10-16-2020

Binary Driver-Customer Familiarity in Service Routing

Marlin Ulmer
Technische Universität Braunschweig

Maciek Nowak
Loyola University Chicago, mnowak4@luc.edu

Dirk Mattfeld
Technische Universität Braunschweig

Bogumił Kaminski
SGH Warsaw School of Economics

Author Manuscript

This is a pre-publication author manuscript of the final, published article.

Recommended Citation

Ulmer, Marlin; Nowak, Maciek; Mattfeld, Dirk; and Kaminski, Bogumił. Binary Driver-Customer Familiarity in Service Routing. European Journal of Operational Research, 286, 2: 477-493, 2020. Retrieved from Loyola eCommons, School of Business: Faculty Publications and Other Works, http://dx.doi.org/10.1016/j.ejor.2020.03.037

Binary Driver-Customer Familiarity in Service Routing

Marlin Ulmer[a,*], Maciek Nowak[b], Dirk Mattfeld[a], Bogumił Kaminski[c]

[a]Technische Universität Braunschweig
38106 Braunschweig, Germany

[b]Loyola University Chicago
Chicago, USA 60611

[c]SGH Warsaw School of Economics
02-554 Warsaw, Poland

*corresponding author
E-mail addresses: *m.ulmer@tu-braunschweig.de, mnowak4@luc.edu, d.mattfeld@tu-braunschweig.de, bogumil.kaminski@sgh.waw.pl*

Abstract

With a growing number of services provided at or to a customer's home, the familiarity of the service provider, or driver, with the customer's location is increasingly important. One prominent example is retail distribution, where familiarity with the delivery location can save the driver time. In contrast to other kinds of familiarity (e.g., tasks, customer needs) that continually increase with a larger number of repetitions, location familiarity is primarily established with a driver's first visit. Thus, familiarity results from operational routing decisions. However, as we show in this paper, there is potential value in considering the tactical value of familiarity and its development over a longer horizon. To this end, we develop and solve a tactical model to specify the long-term implications of improved driver-customer familiarity, introducing a solution methodology for the stochastic and dynamic multi-period routing problem with driver-customer familiarity. Our methodology utilizes a policy that explicitly invests in the familiarity between selected driver-customer pairs, encouraging the development of pairs that are tactically beneficial. We determine the appropriate investment dimensions for each pair, considering which locations a driver has visited and how many drivers have visited a location. We show that under the problem conditions tested this investment policy leads to a reduction in cost compared to a short term, myopic policy, while increasing the overall level of familiarity between drivers and customers and hedging against driver or customer turnover. We also find that focusing only on routing or on exploiting existing familiarity leads to substantial increases in cost.

Keywords: routing, stochasticity, dynamism, familiarity, cost function approximation

---

1. Introduction

With the growth in services provided at a customer's location, including retail distribution, home health care, repair service, and parcel delivery, the relationship between the driver delivering the service and the customer has grown in importance. Logistics service providers are increasingly focused on the familiarity that a driver has with the base of customers regularly visited, as this has a direct impact on service related time and customer retention. Familiarity in home health care has been shown to decrease the probability of hospitalization and improve functioning in daily activities (Russell et al. 2011), while the plentiful stories of friendships between UPS drivers and their customers underline the importance placed on these relationships (Patel and Wright 2018).

Familiarity can be gained with the customers and their needs, the service tasks to fulfill, and the location at which the service takes place. Developing familiarity with a customer's needs is important in home-healthcare, where a caregiver's knowledge regarding a patient can improve service quality. Task familiarity is important in repair services, where the same tasks are conducted at different customer locations. Developed knowledge of tasks can accelerate future service times regardless of the customers requesting them (Chen et al. 2016). The focus of this paper is familiarity with the location of the customer, such as an understanding of the road network, where to find parking, where a dock is located, etc. While familiarity with a customer's needs or the required tasks is established gradually with a larger number of repetitions, location familiarity is rather binary as it is primarily established with a driver's first visit (Zhong et al. 2007). For example, Juan Perez, UPS's chief information and engineering officer, states, "if you plug in the address to the New York Stock Exchange [in a commercial mapping app], that tool will just drop you off at the front of the building, but in the case of UPS, that's not useful for us. We need to know where drivers need to go to make pickups and deliveries" (Premack 2019). While a driver may spend considerable time on the first visit, the time required for a second visit should be considerably shorter.

While familiarity is important in practice, drivers have traditionally had customers assigned to them using methodologies that minimize operating cost, with driver-customer familiarity only an afterthought. Familiarity may develop organically, as drivers visit new customers when the routing solution dictates. However, it is advantageous for service providers to integrate familiarity in their long-term planning approaches. At a tactical level, providers should establish and maintain sufficient driver familiarity with consideration for longer term operational cost (which for the purposes of this paper includes routing and service costs), driver turnover and changes in the customer base. This may be achieved by investing in additional driver-customer familiarity early in the planning horizon, allowing for greater flexibility and more robust planning options. An investment in familiarity is made when a driver visits an unfamiliar customer following an extended path that results in an increased routing cost. While initially costly, intentionally introducing a customer to a new driver provides routing flexibility in the future, reducing long term routing and service costs. However, selecting an appropriate investment is not trivial as the return on this investment is not immediately realized.

We propose that the value of familiarity between driver and customer primarily depends on the *investment magnitude*, the *existing familiarity* that a customer has with the set of drivers and the *geographic proximity* of unfamiliar customers to a driver's known customers. A large investment in familiarity between many drivers and customers leads to considerable flexibility in routing later in the horizon, but that investment may be too large to recover over the long term. Alternatively, a small investment may be insufficient to promote familiarity, resulting in a solution similar to one without any investment. Also, if a customer has existing familiarity with many drivers or if an unfamiliar customer is geographically distant from customers familiar to the driver, an investment may not be of benefit. In this paper, we show that it is vital to incorporate all three of these dimensions when developing an effective investment strategy.

To analyze the impact of tactically developing familiarity, we introduce a stochastic, dynamic and multi-period vehicle routing problem with binary driver-customer familiarity (VRPBDCF). The problem comprises an extended number of periods (days), with a horizon that reaches several months. The problem starts with a fleet of unfamiliar drivers and a pool of potential customers. Each driver has a maximum number of working hours per day. Each day a subset of customers requesting service is revealed and each customer's service time differs depending on familiarity with the assigned driver. The dispatcher determines a corresponding routing solution based on these requests. When a driver serves an unfamiliar customer, the driver gains familiarity with this customer for the remaining periods. The operational objective is to minimize travel and service times across all drivers over all periods.

Solving this stochastic and dynamic problem is challenging. We face a large number of periods and extremely large state, decision, and information spaces. Thus, we draw on heuristics. Our heuristic is based on the idea of targeted "investments" in driver-customer familiarity. We embed the investment strategy in a cost function approximation (CFA), a method of approximate dynamic programming (Powell 2011). A CFA manipulates the cost function of the problem to account for potential future impacts of a decision. In our case, the CFA incorporates the value of familiarities between drivers and customers as a deduction from a daily routing decision's cost.

To understand the role of investments for driver-customer familiarity, we apply our heuristic to a large-scale, retail distribution network in which drivers move goods from a central warehouse (or depot) to retail stores (or customers) distributed around the warehouse. This analysis provides insights into the interplay of investment magnitude, existing familiarity and geographic proximity by revealing the influence that each dimension has on the others. It also indicates the importance of an investment policy in the context of this problem. Our policy leads to a reduction in cost and significant increase in familiarity compared to a myopic policy that only considers routing and service costs in the short term. We find that the benefits of investment change as these dimensions are modified. We show that well targeted investments amortize relatively quickly after a few weeks while poorly targeted investments increase cost significantly with no added benefit. Our study also reveals that ignoring familiarity and solely focusing on routing cost in decision making can result in a considerable increase in total cost. Similarly, prioritizing familiarity over efficient routing leads to high cost in the long run. Finally, we establish that familiarity may be particularly valuable to hedge against driver or customer turnover.

We contribute to the field of research by being the first to develop anticipatory investment policies for a vehicle routing problem with familiarity between drivers and customers. We introduce the VRPBDCF and present a Markov decision process model of the problem. We introduce and motivate three dimensions to consider when making investment decisions. We integrate the dimensions in a cost function approximation. Our computational study is a proof of concept to confirm and analyze the importance of the three dimensions, providing insight into how a manager may take these various dimensions into consideration when solving her specific problem.

The paper is outlined as follows. We present related literature on service routing and customer familiarity in §2. In §3, we define the problem and model it as a Markov decision process. We present our CFA in §4. In §5, we conduct our experimental study. The paper ends with a conclusion and an outlook in §6.

2. Literature

In the following, we present the relevant literature. We first present the work most relevant to ours. We then discuss work on consistency routing and workforce training and learning. Finally, we present a summarizing table.

The most relevant work is presented by Zhong et al. (2007), Chen et al. (2016), and Chen et al. (2017). Zhong et al. (2007) were the first to incorporate driver learning into a routing model. Their two phase algorithm first strategically clusters stops into core areas, with a core area served by the same driver each period. The second, operational phase of the algorithm connected the stops within these core areas along with unclassified stops into daily routes. The stops that are not in core areas may be visited by a different driver each period; however, as a driver learns the neighborhood around a stop, she services that neighborhood more quickly. Our work differs from Zhong et al. (2007) for the following reasons. First, Zhong et al. (2007) consider a different type of problem with many different customers. Instead of driver-customer familiarity, they focus on territory familiarity and strategic assignment of core territorial *areas* to drivers instead of specific customer locations. Second, the assignment is a constraint in the operational problem where routing cost is minimized myopically every day over 30 days. Thus, the dynamic development of familiarity is not considered in their method and drivers are not trained explicitly.

Chen et al. (2016) and Chen et al. (2017) incorporate learning when considering the assignment of technicians serving numerous customers. Over a sequence of days, each day, a set of tasks need to be performed by a technician workforce. Based on the assignment decision in one day, the experience level of the workforce changes and tasks can be performed faster than before. In Chen et al. (2016), the technicians are routed myopically every day based on the most cost-efficient routing. Future developments are ignored. In Chen et al. (2017), no routing is considered. Similar to our work, the authors present anticipatory methods that incorporate the impact of a current assignment decision on the future cost.

Our work differs from Chen et al. (2016) and Chen et al. (2017) for several reasons. First, both papers focus on technician-task familiarity but not the familiarity of drivers and customers. So, a driver may be unfamiliar with a previously visited customer, if the customer requests a different task, and, a driver may be familiar with a task at a customer the driver never visited before. Thus, the routing becomes more a constraint for the assignment problem than a part of strategic considerations. In contrast to Chen et al. (2017), we consider familiarity in the context of a routing problem. In particular, we show that suitable investments in familiarity depend on the familiarity of the driver to other customers in the same neighborhood. In contrast to Chen et al. (2016), we make anticipatory routing decisions. We show that anticipation leads to a reduction in cost and a significant increase in familiarity.

2.1. Consistency Routing

Familiarity is closely related to consistency of service between customers and drivers. Consistency implies that a customer should be served by the same driver for each visit, such as in cases where a driver needs a key or security clearance. To address consistency, Groër et al. (2009) introduce the consistent vehicle routing problem (ConVRP). The ConVRP is an extension of the period vehicle routing problem in which the daily routes are constructed such that a customer is visited by the same driver each day she requests service. Consistency of service has been studied in a variety of settings (Kovacs et al. 2014), including small package shipping, health care, and inventory replenishment. Of note is that the majority of these papers constrain the requirement for consistency, such that every customer must be served by the same driver across all visits. While utilizing constraints guarantees consistency, it likely comes at an increased routing cost. For example, Smilowitz et al. (2013) quantify the tradeoff between minimizing routing cost and maximizing consistency of service, finding that a focus on driver consistency leads to an increase in distance of up to 5.2 percent.

