varies by day, it may be beneficial to study workforce metrics at an aggregated level.

As noted in Section 3.1, the key decisions of the PVRP are (1) vehicle route design (set of customers visited and sequence of visits) and (2) driver assignment to vehicle routes. Figure 2 illustrates how these decisions are addressed in each variation. In the PVRP(TD), vehicle route design is separated from driver assignment. In Phase I, routes are associated with vehicles, but not drivers. The set of customers visited and sequence of visits on a vehicle route are determined. In Phase II, driver assignments for the routes are performed. In the workforce management models, PVRP(CF), and PVRP(RF), there is no distinction between drivers and vehicle routes since the decisions are made simultaneously.

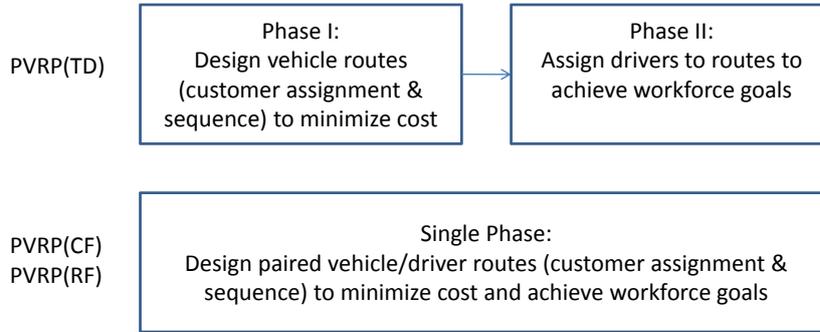


Figure 2 Decision phases of the PVRP by model

3.3. Travel distance: PVRP(TD)

The PVRP(TD) designs vehicle routes in Phase I to minimize travel distances and then assigns drivers to those routes in Phase II to achieve workforce goals.

3.3.1. Phase I: Vehicle routing For each day in the time period, Phase I determines the assignment of customers to routes and the arc routing variables:

$$\begin{aligned}
 y_{ik}^t &= \begin{cases} 1 & \text{if customer } i \in N \text{ is visited by vehicle route } k \in K \text{ on day } t \in T \\ 0 & \text{otherwise} \end{cases} \\
 x_{ijk}^t &= \begin{cases} 1 & \text{if vehicle route } k \in K \text{ traverses arc } (i, j) \in A \text{ on day } t \in T \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

The travel cost objective is:

$$f_{\text{TD}} = \sum_{(i,j) \in A} \sum_{k \in K} \sum_{t \in T} c_{ij} x_{ijk}^t. \quad (1)$$

Given that costs and constraints can be decomposed by day, Phase I of the PVRP(TD) is solved as a series of $|T|$ independent vehicle routing instances, using the following formulation for day $t \in T$:

$$\min \sum_{(i,j) \in A} \sum_{k \in K} c_{ij} x_{ijk}^t \quad (2a)$$

subject to

$$\sum_{k \in K} y_{ik}^t \geq 1 \quad \forall i \in N : d_i^t > 0 \quad (2b)$$

$$\sum_{i \in N} d_i^t y_{ik}^t \leq C \quad \forall k \in K \quad (2c)$$

$$\sum_{j \in N^0} x_{ijk}^t = y_{ik}^t \quad \forall i \in N; k \in K \quad (2d)$$

$$\sum_{j \in N^0} x_{ijk}^t = \sum_{j \in N^0} x_{jik}^t \quad \forall i \in N^0; k \in K \quad (2e)$$

$$\sum_{i,j \in Q} x_{ijk}^t \leq |Q| - 1 \quad \forall Q \subseteq N; k \in K \quad (2f)$$

$$y_{ik}^t \in \{0, 1\} \quad \forall i \in N; k \in K \quad (2g)$$

$$x_{ijk}^t \in \{0, 1\} \quad \forall (i, j) \in A; k \in K \quad (2h)$$

The objective function (2a) minimizes routing costs for day $t \in T$. Constraints (2b) ensure that a customer $i \in N$ is visited by a route if that customer has demand on a given day. If there is no demand on that day, the customer does not have to be visited. Constraints (2c) are the physical capacity constraints for each route. Constraints (2d) link the x and y variables for the customers in set N . Constraints (2e) ensure flow conservation at each node $i \in N^0$. Constraints (2f) are the subtour elimination constraints which also guarantee that each route contains a stop at the depot. Finally, constraints (2g) and (2h) define the binary variables for assignment and routing.

3.3.2. Phase II: Driver assignment Once the routes are constructed, a post-processing phase is employed to improve workforce measures when assigning drivers to routes. The mathematical programming approach from Francis et al. (2007) is used to assign the drivers to the routes of the Phase I solution to minimize a performance metric that counts the regions visited by a route. Their approach is extended here to post-process according to customers as well.

The service area is partitioned into a set R of regions, indexed by r , such that each region contains at least one customer. In Figure 1, the service region is divided into two regions: $r = 1, 2$. Let N_r denote the set of customers contained in region $r \in R$ and let L be the set of drivers. Assume that $|L| = |K|$. If a PVRP(TD) solution does not use all $|K|$ vehicles, the idle vehicles are modeled as routes that remain at the depot with no cost incurred.

The assignment problem from Francis et al. (2007) minimizes the number of regions that each driver covers. A driver covers a region if he visits at least one customer in that region during the time period.

Given a Phase I solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$, one can determine the following parameter:

$$\hat{a}_{rk}^t = \begin{cases} 1 & \text{if vehicle route } k \in K \text{ visits region } r \in R \text{ on day } t \in T; \text{ i.e. if } \sum_{i \in N_r} \hat{y}_{ik}^t \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

The following decision variables are used in the assignment problem:

$$\pi_{lk}^t = \begin{cases} 1 & \text{if driver } l \in L \text{ is assigned to vehicle route } k \in K \text{ on day } t \in T \\ 0 & \text{otherwise} \end{cases}$$

$$\nu_{rl} = \begin{cases} 1 & \text{if region } r \in R \text{ is visited by driver } l \in L \text{ at least once during the period} \\ 0 & \text{otherwise} \end{cases}$$

The region-based assignment problem is formulated as:

$$\min \sum_{r \in R} \sum_{l \in L} \nu_{rl} \quad (3a)$$

subject to

$$\nu_{rl} \geq \sum_{k \in K} \hat{a}_{rk}^t \pi_{lk}^t \quad \forall r \in R; l \in L; t \in T \quad (3b)$$

$$\sum_{k \in K} \pi_{lk}^t \geq 1 \quad \forall l \in L; t \in T \quad (3c)$$

$$\sum_{l \in L} \pi_{lk}^t \leq 1 \quad \forall k \in K; t \in T \quad (3d)$$

$$\pi_{lk}^t \in \{0, 1\} \quad \forall l \in L; k \in K; t \in T \quad (3e)$$

$$\nu_{rl} \geq 0 \quad \forall r \in R; l \in L \quad (3f)$$

The objective (3a) minimizes the number of regions covered by the drivers. Constraints (3b) set ν_{rl} to 1 if driver $l \in L$ is assigned to vehicle route $k \in K$ that visits region $r \in R$ on at least one day. Constraints (3c) ensure that each driver is assigned to a vehicle route on each day and constraints (3d) ensure that only one driver is assigned to a vehicle route on a given day. Given that $|L| = |K|$, both constraints (3c) and (3d) hold at equality to guarantee that each driver is assigned to at least one route each day and each route is served by at most one driver. Constraints (3e) and (3f) define the decision variables (note that ν_{rl} is binary, given binary values for π_{lk}^t).

Alternatively, a customer-based assignment problem is introduced with the following notation.

$$\omega_{il} = \begin{cases} 1 & \text{if customer } i \in N \text{ is visited by driver } l \in L \text{ at least once during the period} \\ 0 & \text{otherwise} \end{cases}$$

The customer-based assignment problem is formulated as:

$$\min \sum_{i \in N} \sum_{l \in L} \omega_{il} \quad (4a)$$

subject to

$$\omega_{il} \geq \sum_{k \in K} \hat{y}_{ik}^t \pi_{lk}^t \quad \forall i \in N; l \in L; t \in T, \quad (4b)$$

$$\sum_{k \in K} \pi_{lk}^t \geq 1 \quad \forall l \in L; t \in T \quad (4c)$$

$$\sum_{l \in L} \pi_{lk}^t \leq 1 \quad \forall k \in K; t \in T \quad (4d)$$

$$\pi_{lk}^t \in \{0, 1\} \quad \forall l \in L; k \in K; t \in T \quad (4e)$$

$$\omega_{il} \geq 0 \quad \forall i \in N; l \in L \quad (4f)$$

The objective (4a) minimizes the number of customers visited by the drivers. Constraints (4b) set ω_{il} to 1 if driver $l \in L$ is assigned to vehicle route $k \in K$ that visits customer $i \in N$ on at least one day. Constraints (4c) - (4e) are the same as (3c) - (3e) in the region model. Constraints (4f) define the decision variables (again ω_{il} is binary, given binary values for π_{lk}^t).

