Building Trust in Home Services - Stochastic Team-Orienteering with Consistency Constraints

Yongjia Song, Marlin W. Ulmer, Barrett W. Thomas, Stein W. Wallace

Abstract

In this paper, we consider service applications where drivers serve subscription customers at their homes on a regular basis and at known times. To build trust with customers, the company requires that subscription customers are consistently served by the same driver. In addition to subscription customer, on-demand customers request delivery on a daily basis. For the companies, the challenge is to consistently serve the subscription customers while simultaneously maximizing the daily profit from the on-demand customers. We model the problem as two-stage stochastic decision problem. The first stage determines the assignment of drivers to subscription customers. The second stage is a team-orienteering problem with time windows and mandatory visits by fixed drivers. We present an anticipatory consistent customer assignment policy (ACCA) based on the multiple scenario approach-framework. Our computational study shows that ACCA significantly outperforms consistency concepts from the literature while increasing costs less than 5%.

Keywords: Service Routing, Assignment Consistency, Stochastic Customers, Team Orienteering Problem, Multiple Scenario Approach

1 Introduction

The home-delivery market is becoming increasingly competitive, and companies are testing different business models to capture and keep customers. In this paper, we study a problem inspired by conversations with a large European grocery chain concerning their delivery operations. This company offers customers a grocery subscription. Customers who subscribe receive not only regular orders, but the delivery drivers enter customers' homes to place groceries into customers' refrigerators. This service requires trust between the customer and the driver, and one way to build trust is to have one driver consistently service a customer's home. Yet, at the same time, the company serves on-demand grocery customers who may or may not request service on any given day. Thus, there is the question of how to provide the customer-driver consistency needed for the subscription customers, but also to maintain the flexibility needed to serve as many on-demand customers as possible.

In this problem, the locations of subscription customers are known in advance, and each has a known hard time window in which service needs to take place. The subscription customers must be served by the same driver on each day. Because of previous requests, the locations of on-demand customers are known. However, we do not know whether or not an on-demand customer will request service on a given day. Further, because they are selected at the time an order is made, on-demand customers' time windows are unknown until they request service. We are not required to serve every on-demand customer, but serving on-demand customer earns revenue. Thus, each day, we must determine which on-demand customers to serve and how to route them given the constraint that we serve the subscription customers with the same driver each day. The objective is to minimize the expected difference between the routing costs and the revenue of on-demand customers. This objective allows us to avoid on-demand customers who are marginally unprofitable given that we are required to serve the subscription customers. We call the problem the two-stage stochastic assignment and team orienteering problem (TSATOP): the first stage fixes the assignments of the subscription customers to drivers; in the second stage, the on-demand customers that request service are selected and routed together with the subscription customers. The first stage problem is a stochastic assignment problem under the uncertainty of on-demand customers. The second stage problem is a deterministic team orienteering problem with time windows and mandatory visits.

The TSATOP is challenging for two reasons. First, the first stage assignment decisions are made under incomplete information. That is, at the time the subscription customers are assigned to specific drivers, we do not yet know what on-demand customers will be served. Second, decisions in the second stage require the solution of a deterministic team-orienteering problem with time windows and mandatory visits. We address the two challenges as follows. For the first stage, we develop an anticipatory consistent customer assignment approach (ACCA) based on the multiple scenario approach (MSA). To facilitate assignments, we introduce a new MSA consensus function that captures the needs of our stochastic assignment problem. For the second stage, we present a set partitioning formulation for solving the deterministic team-orienteering problem with time windows and mandatory visits by a branch-and-price algorithm.

In a comprehensive computational study, we compare our policy with four benchmark heuristics. For one, we compare our policy to a benchmark policy that does not enforce consistency. The solution values derived from this benchmark are a lower bound on the value of ACCA. Second, we split the fleet of vehicles and restrict the subscription customers to one subset of vehicles and the on-demand customer to the other. The final two policies are derived from the literature. As a third benchmark, we enforce a stricter consistency on the subscription customers in which we require the subscription customers to be served by the same driver everyday while also serving the subscription customers in the same relative order. We call this benchmark "master tours." Finally, we introduce a policy that seeks to distribute the subscription customers equally across the available drivers. We refer to this benchmark as "districting."

Using these benchmarks, we analyze the costs of consistency. We derive the following managerial insights:

1. Enforcing consistent assignments increases cost by less than 5% on average.

2. Splitting the fleet to focus on just subscription or just on-demand customers increases costs significantly, by less than 40% in our experiments, relative to just enforcing driver consistency.

3. Adding routing consistency by using master tours of the subscription customers increases the cost over just enforcing consistent driver assignment by about 15%. Given that fixed

sequences are easier for drivers to implement, the result offers companies insight into the additional cost of easing implementation and of potentially increasing driver satisfaction.

4. Enforcing driver consistency balances consistent service for subscription customers with the ability to serve on-demand customers. Relatedly, the results of the analysis also demonstrate that the use of "master tours," or sequence consistency, is most valuable when routing cost is the most important consideration. However, a districting strategy is more effective when the goal is to grow the business.

To derive these insights, we make the following methodological contributions. We introduce a new consensus function tailored to the proposed two-stage stochastic program. The consensus function introduced in Bent and Van Hentenryck (2004) focuses on routing. Our consensus function focuses on assignments and is therefore particularly suited for the problem at hand. This paper also demonstrates that branch-and-price is an effective method for exact solutions to the team orienteering problem with time windows and mandatory customers. Our proposed method may further be transferred to a variety of related business models. While this paper is focused on a particular problem emerging in grocery delivery, the driver-customer consistency is important in many industries that involve routing. Thus, this work has applications beyond that specifically discussed in the paper. For example, in parcel delivery, the consistent use of the same driver allows drivers to learn the area and the customer as well as specific drop-off locations (Wong 2008). A similar case has been made for the routing of drivers in vendor-managed inventory systems (Kovacs et al. 2014). Further, consistent assignments of field service personnel to locations allows the technicians to learn the facilities, have security clearance if necessary, and increase familiarity with the machines being repaired or maintained (Oldland 2017, Spliet and Dekker 2016). Personnel consistency is also important in home healthcare routing where patients and nurses need to develop trust (Lian 2017). Relatedly, customers receiving paratransit services identify driver consistency as an important factor in their satisfaction (Paquette et al. 2012).