Some research enforces consistency by including it in the objective instead of the constraints. Smilowitz et al. (2013), Kovacs et al. (2015b) and Lian et al. (2016) all consider multiple objectives with a particular focus on travel time and driver consistency, but also include region (Smilowitz et al. 2013) and arrival time consistency (Kovacs et al. 2015b). Sungur et al. (2010) use a multi-objective analysis to consider the time required to complete all routes and the similarity of the routes from day to day, improving upon both metrics relative to a real world solution methodology. Focusing on the inventory routing problem, Coelho et al. (2012) relax the requirement that all customers are visited by one driver with a penalty term in the objective function, proportional to the number of extra vehicles assigned to each customer. This is used to evaluate the amount by which the solution cost increases when the penalty is incorporated. Hewitt et al. (2016) present a simple heuristic that anticipates future demand by incorporating dummy service requests located at the vehicle depot. As demand is realized, the dummy customer is replaced by the location of the actual customer and routed accordingly.

Several papers test the impact of consistency on routing cost by parameterizing the number of drivers and adjusting this value for various scenarios (Kovacs et al. 2015a, Luo et al. 2015, Braekers and Kovacs 2016). Battini et al. (2015) use a simulation to compare scenarios in which a fixed set of routes is created and used on a daily basis, thereby serving each customer with the same driver, with routes that are modified on a daily basis depending on customer demand. Consistency in the Traveling Salesman Problem (TSP) is evaluated by Subramanyam and Gounaris (2016) who compare the cost of optimal solutions for the TSPLIB instances with and without consistency.

Our paper extends the understanding of the relationship between the benefits and costs of driver-customer familiarity, revealing how incurring up front costs can lead to greater long term benefits. We note that our problem is stochastic and dynamic, such that methodologies developed for the ConVRP cannot be applied.

2.2. *Workforce Learning and Training*

Several papers address the impact of workforces gaining experience, either by learning or training. The experience that a driver gains by developing familiarity with a customer base can be classified as a hierarchical skill (De Bruecker et al. 2015), such that a driver with more experience can perform more tasks, or perform certain tasks faster or better. While considerable research has focused on scheduling a workforce based on the skills that each worker possesses (De Bruecker et al. 2015), more limited is work that considers how those skills transform over time through training or learning. Training is a planned activity that occurs on the job or externally, with time or resources allocated for a worker to be trained. Learning is something that happens automatically as the worker completes her standard tasks. Our research incorporates aspects of both training and learning literature. Training has benefits but comes at a cost, whether that involves dedicating worker time or paying a worker to learn a new skill, or decreased productivity by assigning a worker to a task so that he or she may gain experience. Learning traditionally does not have an explicit cost associated and happens as a part of a worker's standard daily tasks.

Some attention has been given to driver learning, as seen in Zhong et al. (2007), Chen et al. (2016) and Chen et al. (2017). However, we are not aware of any research that has considered driver training as it relates to routing or operational improvement. There is some literature on training in a general service setting (Fowler et al. 2008, Li and Li 2000, Marentette et al. 2009). Valeva et al. (2017a) and Valeva et al. (2017b) incorporate practice time into their models. When practicing, a worker may be assigned a task in order to learn, but this practice time results in no productive output.

2.3. Summary

Table 2 provides a summary of literature that is relevant to this paper. We differentiate work with respect to the following categories: Period length, which is the planning horizon the paper considers; Stochasticity, whether the paper considers uncertainty in at least one problem component, usually, the customers to serve; Anticipation, whether the proposed method anticipates the impact of a decision on future costs; Routing, whether the paper addresses the routing of vehicles to serve customers; Familiarity: whether the paper considers familiarity ("Constrained" indicates that the customer must be served by the same driver/worker on each visit, "Objective" indicates that familiarity is incentivized by being incorporated in the objective value); Learning/Training, whether the paper considers learning or training ("Learning" indicates that worker performance improves with experience, "Training" indicates that a cost is incurred for workers to gain experience); and Workforce, whether the problem considers a heterogeneous workforce, such as with different skill levels for technicians.

3. Problem Definition

In this section, we present the problem statement. We then model the VRPBDCF as a Markov decision process (MDP, Puterman 2014). Finally, we give an example for the MDP.

3.1. Problem Statement

Over a sequence of days, a provider serves customers with a fleet of drivers. Regulatory guidelines limit the hours available in a driver's working day. Each customer requests service intermittently, with these requests known only at the beginning of each day. Thus, each day, a subset of the overall set of customers requires service. We denote these customers as "active" customers. The dispatcher then determines a set of routes for the drivers to serve the active customers. Cost is determined by the amount of time the drivers

Table 1: Literature Classification

Authors	Year	Period length	Stochastic	Anticipatory	Routing	Familiarity	Learning/Training	Workforce
Zhong et al.	2007	30 days	X		X	Objective	Learning	Homogeneous
Haughton	2008	300 days			X	Constrained		Homogeneous
Groer et al.	2009	5 days			X	Constrained		Homogeneous
Sungur et al.	2010	42 days	X	X	X	Objective		Homogeneous
Coelho et al.	2012	6 days			X	Constrained		Homogeneous
Tarantilis et al.	2012	5 days			X	Constrained		Homogeneous
Smilowitz et al.	2013	5 days			X	Objective		Homogeneous
Feillet et al.	2014	5 days			X	Objective*		Homogeneous
Kovacs et al.	2014	5 days			X	Constrained		Homogeneous
Battini et al.	2015	30 days	X		X	Constrained	Learning	Homogeneous
Kovacs et al.	2015a	5 days			X	Constrained		Homogeneous
Kovacs et al.	2015b	5 days			X	Objective		Homogeneous
Luo et al.	2015	7 days			X	Constrained		Homogeneous
Braekers and Kovacs	2016	5 days			X	Constrained		Homogeneous
Chen et al.	2016	29 days			X	Objective	Learning	Heterogeneous
Hewitt et al.	2016	94 days	X	X	X	Constrained		Homogeneous
Lian et al.	2016	5 days			X	Objective		Homogeneous
Subramanyam and Gounaris	2016	5 days			X	Constrained		Homogeneous
Chen et al.	2017	120 days	X	X		Objective	Learning/Training	Heterogeneous
Subramanyam and Gounaris	2017	5 days			X	Constrained		Homogeneous
Valeva et al.	2017a	10 days	X			None	Learning	Heterogeneous
Valeva et al.	2017b	20 days	X			None	Learning	Heterogeneous
Campelo et al.	2018	5 days			X	Constrained		Homogeneous
Xu and Cai	2018	5 days			X	Constrained		Homogeneous
Our Work	2019	120 days	X	X	X	Objective	Learning/Training	Heterogeneous

spend serving all customers over the course of the day. This time is spent either driving from customer to customer or serving the customer.

The familiarity between driver and customer determines the service time required for each visit, with this time decreasing as familiarity increases. Familiarity is assumed to be binary. If a customer is served by an unfamiliar driver, this driver then gains familiarity with that customer for future visits. All drivers and customers are initially unfamiliar with each other, gaining familiarity when the driver first visits the customer (the model and methodology can be easily extended to instance settings where familiarity is already partially established). This is supported by well established research on learning curves in manufacturing, technician routing and vehicle dispatching (Wright 1936, Zhong et al. 2007, Chen et al. 2016), finding that the largest jump in familiarity is after an individual's first exposure to an activity. The most information is generally gained on the first visit to a customer, such as where to park, which dock door to use, what tools are needed, etc.

The service provider has two goals in mind. They primarily want to serve the customers cost-efficiently. Thus, the objective of this problem is to minimize the working time of all drivers across all periods. However, the provider would also like to gain familiarity between drivers and customers to allow efficient and high quality service. This also allows the provider to prepare for future changes in driver availability, such as with driver turnover. While the objective is to minimize cost, we analyze how our decisions impact familiarity in our computational study.

3.2. Markov Decision Process

The problem at hand is a stochastic and dynamic decision problem. It is stochastic because the set of active customers only becomes known at the beginning of each day. It is dynamic because a sequence of decisions need to be made, one decision per day. Furthermore, current decisions change driver familiarity and therefore impact potential future outcomes.

A stochastic and dynamic decision problem can be modeled as a Markov decision process (MDP). An MDP models the problem as a sequence of decision states. In each state, a decision is made and costs are observed, followed by a stochastic transition that leads to the next decision state. In the following, we define decision states, decisions, costs, and transitions for our problem. But first, we introduce some general notation for the problem.

General Notation

We assume a set of K drivers serve customers from a set N over a sequence of periods (days) $1, \ldots, P$. The drivers start at the depot 0. The depot and customers form the set N^0. The service area is represented by a complete graph $G = (N^0, A)$, with the set of directed arcs A connecting nodes in N^0. Each arc $(i, j) \in A$ is associated with a non-negative travel cost in minutes c_{ij}. If a driver is familiar with a customer, the service at this customer requires t^{familiar} minutes. Otherwise, the service requires $t^{\text{unfamiliar}}$ minutes. The maximum working time per driver and day is denoted by T.

Decision State

The decision state comprises all information necessary to make a decision. We denote the decision state in period p as S_p. For our problem, the decision state S_p consists of two components: the set of active customers to serve in the period and the current familiarity between drivers and customers. The set of active customers is denoted by $N_p \subseteq N$. The set including the depot is denoted $N_p^0 = N_p \cup \{0\}$. Each customer has a probability of ν to request service in a period p. The expected number of customers to request service in a given period is therefore $\nu \times |N|$.

The familiarity between a customer i and a driver k in period p is represented by the binary variable $f_{ikp} \in \{0, 1\}$. If driver k is familiar with customer i in period p, the variable is 1. Else, the variable is 0. The familiarity in the beginning of period p is summarized in matrix $F_p = (f_{ikp})_{i=1,\ldots,N,k=1,\ldots,K}$.

Based on the familiarity matrix F_p, we can also generate a service time matrix t_p with

$$t_{ikp} = \begin{cases} t^{\text{familiar}} & \text{if } f_{ikp} = 1 \\ t^{\text{unfamiliar}} & \text{otherwise.} \end{cases}$$

We set the service time at the depot to 0, such that $t_{0kp} = 0$ for all k and p. A state can be summarized as $S_p = (N_p, F_p)$.

Decision

A decision a_p in period p is made regarding the assignment of active customers to drivers as well as the routing of drivers. Thus, every decision results in the solution of a vehicle routing problem. In the following, we present a mixed-integer formulation of this routing problem.

The decision a_p determines the assignment of customers to drivers and the arc routing variables. The assignment is denoted by variable y_p and the routing by variable x_p as follows:

$$y_{ikp} = \begin{cases} 1 & \text{if customer } i \in N \text{ is visited by driver } k \in K \text{ in period } p \\ 0 & \text{otherwise} \end{cases}$$

$$x_{ijkp} = \begin{cases} 1 & \text{if driver } k \in K \text{ traverses arc } (i,j) \in A \text{ in period } p \\ 0 & \text{otherwise} \end{cases}$$

A decision $a_p = (y_p, x_p)$ is feasible, if the following constraints hold:

$$\sum_{k \in K} y_{ikp} = 1 \qquad \forall i \in N_p \tag{1a}$$

$$\sum_{i,j \in N_p^0} (c_{ij} + t_{ikp}) x_{ijkp} \leq T \qquad \forall k \in K \tag{1b}$$

$$\sum_{j \in N_p^0} x_{ijkp} = y_{ikp} \qquad \forall i \in N_p; k \in K \tag{1c}$$

$$\sum_{j \in N_p^0} x_{ijkp} = \sum_{j \in N_p^0} x_{jikp} \qquad \forall i \in N_p^0; k \in K \tag{1d}$$

$$\sum_{i,j \in Q} x_{ijkp} \leq |Q| - 1 \qquad \forall Q \subseteq N_p; k \in K \tag{1e}$$

$$y_{ikp} \in \{0,1\} \qquad \forall i \in N_p; k \in K \tag{1f}$$

$$x_{ijkp} \in \{0,1\} \qquad \forall i,j \in N_p^0; k \in K \tag{1g}$$

Constraints (1a) ensure that a customer $i \in N_p$ is visited by a driver. Constraints (1b) ensure that the time limit holds for each driver. We note that a tight constraint can theoretically lead to no feasible solutions. Constraints (1c) link the x_p and y_p variables for the customers in set N_p. Constraints (1d) ensure flow conservation at each node $i \in N_p^0$. Constraints (1e) are the subtour elimination constraints which also guarantee that each route contains a stop at the depot. Finally, constraints (1f) and (1g) define the binary variables for assignment and routing.