3.4. Customer familiarity: PVRP(CF)

The PVRP(CF) incorporates the customer-driver relationship into the initial route design, rather than performing a post-processing step. The two-phase approach of the PVRP(TD) is limited in achieving customer familiarity, as travel cost is the only metric considered when designing routes in Phase I. The PVRP(CF) seeks to maximize the frequency with which a driver visits a customer.

Recall the decision variable ω_{il} from Phase II of the PVRP(TD) that matches drivers to customers. The PVRP(CF) incorporates this matching in the routing decisions. Since the workforce management models of the PVRP treat drivers and vehicle routes as a single entity, the notation $l(k)$ is introduced denoting the driver of vehicle route k . Note that only one driver is associated with a vehicle route over the time period, although the customers visited and the sequence of visits may change by day for the driver-route entity. An auxiliary variable for driver-customer assignment is introduced that is distinguished by the frequency with which a driver visits a customer:

$$v_{il(k)}^n = \begin{cases} 1 & \text{if customer } i \in N \text{ is visited by the driver of route } k \in K, \text{ denoted } l(k), \\ & \text{a total of } n \in T \text{ times in the period} \\ 0 & \text{otherwise} \end{cases}$$

Note that the learning effect often dampens as the visiting frequency increases; i.e. the marginal benefit of increasing visiting frequency decreases as the frequency increases (see Zhong et al. (2007) for a discussion of driver learning behavior). Therefore, the customer access cost for a driver is modeled as a convex function of the frequency with which a customer is visited. Let α^n denote the customer access cost per visit if a customer is visited by the driver n times in the period, where the parameter values are chosen such that

$$\alpha^n > \alpha^{n+1}, \quad \forall n \in T.$$

Assume a homogeneous set of customers such that the values are the same for each customer. Further limits on the values ensure that the model does not artificially increase visits to customers to lower access costs:

$$n\alpha^n \leq (n+1)\alpha^{n+1}, \quad \forall n \in T.$$

Combining the two constraints above, the relationship between α^n and α^{n+1} is:

$$\frac{n}{n+1} \leq \frac{\alpha^{n+1}}{\alpha^n} < 1, \quad \forall n \in T. \quad (5)$$

The total customer access cost for the drivers is:

$$f_{\text{CF}} = \sum_{i \in N} \sum_{k \in K} \sum_{n \in T} n\alpha^n v_{il(k)}^n. \quad (6)$$

With this expression, the PVRP(CF) is as follows:

$$\min f_{\text{TD}} + \phi_{\text{CF}} f_{\text{CF}} \quad (7a)$$

subject to (2b)- (2h)

$$\sum_{t \in T} y_{ik}^t = \sum_{n \in T} n v_{il(k)}^n \quad \forall i \in N; k \in K \quad (7b)$$

$$v_{il(k)}^n \in \{0, 1\} \quad \forall i \in N; k \in K; n \in T; \quad (7c)$$

The objective function (7a) minimizes the total customer access cost for the drivers and the travel cost of the resulting solution, weighted by the parameter ϕ_{CF} . Constraints (7b) define the driver visit frequency variables, $v_{il(k)}^n$, based on the assignment variables. Constraints (2b) set the sum of the y_{ik}^t variables over k and t equal to the number of visits requested by customer i over the time period, such that the left hand side of constraints (7b) is the total number of visits by vehicle route k to customer i . Constraints (7c) define the binary variables for visit frequency by driver.

Optimal PVRP(CF) solutions will aggregate visits to a customer to as few drivers as possible during the period, as defined through the following lemmas:

LEMMA 1. *Given that $\alpha^n < \alpha^m$ when $n > m$, then $n = n_1 + n_2$ implies $n\alpha^n < n_1\alpha^{n_1} + n_2\alpha^{n_2}$.*

Proof:

$$n\alpha^n = (n_1 + n_2)\alpha^n = n_1\alpha^n + n_2\alpha^n < n_1\alpha^{n_1} + n_2\alpha^{n_2}$$

□

LEMMA 2. *Given the limits on the α^n values defined by (5), then*

$$\sum_{n \in T} v_{il(k)}^n \leq 1, \quad \forall i \in N; k \in K.$$

Proof: This follows directly from Lemma 1, as for each i, k pair there exists at most one n value such that $v_{il(k)}^n = 1$. \square

LEMMA 3. *If ϕ_{CF} is arbitrarily large, each customer's contribution to f_{CF} is minimized when the number of drivers visiting that customer is minimized.*

Proof: From Lemma 1 and Lemma 2, there exists a driver $l(k)$ such that

$$n\alpha^n v_{il(k)}^n < n_1\alpha^{n_1} v_{il_1(k)}^{n_1} + n_2\alpha^{n_2} v_{il_2(k)}^{n_2} \quad \forall i \in N,$$

where $n = n_1 + n_2$. The cost function is then minimized for customer i with driver $l(k)$ rather than drivers $l_1(k)$ and $l_2(k)$. \square

We show below that the PVRP(CF) minimizes the number of different drivers serving each customer. Let $w_{il(k)} = 1$ if customer $i \in N$ is visited by the driver of vehicle route $k \in K$, denoted $l(k)$, at least once in the time period; and 0 otherwise.

LEMMA 4. *In the optimized case, for every $i \in N$ and $k \in K$, $\sum_{n \in T} v_{il(k)}^n = w_{il(k)}$.*

Proof: When $w_{il(k)} = 0$, then it is optimal to set $v_{il(k)}^n = 0$ for all n . By definition, when $w_{il(k)} > 0$, there exists at least one n such that $v_{il(k)}^n \neq 0$. By Lemma 2, $\sum_{n \in T} v_{il(k)}^n \leq 1$. Therefore, in the optimized case, $\sum_{n \in T} v_{il(k)}^n = w_{il(k)}$. \square

THEOREM 1. *The PVRP(CF) minimizes the number of drivers visiting each node.*

Proof: Given optimal solutions to a PVRP that minimizes the number of drivers and the PVRP(CF) denoted $\{\mathbf{x}^{\text{MinDriv}}, \mathbf{y}^{\text{MinDriv}}, \mathbf{v}^{\text{MinDriv}}, \mathbf{w}^{\text{MinDriv}}\}$ and $\{\mathbf{x}^{\text{CF}}, \mathbf{y}^{\text{CF}}, \mathbf{v}^{\text{CF}}, \mathbf{w}^{\text{CF}}\}$, respectively.

Assume that the PVRP(CF) does not minimize the number of drivers visiting each node; i.e., that

$$\sum_{i \in N} \sum_{k \in K} w_{il(k)}^{\text{MinDriv}} < \sum_{i \in N} \sum_{k \in K} w_{il(k)}^{\text{CF}}.$$

Given that $w_{il(k)} = \sum_{n \in T} v_{il(k)}^n$, then:

$$\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} v_{il(k)}^{n, \text{MinDriv}} < \sum_{i \in N} \sum_{k \in K} \sum_{n \in T} v_{il(k)}^{n, \text{CF}}.$$

However, it is shown next that $\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} n\alpha^n v_{il(k)}^n$ is minimized when $\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} v_{il(k)}^n$ is minimized, which contradicts the original assumption, and the PVRP(CF) must minimize the

number of drivers visiting each node. The expression $\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} n \alpha^n v_{il(k)}^n$ can be decomposed by customer for a feasible routing solution given that all customers incur the same values of α^n for all n . Lemma 3 states that, for a given customer, $\sum_{k \in K} \sum_{n \in T} n \alpha^n v_{il(k)}^n$ is minimized when $\sum_{k \in K} \sum_{n \in T} v_{il(k)}^n$ is minimized. Therefore, summed over all customers, $\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} n \alpha^n v_{il(k)}^n$ is minimized when $\sum_{i \in N} \sum_{k \in K} \sum_{n \in T} v_{il(k)}^n$ is minimized. \square

It can easily be shown by counter-example that minimizing the number of drivers does not necessarily maximize the frequency with which a driver visits a node, see Smilowitz et al. (2009).

3.5. Region familiarity: PVRP(RF)

The PVRP(CF) creates vehicle routes that focus on the relationship between a driver and a customer. The same concept can apply to a driver's familiarity with the network. Practically, when a driver visits an area more frequently, her familiarity with the area increases and hence the travel time decreases. With parcel delivery, drivers often work in the same region with varying customers. By visiting one customer in the region, the driver actually gains familiarity with the entire area. The PVRP(CF) does not consider that drivers become more efficient without paying a direct visit to the same customers repeatedly. This effect is reflected in the PVRP(RF).