The paper is outlined as follows. In Section 2, we present the related literature. In Section 3, we define the TSATOP. We present our method in Section 4. In Section 5, we describe the setup of experiments. The results of our experiments are displayed in Section 6. This paper ends with a conclusion and outlook in Section 7.

2 Literature Review

This literature review first reviews the literature in routing consistency, focusing on the literature that considers stochasticity. To explore the consistency routing questions of interest to this paper, we solve a variant of the team orienteering problem with time windows. As result, we also provide a brief review of related orienteering literature.

We also note that our problem is motivated by a business model in grocery delivery. Grocery delivery has received considerable attention in the literature, particularly with regard to time-slot allocation. In such problems, decisions need to be made regarding what time windows to offer to which customers and at which price. Examples of such work include Campbell and Savelsbergh (2006), Agatz et al. (2011), Ehmke and Campbell (2014), Klein et al. (2017), and Yang and Strauss (2017). None of these papers considers the consistency of customer assignments to drivers. To explore this question, we reduce the second-stage problem to the selection of on-demand customers whose time windows are known.

2.1 Consistency Routing

Explicit consideration of consistency is a recent development, and most existing papers study deterministic problems. A review of the consistency routing literature can be found in Kovacs et al. (2014). Recent papers addressing deterministic variants of consistency routing include Campelo et al. (to appear), Rodríguez-Martín et al. (to appear), Subramanyam and Gounaris (2017), and Xu and Cai (2018). Because of the random, on-demand customers, the problem in this paper is stochastic. There are only a small number of papers that address issues of consistency in the presence of uncertainty. In our review of the literature on consistency, we focus on only those papers that study consistency in routing in the presence of uncertainty.

We summarize the literature in Table 1. The table categorizes each paper by the type of consistency, how consistency is enforced, the stochastic component of the problem, and the structure of the first-stage solution. The categorization considers three types of consistency. In this paper, we focus on the consistency of assigning the same driver to the same customers or a subset of customers. We call this the *assignment consistency*. *Geographical consistency* occurs when a driver serves customers in the same geographical area each day. In addition, in some cases, it is important that the driver visit the customer at approximately the same time every visit. This type of consistency is called the *time consistency*.

Consistency is enforced in two ways. It is either implemented as a constraint when solving the day-to-day problem or it is enforced by valuing it in the objective. In our problem and as is common in many applications, time consistency is enforced through time windows.

The stochastic component of the problem can be characterized by customers, customer demand, or travel times. Finally, the table characterizes the structure of the first-stage solution. We use “districting” to refer to solutions that group customers based on proximity. We use “master tours” to refer to solutions that seek to maintain similarity to some pre-existing route. We use “assignment” to refer to methods that generate assignments without relying on proximity or a master tour.

The closest work to this paper is given by Spliet and Dekker (2016), which considers a problem in which one prefers that a driver consistently serve a customer because of the need to maintain customer specific equipment such as keys. In contrast to this paper though, Spliet and Dekker (2016) consider all customer locations to be known, but that demands are unknown. The idea is to assign these customers to a driver. The assignments are done via a clustering method that can be viewed as geographical districting. Because not all customers assigned to a driver can be served by that driver due to the random demand, and Spliet and Dekker (2016) constrain the percentage of the originally assigned customers who must be served by this primary driver in the daily routing problem. The customers not served by the primary driver are served by a backup driver.

Spliet and Dekker (2016) and this paper are related in the desire to have the same drivers serve the same customers, assignment consistency. However, we require subscription customers to be served by their assigned driver, but allow the on-demand customers to be served by anyone. On the other hand, Spliet and Dekker (2016) allows some deviation from the assignments by constraining the extent to which the customer assignment can differ over time. Further, from the perspective of the solution approach, we consider time windows as means of enforcing consistency in the delivery

Table 1: Literature Classification

Literature	**Consistency**		**Stochasticity**	**Solution**
	Type	Enforcement		
Haughton (2007)	**Assignment**	Objective	**Demand**	**Assignment**
Spliet and Dekker (2016)	**Assignment**	**Constraint**	**Customers**	Districting
Zhong et al. (2007)	Geographical	**Constraint**	**Customers**	Districting
Haugland et al. (2007)	Geographical	**Constraint**	**Customers**	Districting
Sungur et al. (2010)	Geographical	Objective	**Customers**	Master tours
Carlsson (2012)	Geographical	Objective	**Customers**	Districting
Carlsson and Delage (2013)	Geographical	Objective	**Customers**	Districting
Crainic et al. (2012)	Time	**Constraint**	Travel times	Master tours
Jabali et al. (2013)	Time	**Constraint**	Travel times	Master tours
Vareias et al. (2017)	Time	**Constraint**	Travel times	Master tours
Spliet and Gabor (2015)	Time	**Constraint**	Demand	**Assignment**
Spliet and Desaulniers (2015)	Time	**Constraint**	Demand	**Assignment**
Spliet et al. (2017)	Time	**Constraint**	Demand	**Assignment**
Dalmeijer and Spliet (2018)	Time	**Constraint**	Demand	**Assignment**
Our work	**Assignment**	**Constraint**	**Customers**	**Assignment**

time while Spliet and Dekker (2016) do not consider time consistency. As a result, the clustering approach used for assignments in Spliet and Dekker (2016) is not applicable to the problem studied in this paper.

Also related to our work is Haughton (2007) that seeks to maintain consistent assignments between customers and drivers in the case of stochastic customer demand. Instead of constraining the assignment, however, Haughton (2007) creates new routes for each day and then assigns drivers such that each route is assigned to the driver who has made the most previous visits to the customers on the route. The method proposed in Haughton (2007) cannot be applied to the problem studied in this paper because there is no way to guarantee that all subscription customers are served together.