The costs of a state and a decision are the sum of routing and service costs:

$$C(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} (c_{ij} + t_{ikp}) x_{ijkp}. \tag{2}$$

Post-Decision State

A combination of a state S_p and a decision a_p leads to a known post-decision state S_p^a. The post decision state only contains the updated familiarity matrix F_p^a with $f_{ikp}^a = \max\{f_{ikp}, y_{ikp}\}$.

Exogenous Information and Transition

The exogenous information is observed between post-decision state S_p^a and new state S_{p+1}. We denote a realization as ω_{p+1}. The exogenous information provides the set of new customers N_{p+1}^ω. Based on ω_{p+1}, the transition function $W(S_p^a, \omega_{p+1})$ leads to a new state $S_{p+1} = (N_{p+1}^\omega, F_p^a)$.

Solution: Decision Policy

A solution for a Markov decision process is a decision policy π. A policy π assigns a decision $A^\pi(S_p)$ to every state S_p. The overall set of policies is defined as Π. An optimal solution $\pi^* \in \Pi$ minimizes the expected cost:

$$\pi^* = \arg\max_{\pi \in \Pi} \mathbb{E}\left[\sum_{p=1}^{P} (C(S_p, A^\pi(S_p))|S_0)\right]. \tag{3}$$

The optimal policy minimizes the expected costs over all periods starting with state S_0.

3.3. Example

In the following, we give a small example to illustrate the components of the Markov decision process and to motivate our solution methodology. Figure 1 presents a state S. For the sake of simplicity, we omit index p in this formulation. The depot is indicated by the black square. Customers are indicated by circles. In the example, we have an overall set of 7 customers. Four of the customers are active (1, 2, 3, and 4).

Familiarity is represented by the shading of a customer. Customers with vertical stripes (2, 4 and 5) are familiar only with driver 1. Customers with horizontal stripes (1, 3 and 6) are only familiar with driver 2. Customers with both vertical and horizontal stripes (7) are familiar with both drivers. In the example, the service time for a familiar customer is 20 minutes. The service time for an unfamiliar customer is 40 minutes.

For the example, the maximum working time per driver is set to 200 minutes. For the purpose of presentation, we assume drivers travel on a Manhattan grid in this example. The travel time of each segment is 6 minutes.

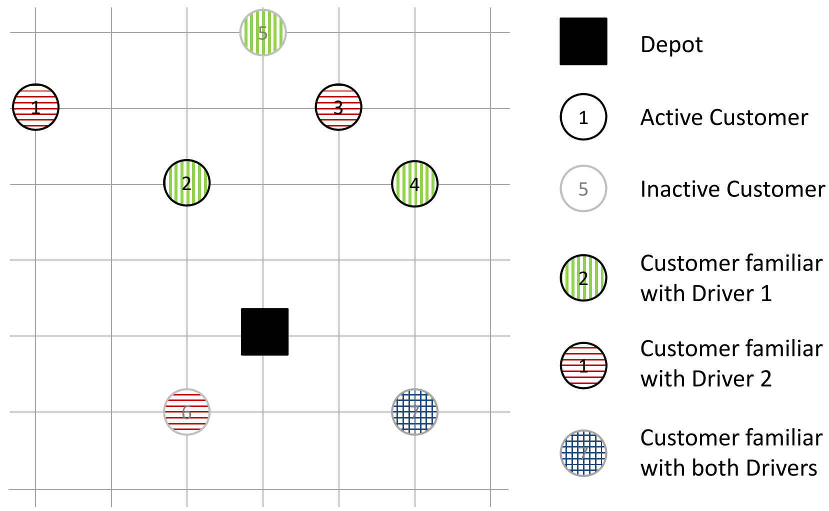

Figure 1: Example for a State S

The state can be summarized as $S = (\{1, 2, 3, 4\}, F)$ with

$$F = \begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

Based on state S, different decisions a can be made. In the following, we present two potential decisions a^1 and a^2 in Figure 2a and b, respectively. Each decision determines a routing for driver 1 and driver 2, represented as a dashed line and dotted line, respectively. In decision a^1, driver 1 serves customers 2 and 4 and driver 2 serves customers 1 and 3. In decision a^2, driver 1 serves customers 3 and 4 and driver 2 serves customers 1 and 2.

Decision a^1 leads to the smallest cost in state S with:

$$C(S, a^1) = \underbrace{10 \times 6 + 2 \times 20}_{\text{Driver 1}} + \underbrace{14 \times 6 + 2 \times 20}_{\text{Driver 2}} = 100 + 124 = 224$$

minutes. Decision a^2 leads to higher cost with:

$$C(S, a^2) = \underbrace{10 \times 6 + 20 + 40}_{\text{Driver 1}} + \underbrace{12 \times 6 + 20 + 40}_{\text{Driver 2}} = 120 + 132 = 252$$

minutes, or 28 additional minutes.

However, the decisions lead to different familiarity in the post-decision states. The familiarity in post-decision state S^{a^1} is the same as in the previous state: $F^{a^1} = F$. The familiarity in post-decision state S^{a^2}

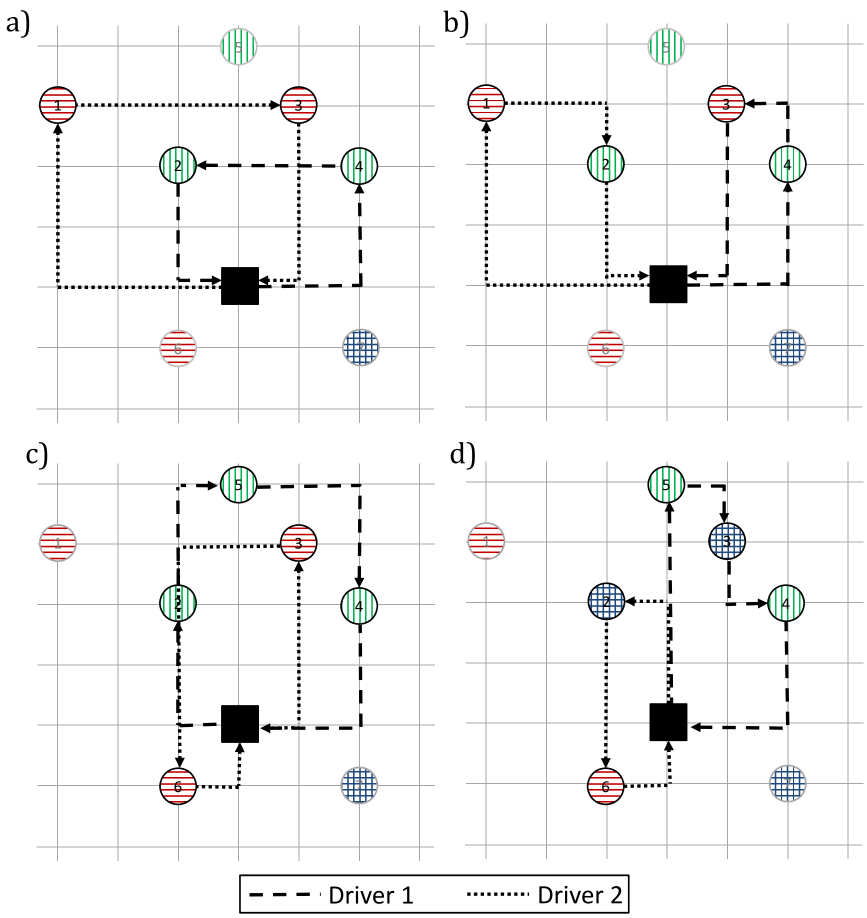

Figure 2: Example of the two potential decisions a) a^1 and b) a^2; and the next states and minimum cost solution for ω resulting from decisions c) a^1 and d) a^2.

is different with

$$F^{a^2} = \begin{pmatrix} 0 & 1 & \mathbf{1} & 1 & 1 & 0 & 1 \\ 1 & \mathbf{1} & 1 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

The entry for driver 2 and customer 2 changed from unfamiliar to familiar, as did the entry for driver 1 and customer 3 (both bolded).

The two different decisions lead to two different new states for the same stochastic realization. In the example, we assume that the realization is $\omega = \{2, 3, 4, 5, 6\}$. Thus, customers 2, 3, 4, 5, and 6 need to be served in the next period. The next states and minimum cost solution for ω resulting from decisions a^1 and a^2 are shown in Figure 2c and d, respectively. The solution for post-decision state S^{a^1} results in a minimum cost of:

$$C(S^{a^1}, a^3) = \underbrace{14 \times 6 + 3 \times 20}_{\text{Driver 1}} + \underbrace{12 \times 6 + 2 \times 20}_{\text{Driver 2}} = 144 + 112 = 256$$

minutes. Post-decision state S^{a^2} has a minimum cost of:

$$C(S^{a^2}, a^4) = \underbrace{12 \times 6 + 3 \times 20}_{\text{Driver 1}} + \underbrace{8 \times 6 + 2 \times 20}_{\text{Driver 2}} = 132 + 88 = 220$$

minutes, with a savings of 36 minutes (if the routing for post-decision state S^{a^2} is used for post-decision state S^{a^1}, the service time cost increases by 40 while the routing cost only decreases by 36). Thus, after only one period the investment, or additional time of 28 minutes, from decision a^2 has been recovered, with an additional savings of $36 - 28 = 8$ minutes. The flexibility afforded by this decision has allowed for both drivers to visit a set of familiar customers at a lower routing cost. Note that in the solution relating to post-decision state S^{a^1} driver 2 bypasses unfamiliar customer 2 as it would result in a greater cost than having driver 1 visit this customer. Therefore, no new familiarity has been gained and the lack of investment in familiarity is likely to incur additional costs in later periods. Although decision a^2 has resulted in a higher cost in the first period, it has reduced the cost in the following state. We use this premise in the generation of our solution method.

4. Solution Method

In this section, we present our solution method. We first describe the motivational considerations for investment in familiarity. We then present how those considerations are incorporated into our cost function approximation. We revisit the example to highlight the impact of our method. Finally, because our computational evaluation requires a fast solution of the VRP for each state visited, we present a runtime-efficient heuristic for the VRP with deadlines and heterogeneous service times.

4.1. Motivational Considerations

Solving the MDP with dynamic programming is very challenging because of the Curse of Dimensionality. The state comprises the period, the set of customers to serve, and the familiarity between customers and drivers. Assuming N potential customers, K drivers, and P periods, the number of potential states is the product of possible familiarity matrices, the subset of customer requests, and the number of periods: $2^{K \times N} \times 2^N \times P$. In our computational study, we test instances with 10 drivers, 120 potential customers, and 120 periods leading to a prohibitively high number of states. The decision space consists of all vehicle routing solutions. Searching the decision space is by itself an NP-hard problem. The information space contains all subsets of the set N leading to 2^N potential transitions per state and decision.

Thus, we draw on approximate dynamic programming, namely, cost function approximation (CFA). The most cost-efficient decision in a state is myopic because it ignores the impact of the decision to future states and costs. A CFA manipulates the cost function in a way that future costs are captured in the selection based on the manipulated cost function. One major advantage of CFAs is that the same conventional routing methodology can be applied without any additional simulations. Instead, future costs are approximated conceptually by identifying important dimensions of the state space. Based on the selected state space parameters, the cost function value of a decision is modified, often artificially increased or decreased. As Powell (2017) notes, CFAs are an "important class of policy that has been overlooked by the research community." CFAs are particularly suited to large-scale optimization problems such as the one presented here because they are able to "easily handle very high-dimensional data."