In the PVRP(RF), we introduce a set of variables similar to $\nu_{r,l}$ from Phase II of the PVRP(TD) that matches drivers to regions; however, a superscript denoting the frequency of visits to the region is added:

$$u_{rl(k)}^n = \begin{cases} 1 & \text{if region } r \in R \text{ is visited by the driver of vehicle route } k \in K, \text{ denoted } l(k), \\ & \text{a total of } n \in T \text{ times in the period} \\ 0 & \text{otherwise} \end{cases}$$

In addition, we use auxiliary variables similar to \hat{a}_{rk}^t from Phase II of the PVRP(TD):

$$a_{rk}^t = \begin{cases} 1 & \text{if route } k \in K \text{ visits region } r \in R \text{ on day } t \in T \\ 0 & \text{otherwise} \end{cases}$$

Similar to the customer access cost, a region access cost is introduced that is a convex function of the frequency with which a driver visits a region within the time period. Recall N_r is the set of customers in region r and let β^n represent the region access cost per visit for each region that a driver visits n times. The access cost parameters reflect the concept that, as a driver visits a region more frequently, the driver becomes more familiar with the region and can travel through the region more efficiently. However, the PVRP(RF) should not create solutions in which regions are visited more frequently than needed to reduce access cost. The limits on β^n are similar to α^n :

$$\frac{n}{n+1} \leq \frac{\beta^{n+1}}{\beta^n} < 1, \quad \forall n \in T \quad (8)$$

The total region access cost for the drivers is:

$$f_{\mathbf{RF}} = \sum_{r \in R} \sum_{k \in K} \sum_{n \in T} n \beta^n u_{rl(k)}^n. \quad (9)$$

The PVRP(RF) is modeled as follows:

$$\min f_{\mathbf{TD}} + \phi_{\mathbf{RF}} f_{\mathbf{RF}} \quad (10a)$$

subject to (2b)- (2h)

$$\sum_{i \in N_r} y_{ik}^t \leq |N_r| a_{rk}^t \quad \forall r \in R; k \in K; t \in T \quad (10b)$$

$$\sum_{t \in T} a_{rk}^t = \sum_{n \in T} n u_{rl(k)}^n \quad \forall r \in R; k \in K \quad (10c)$$

$$u_{rl(k)}^n \in \{0, 1\} \quad \forall r \in R; k \in K; n \in T \quad (10d)$$

The objective function (10a) minimizes the region access cost and the travel distance, weighted by $\phi_{\mathbf{RF}}$. Constraints (10b) link the region assignment variables a_{rk}^t , to the node assignment variables, y_{ik}^t : if at least one node within region k is visited by route k on day t , then $a_{rk}^t = 1$. Given the limits on route access cost, $a_{rk}^t = 0$ if no nodes within the region are visited to minimize the objective function. Constraints (10c) define the region visit frequency variables, $u_{rl(k)}^n$, based on the region assignment variables. Constraints (10d) define the binary variables for visit frequency by driver. Lemmas 1-4 and Theorem 1 can be easily adapted to show that the PVRP(RF) minimizes the number of drivers visiting a region.

3.6. Analysis of the PVRP models

In this section, a general analysis of the PVRP models is presented. As the intention here is to study the basic trade-offs between cost and workforce metrics, assume the weighting factors $\phi_{\mathbf{CF}}$ and $\phi_{\mathbf{RF}}$ are arbitrarily large such that travel costs do not impact the decision making for the workforce models. The empirical study in Section 4 evaluates the weighting factor values.

Figure 3 presents an application of the PVRP models for the example in Figure 1. The performance metrics are compared in Table 1, in which we calculate performance metrics for the solutions derived from each model. The value in bold represents the intended goal for each model. We analyze these solutions in the following subsections, and, where possible, generalize our findings.

3.6.1. Route usage As shown in Figure 3, on day 1, the capacity constraint and demand levels require two routes across all models. However, the total demand on day 2 equals the capacity of one vehicle. The PVRP(TD) uses only a single route to minimize travel cost, while the other

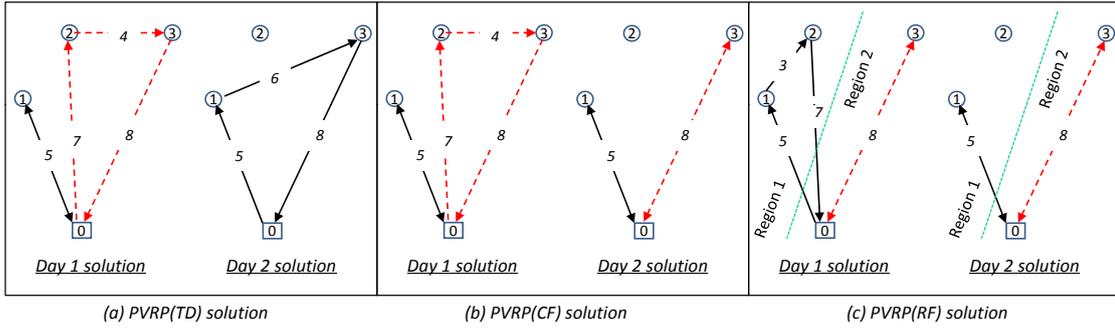


Figure 3 Routing solutions for example from Figure 1

		Performance metric		
Figure 3		f_{TD}	f_{CF}	f_{RF}
PVRP(TD)	3(a)	48	$2\alpha^2+3\alpha^1$	$2\beta^2+2\beta^1$
PVRP(CF)	3(b)	55	$4\alpha^2+\alpha^1$	$2\beta^2+2\beta^1$
PVRP(RF)	3(c)	57	$4\alpha^2+\alpha^1$	$4\beta^2$

Table 1 Comparison of PVRP models

models use two routes to achieve workforce goals. A primary objective of many VRP models is to minimize the number of routes used, due to the fixed cost of each route. A manager may be concerned that incorporating workforce management can require the use of a greater number of routes. Focusing on routing cost may lead to solutions using fewer routes, which can make it difficult to serve a customer with one driver. This is evident in the two phase approach for the PVRP(TD). By focusing on a minimum distance cost solution in Phase I, the number of routes used across days may differ, limiting the workforce management options in Phase II. In the example, only one route is used on day 2 in the PVRP(TD) solution, which limits matching customers with drivers in Phase II. While the focus of PVRP(TD) may result in lower routing costs, the limited flexibility associated with workforce management may be a concern. Alternatively, focusing on workforce management may lead to solutions with more vehicles than needed. It should be noted that minimizing routing cost does not always minimize the number of routes. However, given the triangle inequality, there is often an incentive to use fewer routes if capacity allows such an option.

3.6.2. Relationship between region and customer based models In the example, the PVRP(RF) solution results in the same f_{CF} cost as the PVRP(CF) model. Solving the PVRP(RF) may be simpler computationally given that, unless there is one customer per region, the number of regions is smaller than the number of customers. However, the PVRP(RF) does not always find the same solution as the PVRP(CF). If the regions are very large, containing many customers, it is evident that there will be multiple combinations of customers that may be serviced while drivers

operate continuously within one region. Even as the region size decreases to the point where there are very few customers per region, a counterexample may be found that shows a situation where the PVRP(RF) does not find the same solution as the PVRP(CF).

In Section 4, these observations are extended to larger problem instances.

4. Numerical Analysis

This section presents a comparison of the PVRP models defined in Section 3 using a set of empirical tests. Section 4.1 describes the computational tests. Section 4.2 presents the modified Tabu Search heuristic used to solve the PVRP models. Finally, Section 4.3 presents a full analysis of the results.

4.1. Experimental design

4.1.1. Test case instances The PVRP models are evaluated with the multi-day problem instances used by Groer et al. (2009), adapted from the VRP benchmark problems of Christofides and Eilon (1969). These instances consist of seven sets of customer locations, with the number of customers ranging from 50 to 199. A request for service may come from a customer on any of the five days that compose the time period. A customer has the same demand on any day service is requested. If all customers request service on each of the five days, the problem decomposes to the general VRP, with the same solution used for each day. As with Groer et al. (2009), the probability parameter $p < 1$ denotes the probability that a customer requests service on a given day. If p is too small, customers requesting service on multiple days become too infrequent. In order to create meaningful problem instances, the three values of p selected are: 0.6, 0.7 and 0.8. Each value is tested for each of the seven sets of customer locations, resulting in 21 problem instances.

Five models are tested: (1) PVRP(TD) without Phase II, (2) PVRP(TD) with Phase II optimized at the customer level, denoted as PVRP(TD)+C, (3) PVRP(TD) with Phase II optimized at the region level, denoted as PVRP(TD)+R, (4) PVRP(CF), and (5) PVRP(RF). Note that driver assignments for the PVRP(TD) without Phase II are obtained by setting the driver index l equal to the vehicle route index k from Phase I. Thus, the driver-vehicle route assignments are made arbitrarily without regard to workforce metrics.