Another stream of the consistency literature focuses on consistently serving customers at relatively the same time of day through the determination of time-window assignments. In these problems, customers have stochastic demands. The problem is modeled as having two stages. In the first stage, time windows are communicated before the customer demand is revealed. In the second stage, the demand is realized, and the dispatcher determines a set of routes that serve the customers in their TWs. The problem is introduced in Spliet and Gabor (2015), and Dalmeijer and Spliet (2018) introduce a new formulation leading to significant improvement in computational performance. Spliet and Desaulniers (2015) solve a discrete variant of the problem, and Spliet et al. (2017) extends the work to time-dependent travel times. The key difference between these time-window assignment problems and the problem studied here is that we focus on serving each of the subscription customers with the same driver and take the time windows as given.

Relatedly, Jabali et al. (2013) seek to find routes and according time windows when the set of customers is known and travel times are stochastic. The authors use “time buffers” to account for the uncertainty in travel times. Vareias et al. (2017) extend the existing work to include variable width time windows and to include multiple disruptions in the vehicles’ tours. A related time-window assignment problem is found in the master scheduling problem associated demand adaptive systems in which time windows must be assigned to the required stops (Crainic et al. 2012).

Related to the concept of assignment consistency is the idea that a driver’s performance improves as his or her familiarity with a particularly geographic area. Familiarity is similar to the concepts of consistency discussed in this paper, but it does not directly measure a driver’s relationship with a particular customer. Zhong et al. (2007) introduces driver familiarity and proposes a method for creating “districts” in which each driver operates. While not focused explicitly on consistency, Haugland et al. (2007), Carlsson (2012), and Carlsson and Delage (2013) also implement districting approaches for stochastic vehicle routing problems.

Sungur et al. (2010) note that these districts are difficult to operationalize on a daily basis where fluctuating daily demand requires customers are shifted from one district to another to balance demand. Sungur et al. (2010) suggest an alternative approach in which the region in which a driver is to gain familiarity defined by “master tours.” The solution approach uses the distance of a customer from a driver’s “master tours” to determine if the customer is “too far” from the driver’s region. Our computational study demonstrates that both the districting and master tours approaches are overly restrictive for the problem studied in this paper and lead to higher than necessary costs.

Finally, broadly speaking, a priori routing can be viewed as related to consistency routing. In a priori routing, the sequence of customers, who either have stochastic presence or stochastic demand is fixed in advance. Recourse actions are used to account for the randomness. In the case of stochastic presence, customers who do not need service on a given day are skipped. In the case of random demand, vehicles make return trips to the depot to replenish capacity. Either way, the recourse actions preserve both sequence as well as assignment consistency and the solution method can be categorized as master tours. Kovacs et al. (2014) provide a literature review and discuss the relationship between a priori routing and consistency. In this paper, it is not possible to use a priori routing as we do not know the time windows of the on-demand customers before they request service. However, this paper explores the cost of maintaining sequence consistency for the subscription customers and shows that sequence consistency can be substantially more expensive.

2.2 Orienteering Problems

The work in this paper requires the solution of two variants of the team orienteering problem with time windows (TOPTW). Generally, the team orienteering problem with time windows has been treated extensively in the literature. Gunawan et al. (2016) provide a comprehensive review. The second stage of our stochastic programming model is a specific variant of the TOPTW, the TOPTW with mandatory visits and fixed drivers for the subscription customers. To the best of the authors’ knowledge, the problem does not appear in the literature, and Lin and Vincent (2017) is the only paper to address the TOPTW with mandatory visits. Lin and Vincent (2017) solve the problem with multi-start simulated annealing, while we solve the problem exactly, demonstrating that branch-and-price is an effective method for medium-sized problems. Palomo-Martínez et al. (2017b), Palomo-Martínez et al. (2017a), and Lu et al. (2018) study a single traveler versions of the problem. Further, this second variant can be applied to the application presented in Salazar-Aguilar et al. (2014) in which a road maintenance department must route vehicles among mandatory and optional projects.

Given the stochastic nature of the problem studied in this paper, it is also related to stochastic orienteering. Particularly, Angelelli et al. (2017), Zhang et al. (2016), and Zhang et al. (2017) study a problem in which a single traveler serves a set of customers whose presence is unknown at the time that the tour is constructed. All three propose an a priori approach in which customers who are not present at the time of execution are skipped. The method presented in this paper could be adapted to a team version of the problem presented in Angelelli et al. (2017), Zhang et al. (2016), and Zhang et al. (2017).

3 Model

In this section, we present the model of the TSATOP. We first present a problem statement to give a problem narrative and to introduce the required notation. We then define the problem as two-stage stochastic orienteering problem with time windows.

3.1 Problem Statement

A service provider serves customers on a daily basis with a fleet of K vehicles starting and ending tours at a depot. The shift is limited by the working hours of the drivers. The drivers therefore need to return to depot no later than time $t_{\max}$.

The nodes for the depot and customers form the set V. The depot is represented by node 0. The customers and depot are part of a directed graph $G = (V, A)$ with vertex set V and arc set A. Each arc $(i, j) \in A$ has an associated travel time t_{ij} and travel cost c_{ij}. The customers $V_0 = V \setminus \{0\}$ can be split into two subsets $V_0 = V_S \cup V_D$. Set V_S represents the set of *subscription customers* who must be served regularly, e.g., every day, by the same vehicle. On the other hand, set V_D represents the set of all *on-demand customers*, who randomly request services, and therefore only a random subset of them may actually request services on any particular day. The subset of on-demand customers who requests service on a given day is denoted by $\tilde{V}_D \subseteq V_D$. The request probability of a customer is known. Each subscription customer $i \in V_S$ is associated with a fixed time window $[\ell_i, u_i]$, within which the service at this customer must start, whereas the time window associated with each on-demand customer $i \in \tilde{V}_D$ is assumed to be random $[\tilde{\ell}_i, \tilde{u}_i]$. Each on-demand customer $i \in \tilde{V}_D$ can be optionally served with a revenue ρ_i. We assume that the subset of on-demand customers $\tilde{V}_D$ for a day becomes known at the beginning of the day. Each customer $i \in V_0$ (both subscription and on-demand customers) has an associated service time τ_i, which for simplicity is assumed to be deterministic.