Our CFA is motivated by the concept of "investing" time to increase familiarity, enabling more efficient routing in the future. We consider an investment to be an intentional and state dependent deduction from the cost function that makes what would otherwise be a cost prohibitive move attractive. If a customer and driver are not familiar with each other, we introduce an artificial reduction in cost that makes attractive sending the driver to that customer, a move that might otherwise increase total cost in comparison to a myopic solution. We denote this reduction in state S_p for driver k to visit customer i in period p as I_{ikp}. Our manipulated cost function can be defined as follows:

$$C_I(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} (c_{ij} + t_{ikp} - I_{ikp}) x_{ijkp}.$$

The amount I_{ikp} indicates the value given to an investment in familiarity between driver k and customer i in period p. The CFA needs to determine such a value for every driver-customer pair in every state.

Investing in the "right" familiarity is challenging. Introducing familiarity to every potential driver-customer pair may yield a gain that is small relative to a high cost. Alternatively, sporadic investments will not provide sufficient benefit to justify the cost. Thus, a targeted investment is needed. In the following, we motivate and present the three dimensions that will be used to identify the "right" investments:

- **Investment Magnitude:** An investment in familiarity should not have such a significant cost that we may not be able to recover this investment in future states, while there is a sufficient investment to provide a tangible return. The investment magnitude is the baseline time that is invested in the familiarity between any driver-customer pair.

- **Existing Familiarity:** If a customer is known by only one or two drivers, it may be worthwhile to invest in additional familiarity to gain routing flexibility for future states. However, if many drivers already know a customer, the benefits are likely limited. Thus, the investment decisions should depend on the number of drivers already familiar with a customer.

- **Geographic Proximity:** Familiarity with a customer enables efficient routing of adjacent customers. Thus, investments should consider the proximity of a customer to other customers already known to the driver. This is determined by limiting the distance that a driver may travel from a familiar customer to an unfamiliar one.

In our computational study, we will show that all three dimensions are required for effective investments in driver familiarity.

4.2. Algorithmic Implementation

These dimensions are used to determine the values for I_{ikp}, calculated as follows:

$$I_{ikp} = (1 - f_{ikp}) \times \mathcal{M} \times \left(\frac{1}{|F(i)| + 1} \right)^q \times n_{ikp} \tag{4}$$

where:

- $\mathcal{M}$ is the investment magnitude, or the baseline time that is deducted when a driver visits an unfamiliar customer;

- q is the discount factor associated with the existing familiarity of a customer, or the number of drivers familiar with customer i, $|F(i)|$, $q \geq 0$;

- n_{ikp} is the variable indicating if customer i is in geographic proximity of another customer familiar with driver k, equal to 1 if so and 0 otherwise. Proximity is determined by the time required for driver k to reach customer i when traveling from the nearest customer that the driver is already familiar with. This proximity is limited using a maximum time parameter, t_{max}.

Increasing the investment magnitude leads to increasing values of I_{ikp}, increasing the existing familiarity of a customer decreases I_{ikp} because of the discount factor, and increasing the geographic proximity with the parameter t_{max} increases the number of driver-customer pairs that may be considered for investment. We note that as the problem progresses and p increases, customers become familiar with more drivers, increasing $|F(i)|$ and decreasing I_{ikp}. Because of the existing familiarity factor, the investments diminish once a certain saturation in familiarity is achieved.

In summary, the cost function C_I becomes:

$$C_I(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} \left(c_{ij} + t_{ikp} - \left((1 - f_{ikp}) \times \mathcal{M} \times \left(\frac{1}{|F(i)| + 1} \right)^q \times n_{ikp} \right) \right) x_{ijkp}.$$

We now revisit the example of Section 3 to show the impact of the investment dimensions. We draw on the state in Figure 1. The familiarity matrix in this state is

$$F = \begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

As an investment in relationships where familiarity has already been established does not add benefit, only those driver-customer pairs where $f_{ik} = 0$ are eligible for investment. We chose parameters $t_{max} = 12$, $q = 1.0$, and keep $\mathcal{M}$ abstract. As transit time on any segment is six minutes, parameter $t_{max} = 12$ leads to a neighborhood of two street segments. The only customers that are two segments or fewer from another customer are 3, 4, and 5, such that only $n_{3k}, n_{4k}, n_{5k} = 1$, for both drivers. Parameter $q = 1.0$ leads to a linear decrease in investment with increasing number of drivers knowing a customer. However, because the example only considers two drivers and every customer is known by one driver at this stage, this leads to a discount of $\left(\frac{1}{|F(i)|+1} \right)^q = \frac{1}{2}$ for customers eligible for investment.

In summary, we obtain the following investment matrix:

$$I = \begin{pmatrix} 0 & 0 & \frac{\mathcal{M}}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{\mathcal{M}}{2} & \frac{\mathcal{M}}{2} & 0 & 0 \end{pmatrix}.$$

Taking into consideration the familiarity matrix, we may construct the $t - I$ component of the cost function to understand how this investment impacts decision making. Recall that $t^{familiar} = 20$ and $t^{unfamiliar} = 40$. The matrix $t - I$ becomes:

$$t - I = \begin{pmatrix} 40 & 20 & 40 - \frac{\mathcal{M}}{2} & 20 & 20 & 40 & 20 \\ 20 & 40 & 20 & 40 - \frac{\mathcal{M}}{2} & 40 - \frac{\mathcal{M}}{2} & 20 & 20 \end{pmatrix}.$$

Recall the two potential decisions depicted in Figure 2. In the original example, decision a^1 leads to a cost of 224 and no familiarity gain. Decision a^2 leads to a cost of 252 and familiarity gains between driver

1 and customer 3 as well as driver 2 and customer 2. Our method reduces the cost of visiting customer 3 with driver 1 by $\frac{\mathcal{M}}{2}$ (with no impact on the driver 2 − customer 2 pair as customer 2 is not within t_{max} of another customer). A positive value of $\mathcal{M}$ reduces the element of $t - I$ associated with this pairing, thereby reducing the overall cost. For $\mathcal{M} > 56$, the investment reduces the cost of visiting customer 3 with driver 1 by more than 28 minutes and decision a^2 results in a lower cost than a^1. For $\mathcal{M} = 56$, both decisions are of equal value. For $\mathcal{M} < 56$, decision a^1 has a lower cost because the time to visit customer 3 with driver 1 is not sufficiently reduced. An alternative decision may swap the routes of the drivers from those shown in a^2, such that driver 2 gains familiarity with customer 4 and the investment similarly reduces cost, resulting in the same total cost.

4.3. Routing Heuristic

For our experiments, a routing heuristic is required to determine a decision a_p in every state S_p. This routing heuristic needs to account for heterogeneity in the customer's service time. Because our experiments require applying the heuristic millions of times, the heuristic also needs to be very runtime efficient.

We propose a construction heuristic and a neighborhood search. Conventional VRP-heuristics such as Savings or Sweep do not allow for heterogeneous service times. Alternatively, insertion procedures tend to complete one tour before starting the next. Thus, we apply a compromise between Savings and Insertion procedures in our construction heuristic, combined with an improvement procedure. Our method generates the K tours simultaneously. It randomly selects an unassigned customer and assigns it by means of cheapest insertion. The cheapest insertion considers both heterogeneous service times and potential investments. To keep the initial tours balanced, an insertion in a tour is only feasible if the resulting tour has at most one more assigned customer than the current tour with the fewest assigned customers. We generate a set of initial solutions following this procedure. For each initial solution, we use an improvement procedure by randomly removing customers from one tour and checking for an improved insertion position in the remaining tours. In the improvement steps, we omit the balancing constraint. The best found solution is then applied in our experiments.

The algorithmic details are presented in Algorithm 1 in Appendix 1. In our computational study, we generate a set of 200 initial solutions and cycle through 1000 improvement iterations per initial solution. The runtimes are generally less than one second with experiments performed on an AMD Ryzen Threadripper and 64 GB memory. However, because the heuristic is applied 24,000 times for the evaluation of a policy, this leads to runtimes of about 6 hours for each parametrization of a policy.

5. Computational Evaluation

In this section, we describe the characteristics of the representative problem that we use to depict the interplay of the three investment dimensions. We outline the three policies that benchmark the investment policy. We then identify the values for each of the investment dimensions used for the computational analysis, demonstrating the impact these have on solution quality. Finally, we present the computational analysis.

5.1. Representative Problem

To evaluate the long-term impact of tactical investments in familiarity, we utilize a representative retail distribution problem that offers a real world context for our analysis. Consider a set of retail stores within a conurbation that are to be served from a central warehouse by a fleet of vehicles. Stores request delivery frequently, but not every day. Requesting stores are known just prior to the start of operations on a particular day, with the capacity of each vehicle sufficient to handle all assigned demand. We consider a simplified binary learning pattern of service times motivated by the driver's knowledge of a store's receiving process. This process may include locating the store, finding the correct unloading area/dock, reporting to the store manager, unloading goods and completing delivery review. After completing the receiving process at a particular store, a driver will likely perform the necessary sequence of actions considerably more quickly during the following visit.

We generate instances of this representative problem with a horizon that allows for both investment in familiarity and the opportunity to observe the impact of that investment. Also, in theory, there may be realizations where no solution can be found with the total work time constraints (1b) being satisfied. In our experiments, the instances are created in a way that this is avoided and a feasible solution may always be found. Specifically, the problem includes 120 retail stores uniformly distributed in a square 50 KM by 50 KM service area, with the warehouse located at the center of the square. Each day, each store has a probability of requesting service of $\nu = \frac{70}{120}$. Thus, in expectation, 70 stores are randomly selected to request service from the warehouse. We utilize ten drivers over the entire problem horizon, such that each driver serves on average seven stores per day, a common value in delivery to retail stores (Figliozzi 2007). Each driver has a maximum working time of eight hours per day. For simplicity we consider that drivers move at a pace of 30 KPH in Euclidean space. Literature suggests that Euclidean distances should be multiplied by a factor of 1.3 to approximate road network distances (Boscoe et al. 2012). Thus, we can assume a driving speed of $30 \times 1.3 \approx 40$ KPH within the service area. With respect to service time, a driver requires $t^{\text{unfamiliar}} = 40$ minutes to visit an unfamiliar store. A familiar store requires half of that time, such that $t^{\text{familiar}} = 20$ minutes. The problem horizon P is set at 120 periods, or approximately half a year of operations. Each of these parameters is selected to ensure that a feasible solution may be constructed for each day of the problem horizon. Using these parameters, we generate 200 instance realizations, each comprised of 120 days, such that we determine 24,000 VRP solutions per policy.

5.2. Benchmark Policies

To evaluate the benefit of a tactical investment in familiarity, we compare the algorithm using several different policies. The policies differ only by the cost function that is used within our heuristic. An *investment* policy follows our proposed strategy of routing drivers to visit unfamiliar customers early in the problem horizon to gain later benefit. The cost function is as previously defined $C_I(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} (c_{ij} + t_{ikp} - I_{ikp}) x_{ijkp}$. In addition to this policy, we determine solutions using a:

- *myopic* policy, where $C_M(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} (c_{ij} + t_{ikp}) x_{ijkp}$, which considers routing and service time costs without familiarity investment. This policy balances routing and service, but without consideration for future decision making.

- *routing* policy, where $C_R(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} c_{ij} x_{ijkp}$, which considers routing cost first, and assumes a static service time for all driver-customer pairs when determining total work time to establish feasibility as defined by constraint (1b). As familiarity is not considered, we must assume that all service times are the unfamiliar time of 40 minutes in order to ensure that all routes are feasible. Service time may be reduced after route creation on each day via neighborhood search, iteratively switching the routes of two drivers with a goal of increasing the number of drivers matched with familiar customers. We apply 200 such iterations per found solution.