4.1.2. Access cost functions The workforce management objectives require a quantification of driver learning. In practice, it can be difficult to determine how quickly a driver adapts to a customer or region, whether the biggest drop in cost per visit occurs after one visit or multiple visits. For the PVRP(CF), we develop four sets, $\alpha(1)$, $\alpha(2)$, $\alpha(3)$, and $\alpha(4)$, to represent driver-customer learning curves over a five-day period; i.e., $\alpha(1) = \{\alpha^1(1), \alpha^2(1), \alpha^3(1), \alpha^4(1), \alpha^5(1)\}$. To

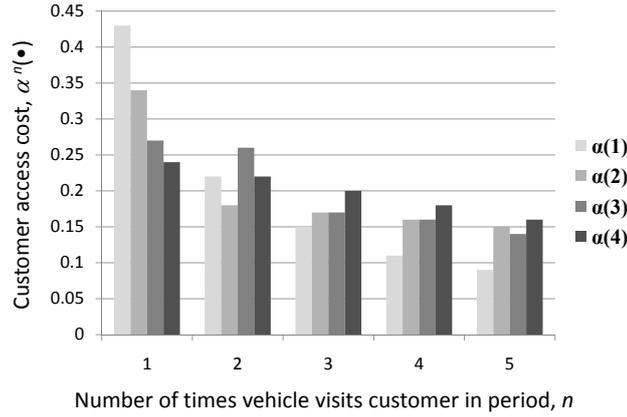


Figure 4 Customer access cost functions

maintain proportions among the sets, the five values within each set sum to one; i.e., $\alpha^1(1) + \alpha^2(1) + \alpha^3(1) + \alpha^4(1) + \alpha^5(1) = 1$. The four sets are shown in Figure 4.

The first set, $\alpha(1)$, is generated by letting $1\alpha^1 = 2\alpha^2 = \dots = n\alpha^n$, such that an objective with this cost function simply minimizes the number of drivers visiting a node. The cost does not change for a driver if she visits the same node multiple times. Each new driver visiting the node incurs an additional cost. The design of this set results in large drops between each of the α values (see Figure 4). The second set, $\alpha(2)$, represents a scenario where the driver gains most familiarity with one visit, with the learning curve dropping off considerably with subsequent visits. Therefore, this set has a high access cost for a driver who visits a customer only once in the period, with a large drop in customer access cost if the driver visits the customer twice and a marginal decrease in cost with additional visits. The third set, $\alpha(3)$, is similar, but the most familiarity is gained after two visits, with the large drop in access cost from two to three visits to a customer within the time period. The costs drop linearly between all visits for the fourth set, $\alpha(4)$. The driver-region learning sets, $\beta(1)$, $\beta(2)$, $\beta(3)$, and $\beta(4)$, are constructed identically for the PVRP(RF).

4.1.3. Multi-objective weighting factors The three costs associated with travel, customer familiarity and region familiarity are of differing orders of magnitude. Travel costs measure distance while customer familiarity costs measure the number of visits per customer. Region familiarity costs also differ from customer familiarity costs as these costs are incurred less frequently (multiple customers can be served in one visit to a region). These differences result in costs that vary by function, with f_{TD} in the range of 2,000-5,000, f_{CF} approximately 20-200, and f_{RF} approximately 10-40. To regulate the balance between workforce metrics and routing costs, the weighting factors ϕ_{CF} and ϕ_{RF} are applied to the workforce cost portion of the objective functions in PVRP(CF) and PVRP(RF), respectively. Placing too much weight on workforce metrics may lead to a considerable

increase in travel costs. Alternatively, a low weight may result in a solution that is no different from one found using PVRP(TD). To find an appropriate balance, approximately 60 values of $\phi_{\mathbf{CF}}$ and $\phi_{\mathbf{RF}}$ are tested. The values of $\phi_{\mathbf{CF}}$ are in the range $[0,500]$, with the majority of tested values below 100. Because the regional access costs are incurred less frequently, the $\phi_{\mathbf{RF}}$ values are extended to the range $[0,1000]$.

4.2. Tabu Search

Tabu Search (TS) is a widely used metaheuristic solution method for the VRP and PVRP. The method begins with a quick construct of an initial solution and then performs an intensive local search by moving from the current step to the best step in the neighborhood. To avoid cycling, certain moves are declared as Tabu and can not be executed for a number of iterations. TS is a flexible algorithm that can be easily modified to suit the needs of specific problems. For example, there are various definitions of search methods for the neighborhood, such as GENI [Gendreau et al. (1992)], which has demonstrated effectiveness in a range of VRP extensions.

4.2.1. Algorithm implementation The TS implementation for the PVRP models defined in Section 3 is similar to that presented by Gendreau et al. (1997). The principal changes to the algorithm here are the definition of a move and the cost functions. In Gendreau et al. (1997) a move is defined as either a change in the vehicle or schedule assignment; we only consider a change in vehicle assignment as the problem does not allow flexibility in visitation schedules. Similar to Gendreau et al. (1997), infeasible solutions are allowed during the search. The cost function $c(s)$ is a combination of the objective function $f(s)$ and an infeasibility penalty $q(s)$ caused by violation of vehicle capacity constraints, weighted by the infeasibility penalty factor, Φ . In Gendreau et al. (1997) the objective function is the travel cost, but here workforce metrics are also added. As an example, the cost function for the PVRP(CF) model is presented:

$$c(s) = f(s) + \Phi q(s) \quad (11a)$$

$$f(s) = \min \sum_{(i,j) \in A} \sum_{k \in K} \sum_{t \in T} c_{ij} x_{ijk}^t + \phi_{\mathbf{CF}} \sum_{i \in N} \sum_{k \in K} \sum_{n \in T} n \alpha^n v_{il(k)}^n \quad (11b)$$

$$q(s) = \sum_{k \in K} \sum_{t \in T} [(\sum_{i \in N} d_i^t y_{ik}^t - C)]^+ \quad (11c)$$

The remaining components of the TS algorithm used by Gendreau et al. (1997) are similar, including the use of an aspiration level, which guides the search to feasible solutions, and a scheme to diversify the local search such that the same attributes are not repeatedly selected for improvement. The search procedure is as follows.

1. For each day construct an initial solution using the sweep algorithm [Gillett and Miller (1974)]. Customers are selected for delivery according to their angle with an arbitrary horizontal line at the start of the algorithm, with the customer closest to the horizon selected first and other customers added as the sweep arm rotates away from the horizon. On each day, adjacent customers are assigned to a driver until the driver reaches its capacity limit, then a new driver is used. For the last driver, all remaining customers are assigned to this driver, even if capacity is exceeded.

2. Construct a set of possible moves. There are $|K| - 1$ potential driver moves for a customer on each day service is requested. Drivers containing one of the customer's g nearest neighbors are considered as a possible move destination.

3. Identify the best move and check its Tabu status from the Tabu list. Calculate the change in the objective function for each candidate move using the GENI heuristic with penalties for capacity infeasibility; diversification is also added into the objective function. The move with the least cost that is not Tabu is used. Perform the best move. Remove the customer from the previous route and perform the insertion.

4. Update the Tabu list to include the attribute representing the implemented move. The move is declared Tabu for θ iterations.

5. Check if the current solution is feasible. If so, update the best feasible solution and decrease the infeasibility penalty factor Φ . Otherwise, increase Φ .

6. Return to Step 2 and repeat for λ_{max} iterations.

The parameters used for this TS follow from Gendreau et al. (1997). The factor used to adjust the intensity of the diversification is 0.015. The initial value of the infeasibility penalty weight, Φ , is set to 10. The parameter used to update Φ is set to 0.1. Tabu tenure, θ , is $7.5 \log_{10} m$, where m is the number of customers. The number of iterations, λ_{max} , is 5000, with neighborhood size in GENI fixed at $g = 3$. All instances are tested on a cluster of 6 computers with quad-core 2.4GHz (4MB L2 cache) 64-bit Xeon processors in each computer for a total of 24 available cores. Each computer has 8GB RAM installed.

4.2.2. Tabu Search performance The performance of the TS algorithm is tested using the seven problem instances from Groer et al. (2009), with $p = 0.7$, shown in Table 2. For the purposes of testing performance against cost-based PVRP solutions, the workforce management metrics are removed. An optimal solution has not been determined for the largest instances. Therefore, the results produced by the TS algorithm are compared with the record-to-record (RTR) travel algorithm used by Li et al. (2005), which achieves the best known solution to the seven instances. The average gap for all instances is 1.17% and the largest gap is below 4%. These low gaps allow

us to analyze workforce management objectives among models. The computational times are not reported by Li et al. (2005) for each problem instance; rather, average computational time over all instances is indicated to be less than half a minute. The average computational time over all seven instances for the TS algorithm is 6.3 minutes.