There are two stages of decision making. The first stage occurs initially, in the beginning of the problem and is tactical. The second stage occurs on a daily basis and is operational. In the first stage, the provider needs to decide how to assign the subscription customers to the drivers. These subscription customers are then served every day by the same assigned driver. An assignment decision is feasible if there exists at least one set of routes that serves all subscription customers using the assigned drivers and does so within the customers' time windows. In the second stage, decisions need to be made about the set of on-demand customers to serve, the assignment of the

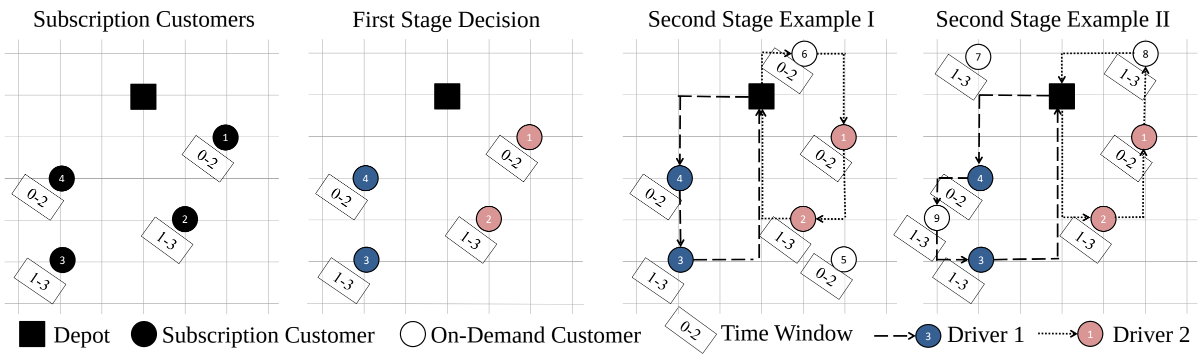

Figure 1: Example for a First Stage Decision and Two Second Stage Decisions for Different Realizations.

on-demand customers to the vehicles, and the routing through the subscription and selected on-demand customers.

The objective is to minimize the expected costs minus the revenue from the chosen on-demand customers. Because the revenue for the subscription customers is constant, we omit it from model and calculation. As a result, the costs can be negative.

3.2 Example

In the following, we give a very simple example to illustrate the two stages of decision making. The example is depicted in Figure 1. It contains four subfigures. From left to right, the subfigures show the set of subscription customers, a potential first stage decision, and two potential second stage decisions, each for a different realization. Each subfigure depicts the depot as the black square and a street network. For the purpose of presentation, we use a Manhattan-grid in the example. Each segment has a travel time of 15 minutes. For simplicity, we omit service times from the example. The time limit is 3 hours.

On the left side of Figure 1, four subscription customers are shown, identified by the black circles. Each customer has a time window that is shown in the adjacent box. For the example, we assume that we have two drivers. In the first stage, a decision needs to be made about the assignment of the subscription customers to the two drivers. A potential assignment decision is shown in the second subfigure from the left. The dark blue circles indicate an assignment to Driver 1. The light red circles indicate an assignment to Driver 2. In the example, Customers 3 and 4 are assigned to Driver 1 and Customers 1 and 2 are assigned to Driver 2. This assignment is feasible because there exists a feasible routing, for example, Depot-4-3-Depot for Driver 1 and Depot-1-2-Depot for Driver 2. In both cases, the four time windows are met, and the vehicles return to the depot within the time limit.

The thirdand fourth subfigures present two potential realizations of the second stage and two corresponding second stage decisions. The first realization provides two on-demand customer requests by Customers 5 and 6. The on-demand customers are indicated by white circles. The decision is to accept Customer 6 and to set the routes for the drivers to Depot-4-3-Depot and to Depot-6-1-2-Depot. The second realization provides requests of Customers 7 and 8. The decision is to accept Customer 8 as well as routes Depot-4-3-Depot and Depot-2-1-Depot. We observe that in the routing of the two realizations, the orders of Customers 1 and 2 are flipped.

3.3 A Two-stage Stochastic Programming Model

Formally, we can model this problem as a two-stage stochastic program. The first-stage decision is to determine the assignment of subscription customers to the drivers. We model this decision with variable $\mathbf{y_S}$. This variable is an assignment matrix with entries $\{y_i^k\}_{k\in[1:K]}^{i\in V_S}$. An entry y_i^k is 1 if subscription customer i is assigned to vehicle k and 0 otherwise.

The second-stage decision follows the realization of on-demand customers $\tilde{V}_D$. In this state, we determine:

(1) For each on-demand customer who has made a request, whether or not we choose to serve them;

(2) For those on-demand customers that we decide to serve, how we use our fleet to serve them, while respecting the fixed assignments of subscription customers to vehicles, as well as the time window constraints for all customers.

Thus, we have the following two-stage stochastic integer program:

$$\min \mathbb{E}\left[\tilde{F}\left(\tilde{V}_D, \mathbf{y_S}\right)\right] \tag{1a}$$