- *service time* policy, where $C_{ST}(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} t_{ikp} x_{ijkp}$, which considers service time cost first. In case of several identical solutions, the one with minimal routing cost is selected. While this policy does not incorporate routing cost, total work time is also limited by constraint (1b).

It may be expected that a policy that considers both routing and service time costs will have a lower total cost than a policy that only considers one of those costs. However, we include both routing and service policies because it is fairly common in many industries for one of these to be adopted and it is of interest to see how they perform relative to the investment or myopic policies.

5.3. Investment Dimension Values

The parameters that most directly influence the investment policy are the dimensions, $\mathcal{M}, q$ and t_{max} (which determines the value of variable n_{ik}), used to calculate the investment matrix I. If I is too large, it is difficult to recover the investment in the later stages of the problem horizon. If I is too small, driver-customer relationships that may be beneficial are not developed. To evaluate the interplay of these dimensions, show how they each impact total cost, and determine the best value of each to use for the analysis comparing policies, we run the 200 test scenarios for each value over a specified range. A range of [0,30] in increments of 5 is tested for the investment magnitude, $\mathcal{M}$, and the parameter defining geographic proximity, t_{max}. The values 0, 0.5, 1 and 2 are tested for the discount factor, q. To find a point of comparison, we first determine the dimension values that generate the solution with the lowest cost. We test every combination of these dimensions, such that $7 \times 7 \times 4 = 196$ combinations are evaluated. The lowest total cost is realized with an investment magnitude of $\mathcal{M} = 20$ minutes, an existing familiarity discount factor of $q = 0.5$, and a maximum travel range of $t_{max} = 15$ minutes.

Appropriately calibrating each of these three dimension values is vital to solution quality and can influence decision making. To determine the sensitivity of the total cost to the dimension values, two are fixed at the value that generates the lowest cost and the third is tested across the range of possible values. Figure 3 shows the increase in total cost as each dimension value is modified from the value with the lowest cost.

Modifying any of the investment dimensions results in an increase in total cost, highlighting the interdependencies between the three dimensions. Cost is most significantly impacted by a change in the investment magnitude when this dimension is greater than 20 minutes. Setting this parameter above 20 minutes creates more familiarity, but the additional time investment does not result in a comparable reduction in routing or service time later in the problem horizon. When $\mathcal{M}$ is less than 20 minutes, the incentive to develop more driver-customer relationships declines and opportunities that decrease cost later in the problem horizon are not explored.

The geographic proximity of unfamiliar customers also impacts the total cost. Expanding t_{max} allows for investments that have a driver gain familiarity with a larger set of customers, but with diminishing returns. When t_{max} is greater than 15 minutes, the cost of routing a driver to visit more distant unfamiliar customers outweighs the benefit of increased familiarity. Limiting the distance the driver may travel to visit an unfamiliar customer to less than 15 minutes has a similar effect to limiting $\mathcal{M}$, as fewer driver-customer relationships may be developed.

The discount factor on the existing familiarity of drivers with a customer has a considerable impact on total cost when set equal to 0, such that this component of I is equal to 1 regardless of the value of $F(i)$ for any customer. By eliminating this component, investments may be made when numerous drivers are already familiar with a customer, generating redundant relationships that do not offer benefit. To a lesser extent, the total cost increases with q above 0.5, with I decreasing as q increases. Thus, the number of drivers that a

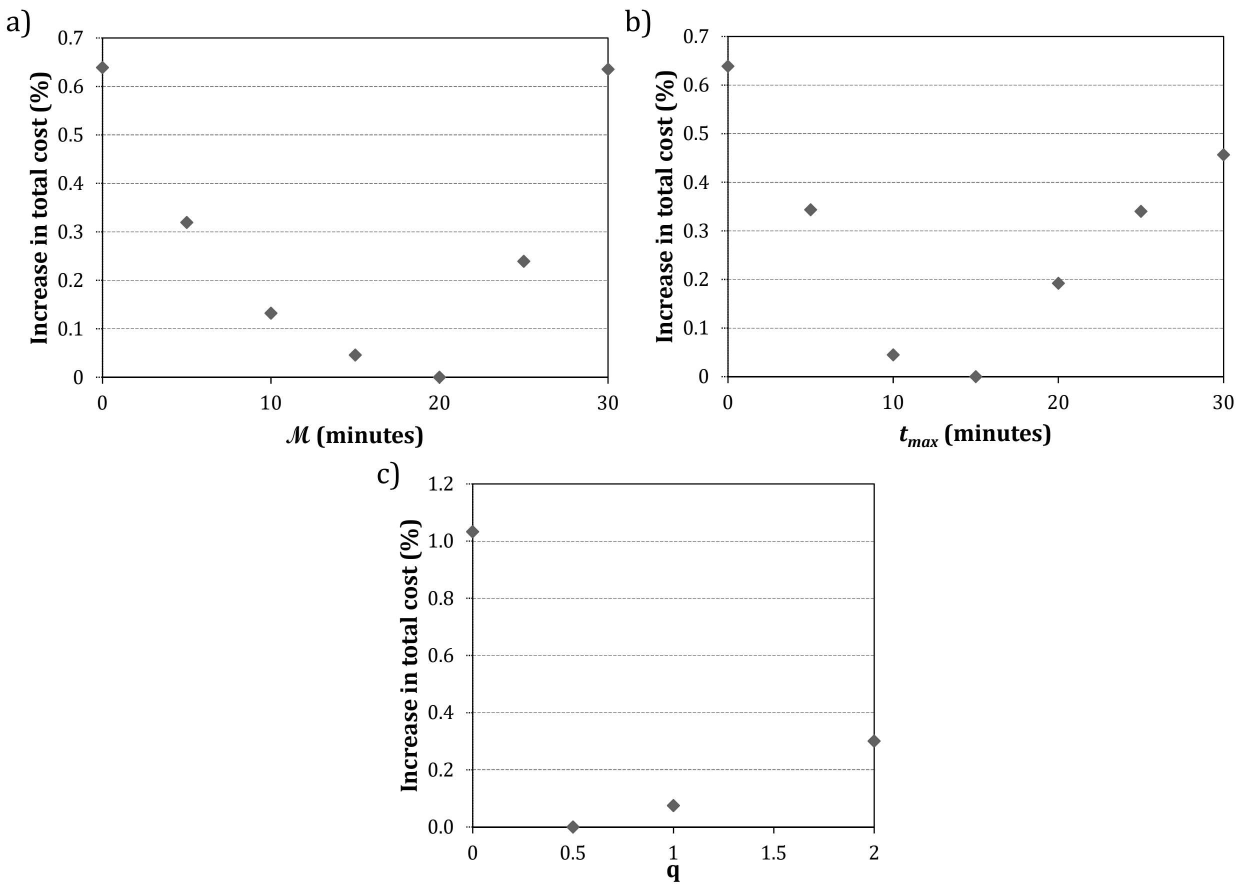

Figure 3: Increase in total cost relative to the solution using the best performing dimension parameter for a) investment magnitude; b) geographic proximity; and c) existing familiarity

customer has existing familiarity with should not be strictly limited. Numerous drivers may be unnecessary, but a second or third driver with whom a customer is familiar offers operational flexibility. This is further explored in Section 5.4. We use the values of $\mathcal{M} = 20$, $q = 0.5$ and $t_{max} = 15$ for the remainder of the analysis.

While the combination of parameters determined here will vary for different problem characteristics, we highlight the importance of correctly calibrating these values to minimize cost. A transportation manager should be aware that the amount of time she is willing to invest in driver's gaining new familiarity, the distance she is willing to allow drivers to travel to visit unfamiliar customers, and the number of drivers she wants familiar with one customer all have a considerable impact on the savings that can be gained through an investment policy.

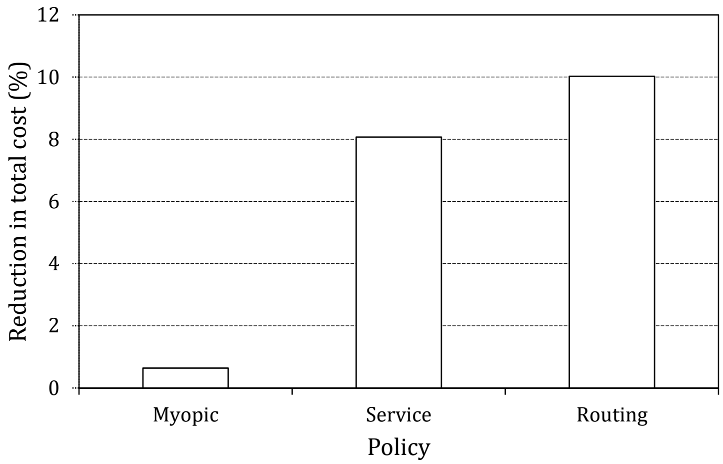

Figure 4: Reduction in total cost of investment policy relative to each of the benchmark policies

5.4. Comparative Analysis

In the following, we compare the results of the investment policy and the benchmark policies. We first analyze the cost of the policies, determining how decision making changes over time by looking at the breakdown in daily costs, and then we observe the impact on familiarity. We note that these results hold for the specific problem conditions that were tested and they may vary as the problem instance varies. However, they do indicate the potential for benefit from the use of an investment policy.

Cost

We first compare the various policies using total cost. When evaluating one policy, we find the standard error between costs across the 200 instances to be below 0.1%. Figure 4 shows the reduction in cost for the investment policy when compared to each of the benchmark policies. This is calculated as $\frac{C_{\text{policy}} - C_I}{C_{\text{policy}}}$ for each policy. By only focusing on service or routing cost, the total cost is 8%-10% greater than with the investment policy. When both costs are minimized using the short-term, myopic focus, the tactical investment policy still maintains a cost advantage (in the trucking industry where profit margins are often in the range of 1%-5%, such savings can be significant). The investment in developing driver-customer relationships over the 120 day problem horizon has not only been recovered, it has led to additional savings.

To provide insight into how the routing and service costs differ by policy, Figure 5 presents the average daily routing and service costs for each policy. The investment policy has a slightly greater service cost than the myopic policy as it develops more familiar driver-customer pairs. However, this increased service cost leads to an even greater reduction in routing cost, as the increased familiarity allows for more routing flexibility. The service policy further reduces the service cost, but the routing cost is much greater as it is not considered in the objective function. Interestingly, the routing policy has a greater routing cost than

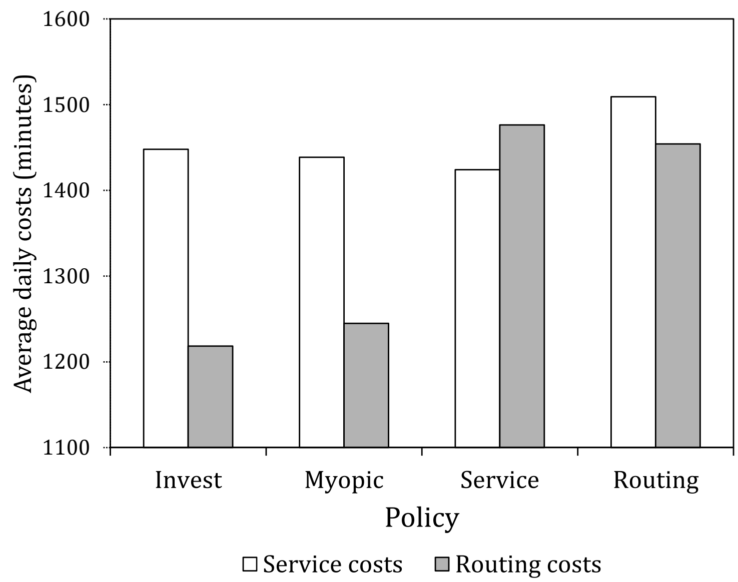

Figure 5: Average daily routing and service costs (in minutes) for each policy

both the investment and myopic policies. By assuming homogeneous service times of 40 minutes for all driver-customer pairs, feasible routes are shorter and have less flexibility than with policies that consider heterogeneous service times. In Appendix 2, we present an additional analysis of more flexible routing policies that take into consideration the length of service time. These further underline the importance of considering both routing and service costs in order to minimize total cost.