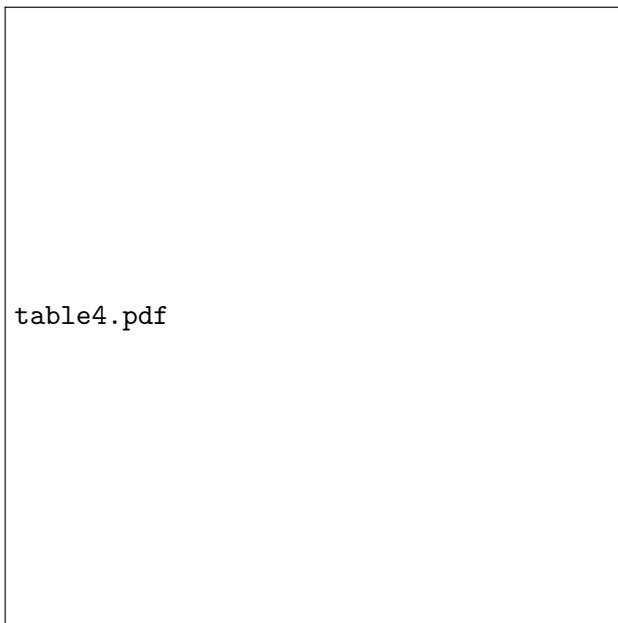


Table 2 Tabu algorithm performance

The computational times for all instances are presented in Table 6 in the Appendix. Most of the variations in times are a reflection of a change in instance size. As the number of customers increases and as p increases, increasing the number of requests per customer, computational times increase. When $\phi_{\mathbf{CF}} = 25$ and $\phi_{\mathbf{RF}} = 50$ the times for the workforce management models are comparable to the times for the PVRP(TD) model. These times are only for Phase I of this model, as Phase II is performed using AMPL. Postprocessing for Phase II requires 15-20 minutes for the smaller instances, but can take several hours for the larger instances.

4.3. Model analysis

In the following subsections, we evaluate the application of workforce management principles to the PVRP. As noted earlier, comparing the objective values of each model does not provide a useful analysis, as these values differ in magnitude based on what they measure. Rather, we compare relative gaps as defined next.

Let $\Delta_{\mathbf{I}}(\text{PVRP}(J_{\phi}))$ represent the relative gap in the value of objective I between a PVRP solution in which J is the primary objective (with designated weight ϕ), denoted $f_{\mathbf{I}}(\text{PVRP}(J_{\phi}))$, and

the minimum value of objective I obtained for the instance, denoted f_I^{MIN} . When $J = TD$, the parameter ϕ is dropped. The gap is calculated as follows:

$$\Delta_I(\text{PVRP}(J_\phi)) = \frac{f_I(\text{PVRP}(J_\phi)) - f_I^{\text{MIN}}}{f_I^{\text{MIN}}}.$$

As an example, the relative gap in travel distance ($I = \text{TD}$) between the $\text{PVRP}(\text{TD})$ and the $\text{PVRP}(\text{CF})$ for weight ϕ is:

$$\Delta_{\text{TD}}(\text{PVRP}(\text{CF}_\phi)) = \frac{f_{\text{TD}}(\text{PVRP}(\text{CF}_\phi)) - f_{\text{TD}}^{\text{MIN}}}{f_{\text{TD}}^{\text{MIN}}}$$

A large value of $\Delta_{\text{TD}}(\text{PVRP}(\text{CF}_\phi))$ indicates that the $\text{PVRP}(\text{CF})$ with weight ϕ produces solutions with large travel costs relative to the solution that has minimum travel cost. A value close to zero indicates that the $\text{PVRP}(\text{CF})$ produces solutions that achieve customer familiarity at a small increase in travel costs.

4.3.1. Impact of the multi-objective weighting factor The relative objective gaps are used to evaluate the multi-objective weighting factors, ϕ_{CF} and ϕ_{RF} , for the respective models. The impact of the weight factors on relative gaps is illustrated in Figure 5 for the test cases. Figure 5(a) presents the trade-off between routing costs and customer familiarity as ϕ_{CF} increases from 0 to 500 when solving the $\text{PVRP}(\text{CF})$. Recall that $\phi_{\text{CF}} = 0$ corresponds to the single-objective $\text{PVRP}(\text{TD})$ model. The x-axis measures $\Delta_{\text{TD}}(\text{PVRP}(\text{CF}_\phi))$, the deviation in travel distance between the minimum travel distance and the travel distance achieved with the $\text{PVRP}(\text{CF})$ for a given weight ϕ . The y-axis measures $\Delta_{\text{CF}}(\text{PVRP}(\text{CF}_\phi))$, the deviation in customer familiarity between the minimum value (over all weights) and the value achieved for a given weight. Each data point represents the average deviation over all seven instances for $p = 0.7$, with $\alpha(\mathbf{1})$. As ϕ_{CF} increases, the relative importance of customer familiarity increases. This figure indicates the amount by which routing costs increase, $\Delta_{\text{TD}}(\text{PVRP}(\text{CF}_\phi))$, and workforce costs decrease, $\Delta_{\text{CF}}(\text{PVRP}(\text{CF}_\phi))$, as ϕ_{CF} increases. Figure 5(b) presents the same analysis for region familiarity, analyzing $\text{PVRP}(\text{RF})$ models with ϕ_{RF} from 0 to 1000 and $\beta(\mathbf{1})$.

In both Figures 5(a) and (b), we observe that as ϕ_{CF} and ϕ_{RF} increase and more emphasis is placed on workforce costs, the Δ_{TD} values increase, although the initial increases from zero are small. Further, Δ_{CF} and Δ_{RF} approach zero while Δ_{TD} continues to increase, indicating that there are diminishing returns in placing a greater emphasis on workforce management. Particularly in the cases of the $\text{PVRP}(\text{CF})$, introducing workforce metrics into routing decisions with a small weight has a high reward with a low cost. Increasing ϕ_{CF} slightly from zero can bring workforce gaps down to single digits, while ϕ_{RF} shows a more gradual drop. The difference between these

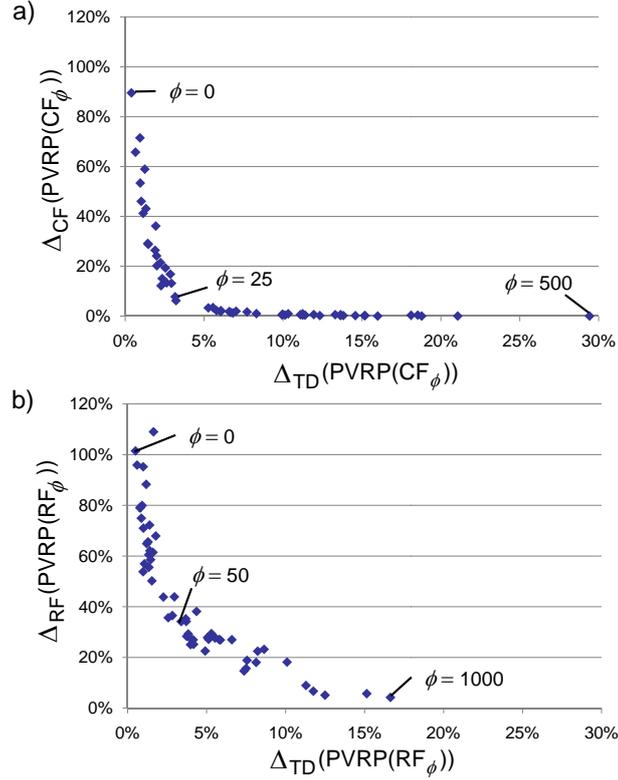


Figure 5 Impact of multi-objective weighting on: (a) customer familiarity; (b) region familiarity

two models may be attributed to the larger cost values for customer familiarity, with a smaller ϕ_{CF} value required in order for these costs to have an impact on total cost.

The trade-off analysis leads to a question of how to determine the appropriate ϕ value to balance between travel and workforce management costs. We introduce a metric which combines the deviations in travel and workforce costs. The Δ values are normalized for the comparison of travel distance and customer familiarity as follows:

$$NORM_\phi = \frac{f_{TD}(PVRP(CF_\phi))}{f_{TD}^{MIN}} + \frac{f_{CF}(PVRP(CF_\phi))}{f_{CF}^{MIN}}.$$

A similar normalization is used for PVRP(RF). Figures 6(a) and (b) present the normalized values for ϕ_{CF} and ϕ_{RF} , respectively. As with Figure 5, data points represent the average over all seven instances for $\alpha(1)$ and $\beta(1)$, respectively, and both have $p = 0.7$. It can again be seen that a slight increase in ϕ from zero has an immediate impact when both travel and workforce management costs are considered. The values that provide the best balance between workforce management and travel costs are $\bar{\phi}_{CF} = 25$ and $\bar{\phi}_{RF} = 50$. As ϕ increases above these values, the emphasis on workforce management becomes too great and travel costs disproportionately increase. We use these weighting factors to further analyze the relationships among the three models next.