$$\text{s.t.} \sum_{k=1}^{K} y_i^k = 1,\ \forall i \in V_S \tag{1b}$$

$$y_i^k \in \{0,1\},\ \forall k \in [1:K],\ \forall i \in V_S, \tag{1c}$$

where $[m\colon n]$ is a shorthanded notation for $\{m, m+1, \ldots, n\}$, and $\mathbf{y_S}$ is a shorthanded notation for $\{y_i^k\}_{k\in[1:K]}^{i\in V_S}$. Function $\tilde{F}\left(\tilde{V}_D, \mathbf{y_S}\right)$ is the second-stage cost given an assignment of subscription customers $\mathbf{y_S}$ and given any realization $\tilde{V}_D$ of the on-demand customers and their associated time windows. In the second stage, given a first-stage decision, $\hat{\mathbf{y_S}}$, these realizations result in an orienteering problem. To model this problem, we use a set partitioning formulation: given a realization of on-demand customers $\tilde{V}_D$, we also input the set of all feasible routes for vehicle k that visit any subset of $\tilde{V}_D$ (including the empty set). We denote this set as $\Omega^k(\tilde{V}_D, \hat{\mathbf{y_S}})$ for each vehicle $k \in [1:K]$. A feasible route r is a route through the customers associated with r such that the time windows of those customers are met. Further, we include $\hat{\mathbf{y_S}}$ as an argument to Ω^k to indicate that feasible routes for vehicle k must visit the subscription customers that are assigned to this vehicle. We let α_{ir}^k be an indicator that identifies customer $i \in \tilde{V}_D$ as belonging to route r associated with vehicle k. Every route r for vehicle k is associated with a cost c_r^k. As discussed previously, these costs include the cost of traversing route r minus the revenue gained in serving the on-demand customers associated with r.

The second-stage decision variable is $z_r^k \in \{0,1\}$ for each vehicle k and each route $r \in \Omega^k(\tilde{V}_D, \hat{\mathbf{y_S}})$, where z_r^k is 1 if route r is selected for vehicle k and 0 otherwise. The set partitioning formulation is defined as follows:

$$\tilde{F}\left(\tilde{V}_D, \hat{\mathbf{y}}_{\mathbf{S}}\right) := \min \sum_{k=1}^{K} \sum_{r \in \Omega^k(\tilde{V}_D, \hat{\mathbf{y}}_{\mathbf{S}})} c_r^k z_r^k \tag{2a}$$

$$\text{s.t.} \sum_{r \in \Omega^k(\tilde{V}_D, \hat{\mathbf{y}}_{\mathbf{S}})} z_r^k = 1, \, \forall k \in [1:K] \tag{2b}$$

$$\sum_{k=1}^{K} \sum_{r \in \Omega^k(\tilde{V}_D, \hat{\mathbf{y}}_{\mathbf{S}})} \alpha_{ir}^k z_r^k \leq 1, \, \forall i \in \tilde{V}_D \tag{2c}$$

$$z_r^k \in \{0, 1\}. \tag{2d}$$

The constraints require that each vehicle k be assigned a feasible route from $\Omega^k(\tilde{V}_D, \hat{\mathbf{y}}_{\mathbf{S}})$, that each realized on-demand customer is served at most once, and that all the decision variables are binary.

4 Anticipatory Consistent Customer Assignments

In this section, we describe our method, the *anticipatory consistent customer assignment policy* (ACCA). We first give a general overview over our method. We then describe how we solve the second stage of the problem and finally how the solutions of the second stage are used to derive a first-stage solution.

4.1 Overview

Before we present methodological details, we first motivate and give a general overview of our method. ACCA consists of two steps. In the first step, we sample sets of on-demand customers. We refer to these sampled sets of customers as *scenarios*. For each scenario, we then solve an independent team orienteering problem for the subscription customers and the sampled set of on-demand customers. These independent team orienteering problems have no requirements regarding the consistency of subscription customers. In the second step, we use information from the individual solutions to determine the assignment decisions of the subscription customers. For the TSATOP, the first step is associated with the second stage, and second step is associated with the first stage.

ACCA is based on the multiple-scenario approach (MSA) framework introduced in Bent and Van Hentenryck (2004). Methods of this framework seek to solve stochastic optimization problems, often routing problems, by solving a series of deterministic problems generated via sampling. The solutions of the individual problems are then used to determine a solution to the stochastic problem, often searching for the solution "most similar" to the other solutions. The motivation behind this selection is that the selected solution is likely able to perform well for different realizations.

The procedure is known as scenario analysis (SCA) in the stochastic programming literature. Methods of SCA extract structures from solutions of deterministic scenarios to guide decision making for the stochastic problem. SCA has experienced broad attention in the stochastic vehicle routing literature, particular with the MSA (Azi et al. 2012, Bent and Van Hentenryck 2004, Ghiani et al. 2012, Hvattum et al. 2006, 2007, Voccia et al. to appear). However, it is not guaranteed that the deterministic structures necessarily hold in the stochastic context (King and Wallace 2012). The SCA methods in stochastic vehicle routing generally extract routing information. We apply such a method as a benchmark heuristic. ACCA utilizes different information, the assignments, for the first-stage solution. Thus, while primarily comparing different heuristics, this paper also sheds light on the qualities of scenario analysis for different information extracted.

While the MSA is a general framework, the individual steps are tailored to the application. In our case, we implement a branch-and-price approach to the team orienteering problem. Then, most importantly, we introduce a new consensus function, the method for combining the individual solutions from the first step into assignments for the subscription customers. This new consensus function is specific to the needs of the assignment problem found in the first stage of our stochastic program.

4.2 Solving the Deterministic Team Orienteering Problem

Because we solve the team orienteering problem for multiple scenarios, our approach depends on an efficient solution method. The set partitioning formulation solved by a branch-and-price algorithms is well-known to be one of the most powerful exact solution approaches for solving deterministic vehicle routing problems and team orienteering problems (Gunawan et al. 2016).

Our formulation of the team orienteering problem includes constraints similar to those found in the second-stage of the TSATOP. However, in solving the orienteering problem associated with each scenario, we do not consider consistency, but solve independent team-orienteering problems for each scenario. Thus, for each scenario $\tilde{V}_D$ of on-demand customers, we relax the route set $\Omega^k(\tilde{V}_D, \hat{\mathbf{y}}_\mathbf{S})$ to a route set $\Omega(\tilde{V}_D)$. The route set $\Omega(\tilde{V}_D)$ does not enforce any assignments related to the subscription customers and all feasible routes through subsets of $V_S \cup \tilde{V}_D \cup \{0\}$.

With the new definition of the set $\Omega(\tilde{V}_D)$, we then redefine the α_{ir}^k parameter and have α_{ir} as an indicator that identifies customer $i \in V_S \cup \tilde{V}_D$ as belonging to route r. Further the route costs can now be defined independent of vehicle k as c_r.