We may dissect the costs by day, with Figure 6 showing the daily cost of the investment policy separated into routing and service costs. The service cost is further divided into the cost when a driver visits an unfamiliar or familiar customer. We first observe that the overall cost decreases over time, which is intuitive since service time decreases as drivers have gained familiarity with more customers in later periods.

The investment policy initially incurs a significant service cost for drivers to visit unfamiliar customers as familiarity is generated. By the 25th day the service cost for unfamiliar driver-customer pairs is below 80 minutes, such that fewer than two new pairs are introduced each day moving forward. This indicates that the investment in familiarity is primarily early in the problem horizon, allowing for the benefits to be realized over a longer period. The routing cost increases over the first several days (for all policies) as detours are made to exploit the short service times for the already familiar driver-customer pairs. This cost quickly stabilizes after sufficient familiarity has been established.

A further evaluation may be made by comparing the daily costs of the investment policy and the myopic policy. We focus here on a comparison with the myopic policy as the difference in cost with the routing and service policies is considerably greater, as shown in Figure 4. Figure 7 shows the difference in daily

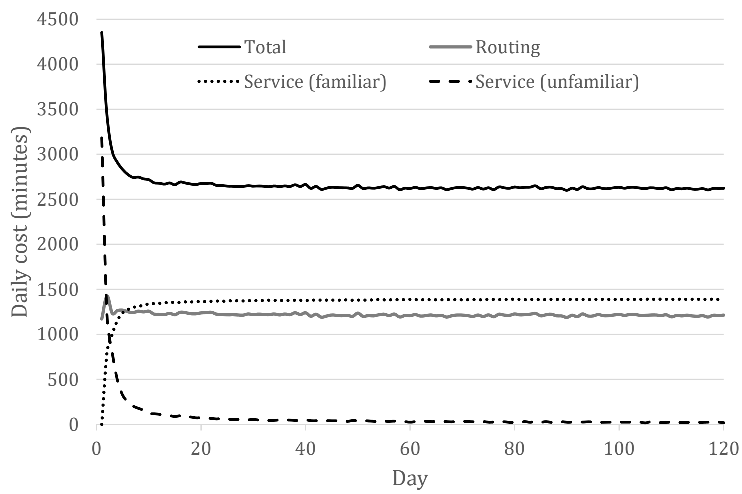

Figure 6: Daily cost for the investment policy separated into routing and service costs for familiar and unfamiliar driver-customer pairs

cost between the investment and myopic policies, separated into routing and service costs. The service cost for the investment policy is considerably greater than for the myopic policy over the first several days as investments are made to develop familiarity between more driver-customer pairs. As there is a difference in service time of 20 minutes between visiting a familiar and unfamiliar customer, we note that over the first eight days at least one additional driver-customer pair gained familiarity each day, with at least two additional pairs over the first four days. The difference in service cost eventually converges close to zero.

The benefit of this investment is reflected in the difference in routing cost. With the additional flexibility offered by more familiar driver-customer pairs, the routing cost is considerably lower for the investment policy early in the problem horizon. As the myopic policy develops familiarity organically over time, the difference in routing cost decreases. However, over the course of the majority of the problem horizon, this difference is greater than the difference in service cost. This can be seen in the total difference in cost. Over the first three days, the investment in familiarity is such that the difference in service cost outweighs the difference in routing cost. By day four, the total cost of the investment policy drops below that of the myopic policy. This behavior shows that the investment in familiarity quickly realizes benefit.

Familiarity

While a cost analysis provides a contrast between the policies, the level of familiarity between customers and drivers is also of interest. Figure 8 shows the amount that familiarity changes across the different policies. Total familiarity is measured by summing the number of driver-customer pairs familiar with each

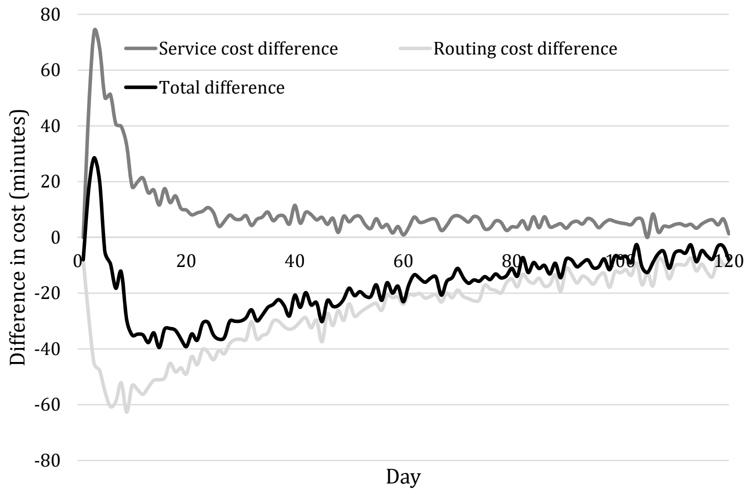

Figure 7: Difference in daily cost between investment and myopic policies separated into routing and service costs

other at the end of the problem horizon. The percentage difference in familiarity is:

$$\frac{\sum_{i\in N}(|F_{policy}(i)| - |F_I(i)|)}{\sum_{i\in N}(|F_{policy}(i)|)}$$

for each policy, where $|F_{policy}(i)|$ and $|F_I(i)|$ are the number of drivers familiar with customer i when using the tested policy and when using the investment policy, respectively. The standard error for familiarity across 200 instances with one policy is below 1%. Familiarity is decreased by almost 25% with the myopic policy and, as seen in Figure 4, this comes at a greater cost. That is, the investment policy has allowed for an increase in driver familiarity not at a cost, but with cost savings. As the service time policy focuses on exploiting existing familiarity, with most customers only gaining familiarity with one driver, the investment policy has an increase in familiarity of more than 100%. Incorporating a homogeneous service cost, the routing policy generates many unnecessary driver-customer relationships as evidenced by an increase in familiarity of almost 60%. More than half of the customers are visited by five or more drivers over the problem horizon, which is likely excessive from a customer perspective.

We may similarly analyze the distribution of familiarity to customers for all policies. It may be desirable to have at least two drivers familiar with a customer in the event that there is driver turnover or a driver is on vacation, but interacting with more than two or three drivers may degrade the service experience for the customer and limit the familiarity benefit per driver. Figure 9a presents the number of customers who are familiar with a certain number of drivers. Under the investment policy, almost all customers are familiar with two or three drivers. The myopic policy has considerably more customers that are only familiar with one driver, while almost all customers are visited by one driver under the service policy. While attractive from

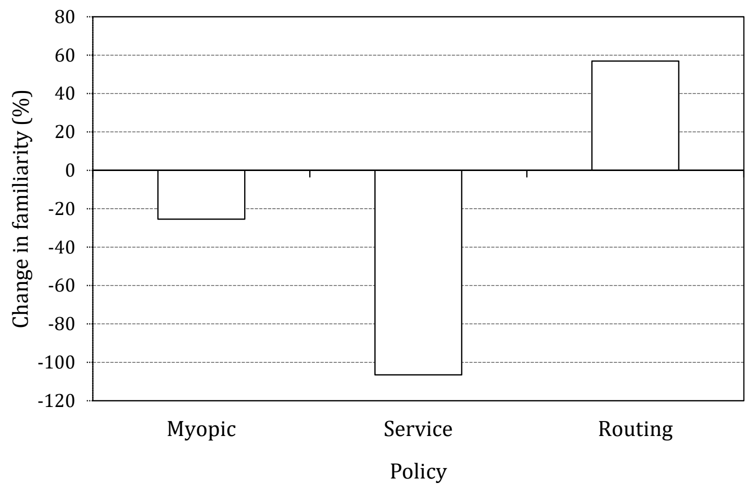

Figure 8: Percentage change in the number of familiar driver-customer pairs for each policy relative to investment policy

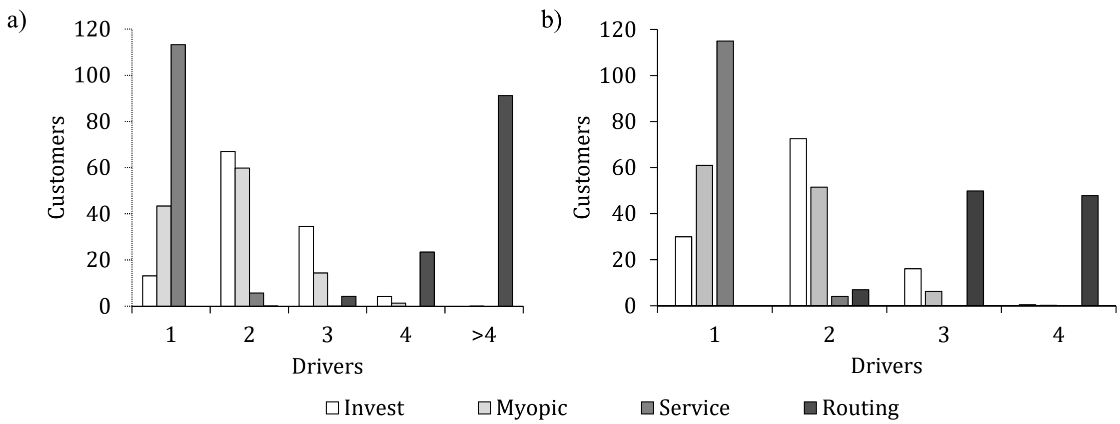

Figure 9: Number of customers who a) are familiar with 1, 2, 3 or 4 drivers; and b) have at least 5 visits by 1, 2, 3 or 4 of the drivers

a service perspective, assigning only one driver to a customer limits operational flexibility. In addition to allowing for driver off days or turnover, having an additional driver that is familiar with a customer provides more routing options. This may slightly increase service cost, but it results in a greater decrease in routing cost. This may ultimately allow for more customers to be served each day and, if demand is sufficient, this could prevent customers from being denied service due to capacity limits. Under these circumstances, it may be necessary to provide some explanation to customers served with two drivers.

In reality, there may be some circumstances where full familiarity is only gained after the driver has made several visits. To evaluate this level of familiarity, we determine how often each customer is visited by a specific driver. Figure 9b shows the number of customers who have been visited by 1, 2, 3 or 4 of the drivers, revealing a similar structure to Figure 9a. With the investment policy most customers are visited more than five times by at least two drivers, while the myopic policy has the majority of customers who are visited more than five times by only one driver. The service policy has almost all customers visited by one driver, while the routing policy has almost all customers visited five times by at least three different drivers.

5.5. The Value of Investments

In the previous section, we show that the investment policy leads to an increase in familiarity and that this familiarity decreases daily operational costs. This increase in familiarity may have another, long-term value for companies by allowing them to hedge against driver or customer turnover. While turnover in drivers or customers results in loss of familiarity, the additional familiarity generated through the investment policy earlier in the problem horizon may allow for more flexibility and a reduction in cost. To quantify this, we take the familiarity solutions obtained by each of the policies and run two additional experiments that consider driver turnover and customer turnover. In reality, turnover is a process with both drivers and customers entering and leaving the system over the course of weeks and months. Because of computational limitation, we approximate the impact of turnover by considering the day after a driver or customer turnover happened.

To analyze driver turnover, we replace an existing driver with a new driver by setting the familiarity vector for this driver to zero at the end of the problem horizon, extending the horizon by one day. We apply the heuristic to all four policies for this additional day to determine the difference in cost when a new driver is introduced into the system. We do this for every driver, resulting in 10 new familiarity matrices for each of the 200 problem instances. For each matrix, we execute 100 evaluation runs of this final day. We then calculate the average value of the 1000 evaluation runs for each of the 200 instances. Customer turnover is analyzed similarly, except we randomly replace approximately 10% of the customers with new customers. The new customers have no familiar drivers and new coordinates in the service area sampled from the same uniform distribution as the original customers. We repeat this procedure to generate 10 new familiarity matrices to test both policies. We again execute 100 evaluation runs for each new matrix and calculate the average value for each of the 200 instances.