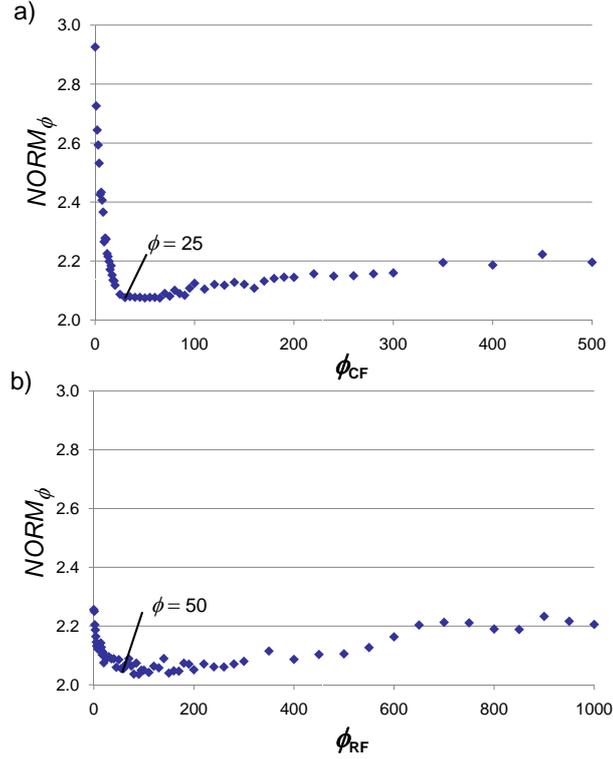


Figure 6 Evaluation of ϕ value to balance routing and workforce management objectives for: (a) customer familiarity; (b) region familiarity

4.3.2. Metric comparisons among all models To compare the PVRP models, we analyze the objective gaps relative to the minimum metric values obtained as defined by $\Delta_I(\text{PVRP}(J_\phi))$. Table 3 presents the values of $\Delta_I(\text{PVRP}(J_\phi))$ for the objective functions ($I = \text{TD}, \text{CF}, \text{and RF}$) obtained with five model variations ($J = \text{TD}, \text{TD}+\text{R}, \text{TD}+\text{C}, \text{CF}, \text{and RF}$). The results are initially aggregated across problem instances for the three values of p . The results in Table 3 are presented for $\phi = \bar{\phi}_{CF} = 25$ for the PVRP(CF) and $\phi = \bar{\phi}_{RF} = 50$ for the PVRP(RF). The CF and RF metrics are evaluated with $\alpha(1)$ and $\beta(1)$, respectively. Other α and β sets are evaluated in Section 4.3.3.

	$p=0.6$			$p=0.7$			$p=0.8$		
	Δ_{TD}	Δ_{CF}	Δ_{RF}	Δ_{TD}	Δ_{CF}	Δ_{RF}	Δ_{TD}	Δ_{CF}	Δ_{RF}
PVRP(TD)	1.3%	64.7%	104.1%	0.6%	89.6%	103.0%	3.2%	85.6%	97.3%
PVRP(TD)+R	1.3%	46.2%	49.1%	0.6%	74.8%	75.0%	3.2%	80.3%	77.3%
PVRP(TD)+C	1.3%	39.2%	52.1%	0.6%	69.9%	82.7%	3.2%	71.4%	83.2%
PVRP(CF ₂₅)	2.9%	5.7%	38.6%	3.4%	7.7%	30.2%	2.0%	4.9%	22.0%
PVRP(RF ₅₀)	3.5%	41.5%	38.8%	3.7%	55.5%	33.3%	5.2%	59.2%	42.7%

Table 3 Objective gaps relative to minimum values for the PVRP models for $\phi_{CF} = 25$ and $\phi_{RF} = 50$

As shown by the values of Δ_{TD} in Table 3, the PVRP models that focus on workforce manage-

ment increase routing costs by no more than 5.2%. These values do not vary significantly by p . On the other hand, the increases in workforce metrics for the PVRP(TD) are considerably greater. This suggests that the models accounting for workforce metrics do a better job of minimizing travel cost than the PVRP(TD) models do with minimizing workforce metrics. As the value of p increases, the values of $\Delta_{\mathbf{CF}}$ increase most dramatically. Increasing p generates more requests per customer over the course of the time period, making it more difficult for the other models to satisfy driver consistency metrics. Note that using ϕ values that balance distance and workforce objectives leads to $\Delta_{\mathbf{I}}(\text{PVRP}(\mathbf{I}_\phi))$ values that are greater than zero.

Comparison of region and customer based models These results show that the relationship between customer and region familiarity models is more complex than that found in the simple example in Table 1. Using ϕ values that balance distance and workforce objectives leads to results such as $\Delta_{\mathbf{RF}}(\text{PVRP}(\mathbf{CF})) < \Delta_{\mathbf{RF}}(\text{PVRP}(\mathbf{RF}))$ for each p . With a smaller disparity between distance and workforce costs, the CF model can put more emphasis on workforce management (both customer and region familiarity) using a smaller $\phi_{\mathbf{CF}}$ value of 25, relative to the $\phi_{\mathbf{RF}}$ value of 50. At the highest ϕ values, which mimic pure workforce models similar to those found in Table 1, this is no longer the case and the RF model leads to the lowest region familiarity cost. However, $\Delta_{\mathbf{RF}}(\text{PVRP}(\mathbf{CF})) > 0$ and $\Delta_{\mathbf{CF}}(\text{PVRP}(\mathbf{RF})) > 0$, regardless of ϕ value, indicating that one model does not achieve the goals of the other. Much of this is driven by the value of p .

As the value of p increases and customers request service more frequently, $\Delta_{\mathbf{RF}}(\text{PVRP}(\mathbf{CF}))$ decreases. As p goes from 0.6 to 0.8, the average number of requests per region per day increases from 2.7 to 3.6. The PVRP(CF) model attempts to serve a customer repeatedly with the same driver, and with more requests per customer within a region, there are more opportunities for a driver to serve a specific customer within a region, decreasing the region familiarity costs. Inversely, as p increases, $\Delta_{\mathbf{CF}}(\text{PVRP}(\mathbf{RF}))$ increases. Increasing the number of requests per region per day from 2.7 to 3.6 results in a wider variety of customers from which to choose for service in the region every day. In this case, there is a greater likelihood that a driver may repeatedly visit the same region without having to visit the same customer each time. Therefore, as customers request service more often, a driver should become more familiar with the region if an attempt is made to have the driver visit the same customers regularly, but the driver will not necessarily become more familiar with a set of customers if assignments are based on regions. For the CF and RF metrics the impact of p is complicated by the α and β structures. This impact is discussed in more detail in Section 4.3.3.

Impact of post-processing The results in Table 3 indicate that it is beneficial to apply Phase II to the PVRP(TD) solutions to improve workforce metrics. Customer-based post-processing is more effective at reducing driver consistency costs (a customer-dependent term) than region-based post-processing. This trend is less strong for customer familiarity. Region-based post-processing is slightly more effective at reducing region familiarity costs than customer-based post-processing.

Neither post-processing method is able to reduce workforce metrics to the levels achieved with PVRP models with workforce objectives. As discussed in Section 3.6, the initial routes obtained with the PVRP(TD) tend to use fewer vehicles than those solutions obtained with the workforce PVRP models. Table 4 presents the average number of vehicles used in routing solutions for each of the seven test cases for $p = 0.7$ with two sets of weighting factors. The PVRP(CF) and PVRP(RF) results are averaged over all α and β values, respectively.

Instance	PVRP(TD)	Low weighting		High weighting	
		$\phi_{CF} = 25$	$\phi_{RF} = 50$	$\phi_{CF} = 500$	$\phi_{RF} = 1000$
		PVRP(CF)	PVRP(RF)	PVRP(CF)	PVRP(RF)
ConVRP1	3.2	3.6	3.6	4.2	4.7
ConVRP2	7.8	7.9	7.8	8.0	8.0
ConVRP3	5.5	5.8	5.7	6.0	6.0
ConVRP4	8.6	9.3	8.6	9.9	9.1
ConVRP5	12.2	12.6	12.2	13.8	12.8
ConVRP11	5.4	6.1	5.2	6.8	5.9
ConVRP12	6.6	7.0	6.9	7.8	7.4

Table 4 Average number of vehicles used

The values in Table 4 confirm that fewer vehicles are more likely to be used in PVRP(TD) solutions than in solutions obtained with workforce management objectives. The difference is more pronounced when a higher weight is placed on workforce metrics. When the difference is large, the opportunities to meet workforce goals with post-processing are more limited. This can further explain the increases in workforce gaps in Table 4.

While fewer vehicles can limit the ability of the PVRP(TD) to achieve workforce goals through post-processing, more vehicles in the workforce PVRP models do not necessarily translate into large increases in distance traveled. With low weightings, although the number of vehicles increases with the workforce models by an average of 6%, the rise in routing costs is small, with an average Δ_{TD} of 2.4%.

4.3.3. Comparison of driver-customer learning models The set of α values can dramatically change routing solutions. The four α sets, $\alpha(1)$, $\alpha(2)$, $\alpha(3)$ and $\alpha(4)$, are differentiated by the reduction in access cost with increased visits to a customer. For example, while $\alpha(2)$ has the largest drop in cost between one and two visits, $\alpha(3)$ has the greatest drop from two to three visits. One might expect that the latter is more effective when customers request more visits, while the former with fewer visits. These differences are quantified by comparing the number of requests made by a customer in a period and the maximum number of those requests served by one driver. For example, if a customer requesting service on all five days of the period has one driver visit three times and another driver visit two times, the resulting metric would be $5/3$. Each customer falls into one of 15 categories: $5/5$, $5/4$, $5/3$... $2/2$, $2/1$, $1/1$.