To facilitate branching, we also add a set of redundant constraints to the model. For that purpose, we introduce $e_{ar} \in \{0, 1\}$ to indicate whether route r traverses arc a. Then, x_a is a binary decision variable indicating that arc a is included in the solution. With the addition of x_a, we can then also relax the integrality of z_r^k.

We have the following integer program:

$$\min \sum_{k=1}^{K} \sum_{r \in \Omega(\tilde{V}_D)} c_r z_r^k = 1 \tag{3a}$$

$$\text{s.t.} \sum_{r \in \Omega(\tilde{V}_D)} z_r^k = 1, \forall k \in [1:K] \tag{3b}$$

$$\sum_{k=1}^{K} \sum_{r \in \Omega(\tilde{V}_D)} \alpha_{ir} z_r^k = 1, \forall i \in V_S \tag{3c}$$

$$x_a = \sum_{k=1}^{K} \sum_{r \in \Omega(\tilde{V}_D)} e_{ar} z_r^k, \ \forall a \in A \tag{3d}$$

$$\sum_{a \in \delta^+(i)} x_a = 1, \ \forall i \in V_S \tag{3e}$$

$$\sum_{a \in \delta^+(i)} x_a \leq 1, \ \forall i \in \tilde{V}_D \tag{3f}$$

$$\sum_{a \in \delta^-(i)} x_a = \sum_{a \in \delta^+(i)} x_a, \ \forall i \in V_S \cup \tilde{V}_D \cup \{0\} \tag{3g}$$

$$z_r^k \in \mathbb{R}_+, \ \forall r \in \Omega(\tilde{V}_D), \forall k \in [1:K] \tag{3h}$$

$$x_a \in \{0,1\}, \ \forall a \in A. \tag{3i}$$

Constraints (3b) enforce that each vehicle k is assigned to only one route, and Constraints (3c) enforce that each subscription customer is served and served at most once. Constraints (3d) activate x_a when a route containing arc a is selected. Constraints (3e) require that an arc enter each subscription customer, and Constraints (3f) enforce that each sampled on-demand customer appear in at most one selected route. Constraints (3g) enforce that any arc entering a node also exit a node. The remaining constraints bound the decision variables.

Formulation (3) contains exponentially many variables z_r^k, motivating a branch-and-price algorithm to solve this formulation: at each node of the branch-and-bound tree, the linear programming relaxation problem needs to be solved by column generation. Given dual multipliers δ_i for Constraints (3c), and σ_a for Constraints (3d), the corresponding pricing problem for column generation can be formulated as:

$$\min_{r \in \Omega} \ \bar{c}_r := c_r - \sum_{i \in V_S} \alpha_{ir} \delta_i - \sum_{a \in A} e_{ar} \sigma_a \tag{4}$$

The pricing problem (4) can be solved via an extension of a standard labeling algorithm for solving VRPTW, by incorporating the profits collected from on-demand customers visited along each route. We summarize the labeling algorithm in Algorithm 1 as follows. In particular, each label L is a partial path, and is associated with: (i) a set of vertices visited by the label $M(L)$; (ii) cumulative reduced cost $c(L)$; and (iii) earliest possible service start time at the last vertex $t(L)$. We say that a label L^2 is dominated by another label L^1, if: (i) L^1 and L^2 end at the same vertex; (ii) $M(L^1) \subseteq M(L^2)$; (iii) $t(L^1) \leq t(L^2)$, $c(L^1) \leq c(L^2)$, and at least one of these two inequalities is strict.

Algorithm 1: The labeling algorithm.

<u>**Label Initialization**</u>**.** We maintain a list of labels $\mathcal{L}(i)$ for each customer $i \in V_0$. $\mathcal{L}(i) \leftarrow \emptyset$ and $\delta \leftarrow 1$.

foreach $i \in V_0$ **do**

  Initialize labels $L^i = ((0,i)), c_{L^i} = \bar{c}_{0i}$ if $i \in V_S$ and $c_{L^i} = \bar{c}_{0i} + \rho_i$ if $i \in \tilde{V}_D$, and $t_{L^i} = min\{t_{0i}, \ell_i\}$. Add label L^i to $\mathcal{L}(i)$.

end

<u>**Label Extension**</u>**.**

repeat

  foreach $j \in V_0$ **do**

    foreach $L = (P)$ *from* $\mathcal{L}(j)$ *where* P *contains* δ *vertices* **do**

      foreach $i \in V_0$ **do**

        if *customer* i *has not been visited by* P **then**

          if $t_L + t_{ji} < u_i$ **then**

            Create labels $L^n = (P \oplus (j,i))$, $c_{L^n} = c_L + \bar{c}_{ji} + \rho_i$ if $i \in \tilde{V}_D$ and $c_{L^n} = c_L + \bar{c}_{ji}$ if $i \in V_S$. Set $t_{L^n} = min\{t_L + t_{ji}, \ell_i\}$.

            if L^n *is not dominated* **then**

              Add L^n to $\mathcal{L}(i)$

            end

          end

        end

      end

    end

  end

  $\delta \leftarrow \delta + 1$.

until $\delta > |V_0|$;

<u>**Label Termination**</u>**.**

foreach $i \in V_0$ **do**

  foreach $L = (P)$ *from* $\mathcal{L}(i)$ **do**

    Extend label L to the depot and calculate the optimal reduced cost.

  end

  Output label L from $\mathcal{L}(i)$ with a negative reduced cost.

end

4.3 Measuring Similarity based on Hamming Distances

Given a solution to Formulation (3), we can extract the assignments of the subscription customers implied by the solution. We let $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}$ be the assignments associated with scenario $s \in \mathcal{S}$, where $\mathcal{S}$ is the set of all scenarios. We then seek to find the $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}$ that is "most similar" to the other solutions. In the following, we recall the concept of "consensus functions" to identify a solution from a set of candidates. We further define our consensus function and give an example of its functionality.