Table 2 presents the results of this analysis. With driver turnover, the average cost reduction for the investment policy relative to the myopic policy is 1.8%, while it is 0.6% for customer turnover (as indicated earlier, a cost savings of 1% can be significant for a trucking firm). There are more opportunities to save costs with driver turnover as the additional driver-customer relationships generated by the investment policy allow for a driver who leaves to be more easily replaced by an existing driver. The myopic policy is more likely to have to assign a new driver to a customer as there is a greater probability that customer was only familiar with the one driver who was replaced. The benefit of the investment policy is even greater relative to the service and routing policies. With the service policy, almost every customer is familiar with only one driver, so it is almost guaranteed that the customers who were served by the driver who leaves will be visited by a new driver. Even with the redundant familiarity developed over the problem horizon, the routing policy performs very poorly as it considers homogeneous service times when creating routes. The alternate routing policies presented in Appendix 2 perform considerably better, with both showing a 2.9% cost reduction relative to the investment policy. This is because both have established familiarity between many more driver-customer pairs, such that it is less costly to replace the driver who leaves, and heterogeneous service times are considered during route development.

With customer turnover, the investment policy still offers a reduction in cost as the additional familiarity allows for more routing flexibility, but there is no reduction in service cost as it is not possible to have established familiarity with a new customer. Thus, there is a slight reduction in the benefit of the investment policy across all other policies with customer turnover. The same is true of both alternate routing policies found in Appendix 2, as having many drivers familiar with many customers does not have the same advantage when there are new customers. Alternate routing 1 results in a 0.1% increase in cost relative to the investment policy and alternate routing 2 results in a 0.1% decrease in cost.

We also observe the range of values for the percentage difference in cost across the 200 instances tested. Because demand patterns differ for each evaluation run, the cost reduction varies from run to run and some results lead to a cost reduction for the myopic policy relative to the investment policy. The range in reduction in cost for the investment policy is fairly narrow relative to the myopic and routing policies with both driver and customer turnover, while at least 97% of runs result in a positive cost reduction relative to all policies. The range is greater for the service policy as turnover guarantees that new familiarity will have to be established, but the cost of this can vary dependent on which driver or customer has been replaced. While this analysis only considers one day of turnovers, the observations indicate that the amount of investment should be coupled to the problem horizon and turnover frequency. For a long horizon, establishing more familiarity may lead to higher savings. For higher turnover frequency, more familiarity leads to both a decrease in the familiarity lost by drivers and customers leaving the system and more flexibility when turnovers happen.

6. Conclusion

In this research, we have addressed the value of familiarity between drivers and customers in a service routing application. To this end, we have introduced a stochastic and dynamic multi-period routing problem with binary driver-customer familiarity (VRPBDCF). Considering a large problem horizon of several months, a subset of customers requires service by a fleet of drivers on each day. Once a driver has visited a customer, the driver gains familiarity with this customer. The familiarity leads to more efficient service,

Table 2: Results for benchmark policies relative to investment policy with driver or customer turnover

Type of turnover	Policy	Average cost reduction	Lower range	Upper range	% of runs with positive cost reduction
Driver	Myopic	1.8%	0.1%	3.8%	99.5%
	Service	8.3%	5.1%	13.4%	100%
	Routing	22.2%	20.2%	24.1%	100%
Customer	Myopic	0.6%	-0.7%	2.2%	97%
	Service	7.0%	4.0%	11.9%	100%
	Routing	19.9%	18.4%	21.8%	100%

with lower costs for serving the customer with that driver.

In a computational study focusing on a retail distribution problem, we have shown the importance of considering the long-term impact on familiarity in day-to-day routing decisions. To this end, we have presented a method that explicitly *invests* in establishing familiarity for selected driver-customer pairs. An investment of reasonable size is made when the driver already knows customers in the neighborhood of the customer under consideration and when there is not already a sufficient number of drivers familiar with the customer. We have shown that under these problem conditions investment decisions can reduce operational cost in the long run while substantially increasing the overall level of familiarity between drivers and customers. This increased familiarity further reduces cost when driver or customer turnover is considered. We have also shown that focusing only on routing or on exploiting existing familiarity leads to substantial increases in operational cost. Most importantly, we present insights that are highly relevant to a transportation manager, that the familiarity gains of an investment policy are contingent on the proper calibration of three dimensions: investment magnitude, existing familiarity between drivers and customers, and geographic proximity of unfamiliar customers to a driver's known customers.

There are several avenues for future research. Because of the tremendous computational burden, our study focuses on one problem setting, a retail distribution network, to analyze the impact of familiarity and investment decisions. This instance setting also reflects characteristics of less-than-truckload, technician, or service routing where customers are served repeatedly and the time for serving customers is significant. Future research may focus on variants of the VRPBDCF that integrate additional, application-specific characteristics such as time-windows or capacity constraints. In our study, we assumed binary familiarity, meaning that drivers gain familiarity with the first visit of the customer, which, while reasonable in the context of our problem, is a simplifying assumption. Future research may focus on more elaborate models of learning and forgetting. We note that both our Markov-model and our method can be transferred to such learning models by suspending the integrality assumption within the familiarity matrix.

Finally, our policy establishes a more familiar driver workforce at no cost (even reducing cost). Such an increase in familiarity is shown to hedge against changes in the workforce and customer base. While our turnover experiments were limited due to the computational burden, future research may take a closer look at turnovers (and their reasons) in both of these populations. The amount of investment should be tailored to the turnover rate when the relationship between the two is not trivial. With low turnover, investments remain in the system and provide long term gain. However, when turnovers occur, some familiarity investments may be lost, but investments in drivers or customers that remain allow for more flexible reactions to such turnovers. Additionally, while consistency in customer service may come at an increased cost, it generally leads to happier customers and drivers and, therefore, fewer turnovers.

References

Battini, Daria, Maurizio Faccio, Alessandro Persona, Stefan Røpke, Giorgia Zanin. 2015. Routing strategy in a distribution network when the driver learning effect is considered. *International Journal of Logistics Systems and Management* **21**(3) 385–411.

Boscoe, Francis P, Kevin A Henry, Michael S Zdeb. 2012. A nationwide comparison of driving distance versus straight-line distance to hospitals. *The Professional Geographer* **64**(2) 188–196.

Braekers, Kris, Attila A Kovacs. 2016. A multi-period dial-a-ride problem with driver consistency. *Transportation Research Part B: Methodological* **94** 355–377.

Campelo, Pedro, Fábio Neves-Moreira, Pedro Amorim, Bernardo Almada-Lobo. 2018. Consistent vehicle routing problem with service level agreements: a case study in the pharmaceutical distribution sector. *European Journal of Operational Research* .

Chen, Xi, Barrett W Thomas, Mike Hewitt. 2016. The technician routing problem with experience-based service times. *Omega* **61** 49–61.

Chen, Xi, Barrett W Thomas, Mike Hewitt. 2017. Multi-period technician scheduling with experience-based service times and stochastic customers. *Computers & Operations Research* **82** 1–14.

Coelho, Leandro C, Jean-François Cordeau, Gilbert Laporte. 2012. Consistency in multi-vehicle inventory-routing. *Transportation Research Part C: Emerging Technologies* **24** 270–287.

De Bruecker, Philippe, Jorne Van den Bergh, Jeroen Beliën, Erik Demeulemeester. 2015. Workforce planning incorporating skills: State of the art. *European Journal of Operational Research* **243**(1) 1–16.

Feillet, Dominique, Thierry Garaix, Fabien Lehuédé, Olivier Péton, Dominique Quadri. 2014. A new consistent vehicle routing problem for the transportation of people with disabilities. *Networks* **63**(3) 211–224.

Figliozzi, Miguel Andres. 2007. Analysis of the efficiency of urban commercial vehicle tours: Data collection, methodology, and policy implications. *Transportation Research Part B: Methodological* **41**(9) 1014–1032.

Fowler, John W, Pornsarun Wirojanagud, Esma S Gel. 2008. Heuristics for workforce planning with worker differences. *European Journal of Operational Research* **190**(3) 724–740.

Groër, Chris, Bruce Golden, Edward Wasil. 2009. The consistent vehicle routing problem. *Manufacturing & service operations management* **11**(4) 630–643.

Haughton, Michael A. 2007. Assigning delivery routes to drivers under variable customer demands. *Transportation Research Part E: Logistics and Transportation Review* **43**(2) 157–172.

Haughton, Michael A. 2008. The efficacy of exclusive territory assignments to delivery vehicle drivers. *European Journal of Operational Research* **184**(1) 24–38.

Hewitt, Mike, Maciek Nowak, Nisha Nataraj. 2016. Planning strategies for home health care delivery. *Asia-Pacific Journal of Operational Research* **33**(05) 1650041.

Kovacs, Attila A, Bruce L Golden, Richard F Hartl, Sophie N Parragh. 2014. Vehicle routing problems in which consistency considerations are important: A survey. *Networks* **64**(3) 192–213.

Kovacs, Attila A, Bruce L Golden, Richard F Hartl, Sophie N Parragh. 2015a. The generalized consistent vehicle routing problem. *Transportation Science* **49**(4) 796–816.

Kovacs, Attila A, Sophie N Parragh, Richard F Hartl. 2015b. The multi-objective generalized consistent vehicle routing problem. *European Journal of Operational Research* **247**(2) 441–458.

Li, Ninghua, Ling X Li. 2000. Modeling staffing flexibility: A case of china. *European Journal of Operational Research* **124**(2) 255–266.

Lian, Kunlei, Ashlea Bennett Milburn, Ronald L Rardin. 2016. An improved multi-directional local search algorithm for the multi-objective consistent vehicle routing problem. *IIE Transactions* **48**(10) 975–992.

Luo, Zhixing, Hu Qin, ChanHou Che, Andrew Lim. 2015. On service consistency in multi-period vehicle routing. *European Journal of Operational Research* **243**(3) 731–744.

Marentette, Kenneth A, Alan W Johnson, Lisa Mills. 2009. A measure of cross-training benefit versus job skill specialization. *Computers & Industrial Engineering* **57**(3) 937–940.

Patel, Khyati, Jeff Wright. 2018. Jasper girl to be featured in UPS short film after friendship with delivery driver. Retrieved from http://www.ktre.com/2018/10/23/jasper-girl-be-featured-ups-commercial-after-friendship-with-delivery-driver/.

Powell, Warren B. 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, *Wiley Series in Probability and Statistics*, vol. 842. John Wiley & Sons, New York.

Powell, Warren B. 2017. A unified framework for stochastic optimization. Available from `http://castlelab.princeton.edu/Papers/Powell-UnifiedFrameworkStochasticOptimization_July222017.pdf`, accessed on August 1, 2017.

Premack, Rachel. 2019. UPS drivers are ditching commercial navigation apps for an in-house tool. The chief engineering and information officer explains how it's already saving the company millions. *Business Insider* July 30, www.businessinsider.com/ups-drivers-navigation-tool-saves-millions-2019-7.

Puterman, Martin L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, New York.

Russell, David, Robert J Rosati, Peri Rosenfeld, Joan M Marren. 2011. Continuity in home health care: Is consistency in nursing personnel associated with better patient outcomes? *Journal for Healthcare Quality* **33**(6) 33–39.

Smilowitz, Karen, Maciek Nowak, Tingting Jiang. 2013. Workforce management in periodic delivery operations. *Transportation Science* **47**(2) 214–230.

Subramanyam, Anirudh, Chrysanthos E Gounaris. 2016. A branch-and-cut framework for the consistent traveling salesman problem. *European Journal of Operational Research* **248**(2) 384–395.