Figure 7 presents results for the four α sets for $\phi_{CF} = 25$ along with PVRP(TD)+C for comparison. Good workforce management is achieved when the number of visits requested is equivalent to

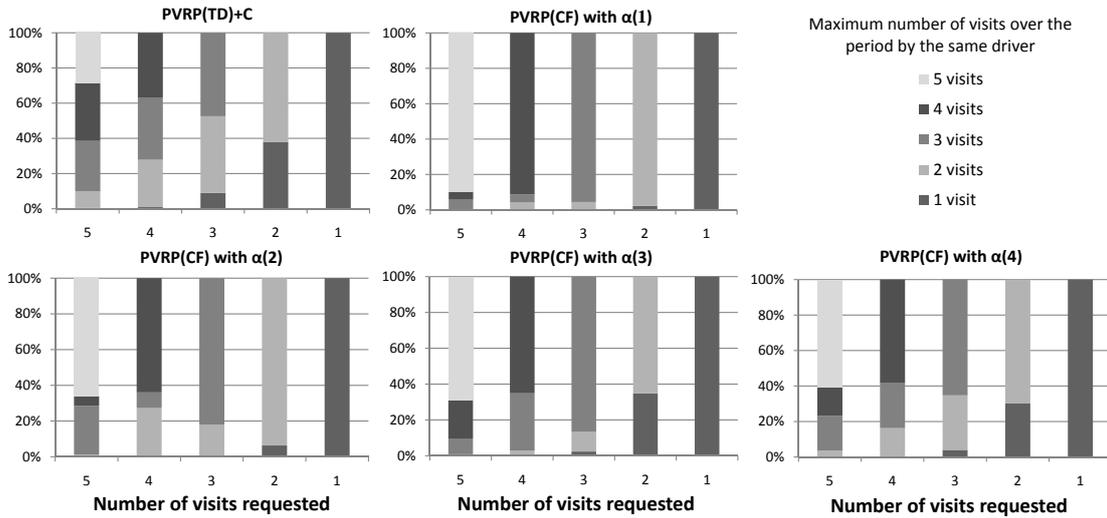


Figure 7 Maximum number of visits made by one driver to a customer ($\phi_{CF} = 25$)

the number of visits made by one driver. In the figure, this is reflected when the uppermost colored bar is the predominant one, as in the graph for $\alpha(1)$. The PVRP(TD)+C model results in a mix of the number of drivers visiting each customer. The PVRP(CF) with $\alpha(1)$ clearly outperforms the other PVRP(CF) models in serving most customers with the same driver. When revisiting the α sets shown in Figure 4, these results are not surprising. Given that $1\alpha^1 = 2\alpha^2 = \dots = n\alpha^n$ for the $\alpha(1)$ set, there is a more significant drop in cost between values. For example, while $\alpha(1)$ and $\alpha(2)$ have similar relative drops in value from one visit to two visits, the absolute drop in $\alpha(1)$ is considerably larger (0.22 vs. 0.17). These differences have an impact, even at small values of ϕ_{CF} .

When using $\alpha(2)$, the model finds the best results from a workforce management perspective when two visits are requested, which is where the largest drop in cost occurred. Similarly, when the $\alpha(3)$ parameters are applied, the model finds the best results when three visits are requested. With relatively little drop between values, $\alpha(4)$ provides far less incentive for repeated visits. These results further underline the importance of the ratio between consecutive α^n values.

Consider the constraint placed on the α^n values when constructing the PVRP(CF) model:

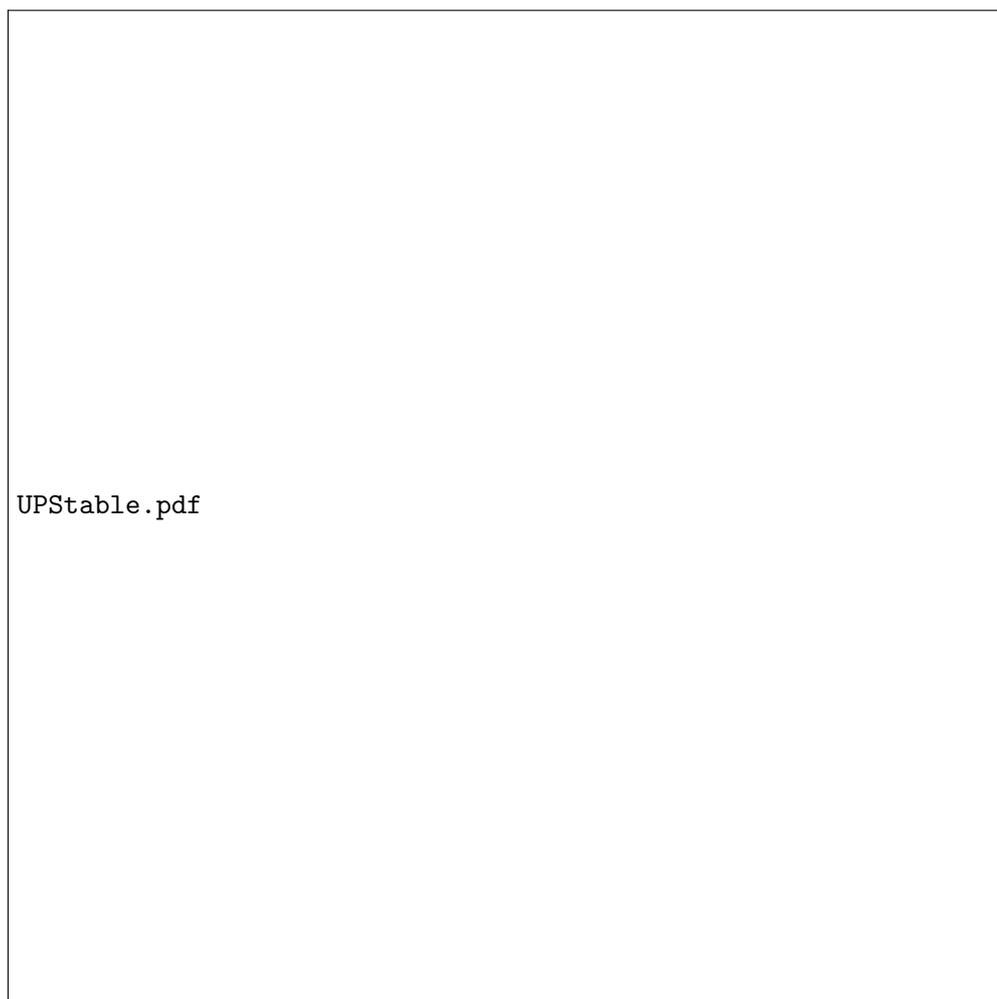
$$\frac{n}{n+1} < \frac{\alpha^{n+1}}{\alpha^n} < 1, \quad \forall n \in T$$

As the ratio $\frac{\alpha^{n+1}}{\alpha^n}$ decreases and approaches $\frac{n}{n+1}$, the benefit from reducing the number of different drivers visiting a customer increases. It is clear that the heuristic is very responsive to this ratio for all values of n .

4.3.4. Industry Analysis The preceding results indicate that a slight emphasis on workforce management leads to a reduction in related costs with a minimal increase in travel costs. In addition, post processing of distance minimizing routes is limited in lowering workforce management costs. We evaluate the extent to which these insights are observed in practice with data provided by a major package delivery carrier. The data set represents five days of operation consisting of 1800 deliveries, serving over 1100 unique customers with three routes per day. Of these 1800 deliveries, only 22% have time windows. Therefore, consistent with our PVRP models, time windows are ignored.

In order to apply the heuristic to this instance, the large data set is aggregated to a set of 197 customers requesting 737 deliveries over the five days with the following customer aggregation criteria. Customers are clustered with their neighbors, such that the average distance from any customer in a cluster to the center of the cluster is less than or equal to 1 distance unit and the total number of customers in a cluster is less than 11. The center is defined by the average x and y coordinates of all customers in a cluster. Customers are combined only if they appear on the same route on each of the days that they request service. As such, we can create a base case for comparison in which each customer cluster is assigned to the same route as in the original assignment. The sequence of cluster visits within each route is improved using a simple k-opt algorithm. We compare this base case with results from the PVRP(TD) and PVRP(RF) models. We focus on region familiarity rather than customer familiarity for two reasons; first, the data set is already an aggregation of true customers as described above, and second, the post-processing of PVRP(TD) routes is performed based on regions due to the size of the problem instance. In all cases, the number of routes is fixed at three per day.

Table 5 provides a summary of the results when the PVRP models are tested with the aggregated package delivery carrier data. For the PVRP(RF) models, the values of ϕ_{RF} range from 1 to 4, to achieve varying balances between travel distance and region familiarity given the relative magnitude of the metrics for this large instance. For each model, the table presents the absolute values of f_{TD} and f_{RF} , as well as the percent improvement in these values compared with the base case.