4.3.1 Consensus Function

We define a measure quantifying the differences between two assignments variables $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}$ and $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}'}$, corresponding to scenarios s and s'. This measure is called the "consensus function" in Bent and Van Hentenryck (2004). The consensus function introduced in Bent and Van Hentenryck (2004) compares the first customers that the vehicles are scheduled to visit in their route. Because for our problem a solution is not a route but an assignment decision, this consensus function does not apply. However, we use a generalization of the original consensus function as a benchmark.

We propose a consensus function based on the Hamming distance between different scenario-based solutions. The Hamming distance measures the difference in vectors or matrices by comparing pairs of individual entries. Specifically, given the solutions for two scenarios s and s', $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}$ and $\mathbf{y}_{\mathbf{S}}^{\mathbf{s}'}$, the Hamming distance between the two solutions is given by:

$$H(\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}, \mathbf{y}_{\mathbf{S}}^{\mathbf{s}'}) := \sum_{k=1}^{K} \sum_{i \in V_S} |(\mathbf{y}_{\mathbf{S}}^{\mathbf{s}})_i^k - (\mathbf{y}_{\mathbf{S}}^{\mathbf{s}'})_i^k|$$

Then the candidate solution $\mathbf{y}_{\mathbf{S}}$, is chosen as the scenario-based solution that is closest to other scenario-based solutions in terms of the Hamming distance, i.e., the scenario-based solution $\mathbf{y}_{\mathbf{R}}^{\bar{\mathbf{s}}}$ corresponding to scenario $\bar{s}$, where

$$\bar{s} \in \arg\min_{s \in \mathcal{S}} \left\{ \sum_{s' \in S, s' \neq s} H(\mathbf{y}_{\mathbf{S}}^{\mathbf{s}}, \mathbf{y}_{\mathbf{S}}^{\mathbf{s}'}) \right\}. \tag{5}$$

4.3.2 Example

In the following, we give an example to illustrate the process. The example comprises 8 subscription customers and three vehicles. We assume that our second stage created three different solutions for scenarios $s1, s2$, and $s3$. The solution for the first scenario $s1$ assigns Customers 1, 2, and 3 to Vehicle 1, Customers 4, 5, and 6 to Vehicle 2, and Customers 7 and 8 to Vehicle 3. The assignment matrix is therefore

$$\mathbf{y}_{\mathbf{S}}^{\mathbf{s1}} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}.$$

The solution for $s2$ assigns Customers 1, 2, and 3 to Vehicle 1, Customers 4, 6, and 7 to Vehicle 2, and Customers 5 and 8 to Vehicle 3:

$$\mathbf{y_S^{s2}} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}.$$

The solution for $s3$ assigns Customers 1, 2, and 4 to Vehicle 1, Customers 3, 6, and 7 to Vehicle 2, and Customers 5, 8 to Vehicle 3:

$$\mathbf{y_S^{s3}} = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}.$$

For these three solutions, the distances are $H(s1, s2) = 4$, $H(s1, s3) = 8$, and $H(s2, s3) = 4$. In that case, the solution for $s1$ has a value of 12 in Equation 5, the solution for $s2$ has a value of 8, and the solution for $s3$ has a value of 12. Thus, the solution for $s2$ is the "most similar" solution to all others and is therefore selected by the algorithm as the consensus solution.

4.3.3 Symmetry Identification

One challenge in our procedure is that solution symmetries may lead to high differences even though the solutions are very similar. As an example, we take the solution for $s2$ assigning Customers 1, 2, and 3 to Vehicle 1, Customers 4, 6, and 7 to Vehicle 2, and Customers 5 and 8 to Vehicle 3. We now have another scenario $s4$. This solution for $s4$ assigns Customers 4, 6, and 7 to Vehicle 1, Customers 5 and 8 to Vehicles 2, and Customers 1, 2, and 3 to Vehicle 3:

$$\mathbf{y_S^{s4}} = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

The two solutions result in the same assignment. However, the Hamming distance between the solutions of $s2$ and $s4$ is maximal with $H(s2, s4) = 48$ because no entry is the same. Thus, the consensus function may discard both solutions even though they are identical. Thus, measures are needed to identify symmetries in the solutions.

The brute force way to identify symmetry is by enumerating all orders of the rows in one of the solution matrices and searching for the setting with minimal distance. However, the number of row-orders for the comparison of two solutions is $K!$ with K being the number of vehicles. Given $|\mathcal{S}|$ scenarios, this leads to about $K!|\mathcal{S}|^2$ comparisons for one instance setting. Thus, we draw on a heuristic to reduce the computational burden.

Our heuristic follows the idea to first sort the vehicles by number of customers and, if the numbers are the same, sort the vehicles by the lexicographic order using the smallest index of the assigned customers. In the example, both solutions would be altered to

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

The new "first" vehicle would be Vehicle 3 in the solution of $s2$ and Vehicle 2 in the solution of $s4$, both serving Customers 5 and 8. The new second vehicle would be Vehicle 1 for the solution of $s2$

and Vehicle 3 for the solution of $s4$ serving Customers 1, 2, and 3. Finally, the new third vehicle would be Vehicle 2 for the solution of $s2$ and Vehicle 1 for the solution of $s4$ serving Customers 4, 6, and 7. The Hamming distance between these two new solutions is zero. Thus, our heuristic is able to identify the symmetry in the solutions of $s2$ and $s4$.

5 Design of Experiments

In this section, we describe the design of experiments. We first describe our instances. We then present the benchmark policies and finally give implemenation details.

5.1 Instance Generation

To analyze the different dimensions of our problem, we test our method for a variety of instances. In generating the instances, we vary the percentages of subscription and on-demand customers, the customer distribution, and the number of vehicles.

We create instances with $|V_S| = 20$ subscription and $|V_D| = 10$ potential on-demand customers, 15 subscription and 15 potential on-demand customers, and 10 subscription and 20 potential on-demand customers. We vary the geography of customers drawing on the well-known Solomon's VRPTW instances (Solomon 1987). We select 10 Solomon geographies: R101 to R105 and C101 to C105. Each of the geographies has 100 potential locations. Because each instance has 30 potential locations, we sample 30 locations from each geography. The first $|V_R|$ locations are the locations of the subscription customers. The remaining locations are the on-demand customers. We repeat this procedure five times for each geography. This leads to 50 different customer distributions.