Subramanyam, Anirudh, Chrysanthos E Gounaris. 2017. A decomposition algorithm for the consistent traveling salesman problem with vehicle idling. *Transportation Science* **52**(2) 386–401.

Sungur, Ilgaz, Yingtao Ren, Fernando Ordóñez, Maged Dessouky, Hongsheng Zhong. 2010. A model and algorithm for the courier delivery problem with uncertainty. *Transportation Science* **44**(2) 193–205.

Tarantilis, Christos D, F Stavropoulou, Panagiotis P Repoussis. 2012. A template-based tabu search algorithm for the consistent vehicle routing problem. *Expert Systems with Applications* **39**(4) 4233–4239.

Valeva, Silviya, Mike Hewitt, Barrett W Thomas. 2017a. A matheuristic for workforce planning with employee learning and stochastic demand. *International Journal of Production Research* **55**(24) 7380–7397.

Valeva, Silviya, Mike Hewitt, Barrett W Thomas, Kenneth G Brown. 2017b. Balancing flexibility and inventory in workforce planning with learning. *International Journal of Production Economics* **183** 194–207.

Wright, Theodore P. 1936. Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences* **3**(4) 122–128.

Xu, Zefeng, Yanguang Cai. 2018. Variable neighborhood search for consistent vehicle routing problem. *Expert Systems with Applications* .

Zhong, Hongsheng, Randolph W Hall, Maged Dessouky. 2007. Territory planning and vehicle dispatching with driver learning. *Transportation Science* **41**(1) 74–89.

Appendix 1

In Appendix 1, we present the details of our routing heuristic. The procedure is described in Algorithm 1. Input for the algorithm is the set of active customers N_p, the familiarity matrix F_p, and the investment matrix I. Output is the routing decision a_p. The algorithm creates m starting solutions as follows. A solution starts with an empty set of routes a_{temp}. Subsequently, a customer n_{temp} is selected with function *Random*(N_{temp}) from the set of unassigned customers N_{temp}. This customer is inserted by function

BalancedInsertion$(a_{\text{temp}}, n_{\text{temp}}, F_p, I)$. This function maintains the current assignments and sequences in a_{temp}. It then checks for routes in a_{temp} whether insertion is feasible because it does not cause an imbalance between the routes. To this end, the function determines the route in a_{temp} with the minimal number of assignments. Every route in a_{temp} with the same number of assignments is eligible for insertion. For each of the eligible routes, the algorithm determines the insertion position that increases the current routing cost by the least possible amount. The calculation considers service times based on F_p and investments by I. Because of the investment, the calculated routing cost may be smaller than the actual cost. If the duration of the route exceeds the time limit, the route is not considered further for this customer. For all feasible routes the algorithm then compares the calculated increase in routing cost and inserts customer n_{temp} in the route with the smallest value. Then, customer n_{temp} is removed from the set of unassigned customers N_{temp}.

Eventually, the algorithm obtains a starting solution a_{temp}. This solution is now improved h times. The improvement step comprises randomly selecting a customer n_{temp}, removing this customer from the routes a_{temp}, and inserting the customer with function *Insertion*$(a_{\text{temp}}, n_{\text{temp}}, F_p, I)$. This function works as *BalancedInsertion* but without the balancing constraint. Eventually, a solution candidate is obtained. The candidate is compared to the current best solution. For comparison, function *Cost*$(a_{\text{temp}}, F_p, I)$ is applied calculating the overall routing cost with respect to the investment matrix I. The best solution candidate over all m runs is returned by the algorithm. In our computational study, we set $m = 200$ and $h = 1000$.

Appendix 2

As shown in Sections 5.4 and 5.5, the routing policy performs quite poorly relative to the other policies because heterogeneous service times are not considered. While this is reflective of many existing commercial routing software packages, in which service time is static, it is of interest to determine if this policy can be competitive with consideration for heterogeneous service times. To evaluate this, we test two alternate routing policies, both of which still use the cost function $C_R(S_p, a_p) = \sum_{i,j \in N_p^0, k \in K} c_{ij} x_{ijkp}$. These policies are:

- Alternate routing 1 - considers routing cost first, incorporating heterogeneous service times that are adjusted based on previously developed familiarity when determining total work time feasibility of routes.
- Alternate routing 2 - same as alternate routing 1, with swapping of drivers after route development. Service time may be reduced after route creation on each day via neighborhood search, iteratively

Algorithm 1: Routing Heuristic

Input : Active Customers N_p, Familiarity Matrix F_p, Investment Matrix I
Output : Routing Solution a_p

$C^* \leftarrow \infty$ `// Best Found Solution Value`
for $(i = 1, \dots, m)$ `// Initial Starting Solutions`
do
  $N_{\text{temp}} \leftarrow N_p$ `// Set of Customers`
  $a_{\text{temp}} \leftarrow$ empty routes `// Starting Solution`
  while $(N_{temp} \neq \emptyset)$ `// Insertion Procedure`
  do
    $n_{\text{temp}} \leftarrow \mathit{Random}(N_{\text{temp}})$ `// Select Random Customers`
    $a_{\text{temp}} \leftarrow \mathit{BalancedInsertion}(a_{\text{temp}}, n_{\text{temp}}, F_p, I)$ `// Insert Customer`
    $N_{\text{temp}} \leftarrow N_{\text{temp}} \backslash \{n_{\text{temp}}\}$ `// Set of Customers`
  end
  for $(j = 1, \dots, h)$ `// Improvement Step`
  do
    $n_{\text{temp}} \leftarrow \mathit{Random}(N_p)$ `// Select Random Customers`
    $a_{\text{temp}} \leftarrow \mathit{Remove}(a_{\text{temp}}, n_{\text{temp}})$ `// Remove Customer`
    $a_{\text{temp}} \leftarrow \mathit{Insertion}(a_{\text{temp}}, n_{\text{temp}}, F_p, I)$ `// Insert Customer`
  end
  $C_{\text{temp}} \leftarrow \mathit{Cost}(a_{\text{temp}}, F_p, I)$ `// Evaluate Solution`
  if $C_{temp} < C^*$ `// If Better Solution`
  then
    $a_p \leftarrow a_{\text{temp}}$ `// Update Best Solution`
    $C^* \leftarrow C_{\text{temp}}$ `// And Cost`
  end
end
return a_p

switching the routes of two drivers with a goal of increasing the number of drivers matched with familiar customers. We apply 200 such iterations per found solution.

These routing policies were both tested in the same fashion as the other policies. We evaluate them relative to the investment policy and include the original routing policy for comparison. Figure 10 presents the reduction in total cost found using the investment policy relative to the three routing policies. Alternate routing 1 finds considerable improvement by considering heterogeneous service times, creating more efficient, feasible routes using the more precise service information. Alternate routing 2 further reduces cost by swapping drivers across routes to increase familiarity.

Figure 11 presents the routing and service costs for the investment, routing and alternate policies. Alternate routing 1 has the lowest routing cost of all policies, but the greatest service costs. It finds the shortest routes on a daily basis without taking into consideration which drivers should serve which customers. While the route lengths are minimized, the service cost is much greater because of the multiple drivers visiting each customer. Alternate routing 2 reduces the service costs by swapping drivers between routes to find more familiarity. However, by doing this early in the problem horizon, some routing flexibility is sacrificed later

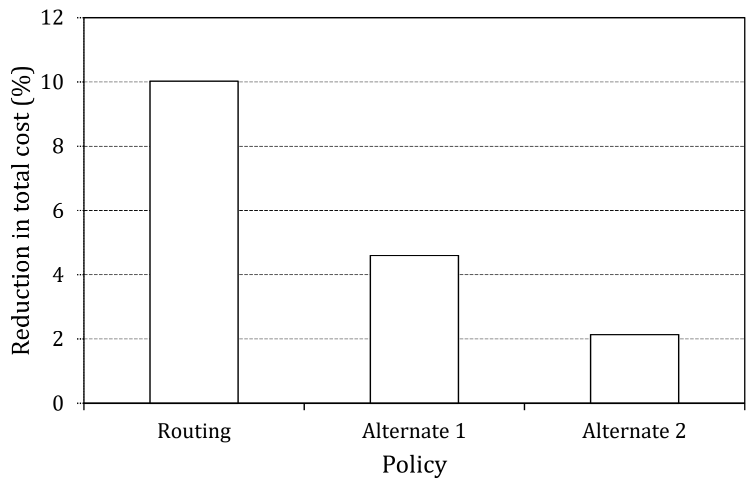

Figure 10: Reduction in total cost of investment policy relative to routing and alternate policies

and the routing cost is greater than with alternate routing 1 (and the investment policy). This emphasizes the importance of incorporating service costs into the routing, as done with both the investment and myopic policies.

The increase in familiarity across driver-customer pairs for alternate routing 1 is also shown in Figure 12. With this policy, each driver visits almost all customers over the problem horizon. Swapping drivers between routes does reduce the redundancy, but almost all customers are still visited by at least three drivers, as with the original routing policy. This indicates that even when considering heterogeneous service times and using a post-routing improvement step to reduce service times, a policy that focuses on routing results in redundancies in familiarity, increases in service cost and even an increase in routing cost.

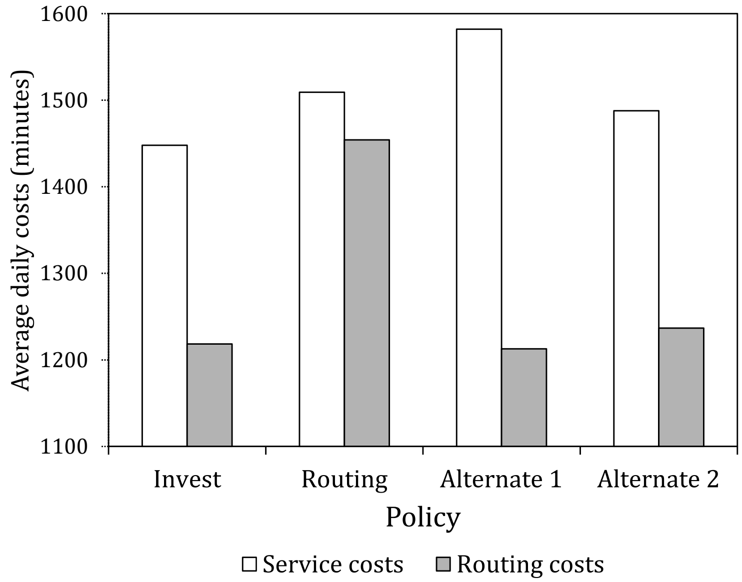

Figure 11: Average daily routing and service costs (in minutes) for investment, routing and alternate policies

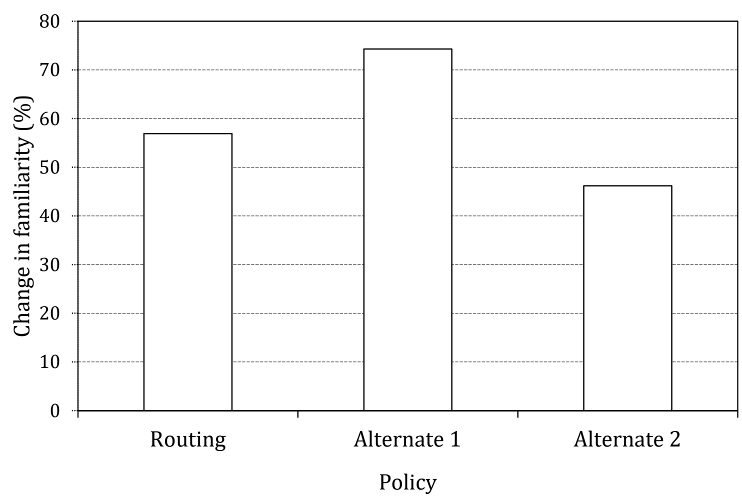

Figure 12: Percentage change in the number of familiar driver-customer pairs for routing and alternate policies relative to investment policy