UPStable.pdf

Table 5 Evaluation of model costs relative to package delivery carrier data

The difference in travel costs for the carrier result and the PVRP(TD) results is likely due, in part, to the omission of time windows in this study. We attempt to improve the sequence of cluster visits in the aggregated carrier solution to make a more fair comparison (i.e., to reduce the impact of time windows). However, recall that customers are only clustered together if they appear on the same disaggregated route, and the disaggregated routes are designed by the carrier with consideration of time windows. All PVRP models ignore time windows. Additionally, the larger travel distances

may arise from a consideration of workforce management costs in the carrier solution that are not present in the PVRP(TD) model. Notably, the value of f_{RF} for the PVRP (TD) does not improve with post-processing. With only three routes, there is little room for improvement, unless the initial assignment of drivers to routes is very poor. The result for the PVRP(RF) model is similar to the results found using the random problem instances, in that workforce management cost savings may be found with little effect on travel costs. In fact, in this case there is no effect on travel cost. Increasing β values beyond those reported led to a minimal reduction in workforce management costs.

The three PVRP model variations provide practitioners with a framework to analyze a wider range of data sets. Importantly, given the results found in the previous section, one may wish to study the impact of p in practice. The likelihood of customer service requests will differ significantly between residential and commercial customers, as will node density. Our modeling framework facilitates the analysis of how workforce management strategies may differ between these two customer classes.

5. Conclusions and Extension

In this paper, we examine how the incorporation of workforce management metrics influences operational decisions for periodic vehicle routing problems. Two PVRP models are proposed that include workforce management metrics as a part of the objective function for comparison with a base PVRP model focused on minimizing travel cost. It is shown that solving the traditional PVRP to minimize travel cost can lead to solutions that are less desirable from the perspective of workforce management, even with post-processing. For example, the traditional PVRP generally reduces the number of routes used, which limits the available options when assigning drivers to customers. However, by adding workforce management metrics to the objective function, an appropriate balance can be obtained between travel cost and workforce management goals.

We explore approaches to model driver familiarity with both customers and regions, finding that workforce management principles can be applied with a minimal impact on travel costs. These results are found through a computational study with both standard PVRP test cases, modified for our problem setting, and a data set provided by a major package delivery carrier. Fundamentally, this paper presents transportation managers with best practice methods for applying workforce management without sacrificing other operational objectives.

As an extension to this work, one can incorporate the dynamic nature of driver learning more explicitly. Over time, as drivers become more familiar with their routes, it should be possible to add more customers or a larger service territory. Therefore, a next step should incorporate variable capacity constraints and variable demand to reflect the improved service level of drivers.

Acknowledgments

This research has been supported by grant DMI-0348622 from the National Science Foundation. The authors would like to thank our contacts at the major package delivery company for providing data and valuable insights on periodic routing in practice. The authors also thank the referees and associate editor for their helpful comments that improved the paper considerably.

References

- Bartlett, C.A., S. Ghoshal. 2002. Building competitive advantage through people. *MIT Sloan Management Review* **43**(2) 34.
- Beasley, J. E. 1984. Fixed routes. *Journal of the Operational Research Society* **35** 49–55.
- Beltrami, E., L. Bodin. 1974. Networks and vehicle routing for municipal waste collection. *Networks* **4** 65–94.
- Christofides, N. 1971. Fixed Routes and Areas for Delivery. *International Journal of Physical Distribution* **1** 87–92.
- Christofides, N., S. Eilon. 1969. An algorithm for the vehicle-dispatching problem. *Operations Research* **20**(3) 309–318.
- Cordeau, J-F, M. Gendreau, G. Laporte. 1997. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks* **30** 105–119.
- Francis, P., K. Smilowitz, M. Tzur. 2006. The period vehicle routing problem with service choice. *Transportation Science* **40** 439–454.
- Francis, P., K. Smilowitz, M. Tzur. 2007. Flexibility and complexity in periodic distribution problems. *Naval Research Logistics* **54** 136 – 150.
- Francis, P., K. Smilowitz, M. Tzur. 2008. The period vehicle routing problem and its extensions. *The Vehicle Routing Problem* 239–261.
- Gendreau, M., J. F. Cordeau, G. Laporte. 1997. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks* **30**(2) 105–119.
- Gendreau, M., A. Hertz, G. Laporte. 1992. New insertion and postoptimization procedures for the traveling salesman problem. *Operations Research* **40**(6) 1086–1094.
- Gillett, B.E., L.R. Miller. 1974. A heuristic algorithm for the vehicle-dispatch problem. *Operations Research* **22** 340–349.
- Groer, C., B. Golden, E. Wasil. 2009. The consistent vehicle routing problem. *Manufacturing and Service Operations Management* **11** 630–643.
- Hopp, W.J., E. Tekin, M.P. Van Oyen. 2004. Benefits of skill chaining in serial production lines with cross-trained workers. *Management Science* 83–98.

- Iravani, S.M.R., B. Kolfal, M.P. Van Oyen. 2007. Call-Center Labor Cross-Training: It's a Small World After All. *Management Science* **53**(7) 1102.
- Jennings, O.B., A. Mandelbaum, W.A. Massey, W. Whitt. 1996. Server staffing to meet time-varying demand. *Management Science* 1383–1394.
- Jozeffowicz, N., F. Semet, E.G. Talbi. 2007. Target aiming Pareto search and its application to the vehicle routing problem with route balancing. *Journal of Heuristics* **13**(5) 455–469.
- Levy, L., L. Bodin. 1988. *Vehicle routing: Methods and studies*, chap. Scheduling the postal carriers for the United States postal service: an application of arc partitioning and routing. North Holland: Elsevier Science Publishers BV, 359–394.
- Li, F., B. Golden, E. Wasil. 2005. Very large-scale vehicle routing: New test problems, algorithms, and results. *Computers and Operations Research* **32** 1197–1212.
- Partyka, J., R. Hall. 2010. On the road to connectivity. *OR/MS Today* **37**(1).
- Pfeffer, J. 1998. *The human equation: Building profits by putting people first*. Harvard Business School Press.
- Russell, R., W. Igo. 1979. An assignment routing problem. *Networks* **9** 1–17.
- Smilowitz, K., M. Nowak, T. Jiang. 2009. Workforce management in periodic delivery operations. IEMS Working Paper No. 09-04.
- UPS Corp. 2006. Driving success: why the UPS model for managing 85,000 drivers is a competitive advantage [Http://pressroom.ups.com/mediakits/popups/factsheet/0,1889,120100.html](http://pressroom.ups.com/mediakits/popups/factsheet/0,1889,120100.html).
- Whitt, W. 2005. Engineering solution of a basic call-center model. *Management Science* **51**(2) 221–235.
- Wong, K., J. E. Beasley. 1984. Vehicle routing using fixed delivery areas. *Omega International Journal of Management Science* **12**(6) 591–600.
- Zhong, H., R. W. Hall, M. Dessouky. 2007. Territory planning and vehicle dispatching with driver learning. *Transportation Science* **41** 74–89.

Appendix

Problem instance				ConVRP1	ConVRP2	ConVRP3	ConVRP4	ConVRP5	ConVRP11	ConVRP12
Number of customers				50	75	100	150	199	120	100
p	ϕ_{CF}	ϕ_{RF}	Model	Computational time (in seconds)						
0.6	-	-	PVRP(TD)	62	133	206	479	766	254	214
	25	-	PVRP(CF)	58	120	207	483	753	288	218
	-	50	PVRP(RF)	64	139	209	490	768	290	226
	250	-	PVRP(CF)	81	175	249	680	1108	380	276
	-	500	PVRP(RF)	80	176	254	698	1114	302	288
	500	-	PVRP(CF)	80	176	257	688	1337	398	291
	-	1000	PVRP(RF)	78	173	255	701	1331	396	288
	-	-	PVRP(TD)	81	178	251	612	1177	335	250
	25	-	PVRP(CF)	71	173	239	614	1162	336	255
	-	50	PVRP(RF)	74	187	252	635	1169	364	261
0.7	250	-	PVRP(CF)	92	238	308	854	1710	463	328
	-	500	PVRP(RF)	93	240	310	871	1730	370	339
	500	-	PVRP(CF)	94	245	324	879	1800	490	343
	-	1000	PVRP(RF)	90	237	311	870	1743	473	338
	-	-	PVRP(TD)	102	220	333	700	1341	423	335
	25	-	PVRP(CF)	100	213	313	690	1330	434	342
0.8	-	50	PVRP(RF)	98	234	335	701	1342	422	346
	250	-	PVRP(CF)	122	296	421	947	1930	546	436
	-	500	PVRP(RF)	122	301	435	952	1961	431	436
	500	-	PVRP(CF)	125	299	423	946	1925	540	434
	-	1000	PVRP(RF)	120	306	442	944	1956	548	433

Table 6 Comparison of the computational times (in seconds)