In our instance generation, we focus on delivery in the afternoon between 4pm and 8pm. Thus, we set the time limit to $t_{\text{max}} = 240$ minutes. The time window sizes are set to 60 minutes. For each customer distribution, we randomly generate the time windows for the subscription customers by sampling the beginning of the time window from the uniform distribution $U(0, 180)$. These time windows are known on the first stage of our problem.

For each customer distribution, we next generate 350 sampled second stage realizations of on-demand customers. We use 50 realizations as in-samples within our algorithm. The remaining 300 realizations are used as out-of-sample realizations for the evaluation of solutions provided by our algorithm. For each realization, the subset of on-demand customers and their corresponding time windows is generated. Each potential on-demand customer requests service with probability of 0.7. We choose this value based on the instance generation in (Sungur et al. 2010). The beginning of the 60-minute time windows is sampled from the uniform distribution $U(30, 180)$ where the start of the distribution accounts for the fact that the customers become known with little notice.

We test instances with fleet sizes of 3 and 5 vehicles. The vehicles travel on Euclidean paths between depot and/or customers. Based on preliminary tests, we set the travel speed to two units per minute. We set the revenue per on-demand customers to be equivalent to 10 minutes of travel time. We set the service time of a customer to 5 minutes.

Overall the combination of subscription percentages, customer distributions, and number of vehicles leads to $3 \times 50 \times 2 = 300$ different instances. Thus, we solve overall 15,000 in-sample orienteering problems and evaluate the resulting solutions using 90,000 out-sample orienteering problems.

5.2 Benchmark Policies

We create four benchmark policies to evaluate the performance of our algorithm and to generate managerial insight. We include the two most prominent concepts from the literature, master tours and districting of the service area. We also analyze a policy splitting the fleet. We finally present a policy without consistency constraints to measure the costs of consistency. In the following, we describe the functionality of these four policies.

Master Tours.

As our literature review discusses, some methods use master tours. These master tours are then adapted to daily demand. ACCA differs because we determine only the assignment, but allow flexibility in the routing on the second stage. To analyze the benefit of this flexibility, we apply a method similar to our policy, but with a fixed sequence of subscription customers. We denote this policy "Master Tours." To generate solutions for this method, we follow the steps of ACCA, but we change the consensus function to account for sequences. Preliminary tests with the consensus function of Bent and Van Hentenryck (2004) provided inferior results because their consensus function considers only the first customer in each tour. Thus, we generalize the idea of Bent and Van Hentenryck (2004).

To do so, we remove all on-demand customers from a solution, to obtain a sequence of subscription customers. We then use the Hamming distance to calculate the number of different edges used in two solutions and select the "most similar" sequence. Similar to ACCA, we use preprocessing steps to identify symmetries. In the second stage, on-demand customers are then inserted in the sequence of subscription customers.

Districting.

As the survey by Kovacs et al. (2014) indicates, another way of ensuring consistency is by grouping customers to districts of approximately the same size. Drivers are then assigned to districts. This leaves every driver some slack to serve additional customers in this district. Another advantage of this procedure is that no master tours are needed and flexibility in daily routing is maintained. We denote this policy "Districting." For our problem, districting only on geography is challenging because of the customers' time windows. Thus, we present a method that districts customers with respect to both geography and time windows. Districting distributes subscription customers "relatively" equally to the drivers in the fleet. This distribution is generated by minimizing the overall routing costs under the condition that the number of customers per vehicle should not differ by more than one customer. For this policy, we solve the well-known VRPTW for all the subscription customers and with the additional constraint that numbers of customers per vehicle are only allowed to differ by one.

Split.

The easiest way to ensure consistency is by splitting the fleet. Vehicles serve subscription or on-demand customers but not both. Thus, vehicles serving subscription customers perform the same routes every day, while the remaining percentage of vehicles is free to serve daily on-demand customers. We denote this policy "Split." To determine the number of vehicles and the corresponding routes, for the subscription customers, we consider the case of allocating $1, 2, \ldots, K$ vehicles to serve the subscription customers. For each possible allocation of vehicles to the subscription customers, we solve a deterministic team orienteering problem for the remaining "on-demand" vehicles and the on-demand customers for each in-sample scenario. We then pick the optimal number of subscription vehicles based on the number leading to the least expected costs. The out-of-sample evaluation for a candidate solution can be done by ignoring the subscription customers and solving the deterministic orienteering problem with the remaining on-demand vehicles separately for each out-of-sample scenario .

No Consistency.

To analyze the costs of consistency, we finally implement a policy that entirely ignores consistency. This policy solves the deterministic team orienteering problem with mandatory visits every day without consistency constraints. Thus, subscription customers can be served by any vehicle of the fleet. This policy provides a lower bound for the costs of our policies. We denote this policy "No Consistency."

5.3 Implementation Details

To solve various variants of deterministic team orienteering problems encountered during the proposed solution framework, we use the SCIP 4.0.1 with CPLEX 12.5.1 as the linear programming solver to implement the branch-cut-and-price framework in C++. All tests are conducted on a Linux workstation with four 3.00GHz processors and 8Gb memory. The number of threads is set to be one. In in- and out-of-sample evaluations, we solve each deterministic team orienteering problem to optimality. All instances terminate at optimality within 300 seconds.

6 Computational Analysis

In this section, we present our computational results. We first compare the solution qualities of the policies. We then analyze the impact on routing costs and customer services. Finally, we analyze how growing our subscriber base impacts the costs of consistency.

6.1 Solution Quality

We first analyze the solution quality of our policies. In this section, we present aggregated results. For the individual results, we refer to the Appendix of the paper. To get a better estimate of the differences between the policies, we also calculate the gap between a policy and the No Consistency policy for each instance as

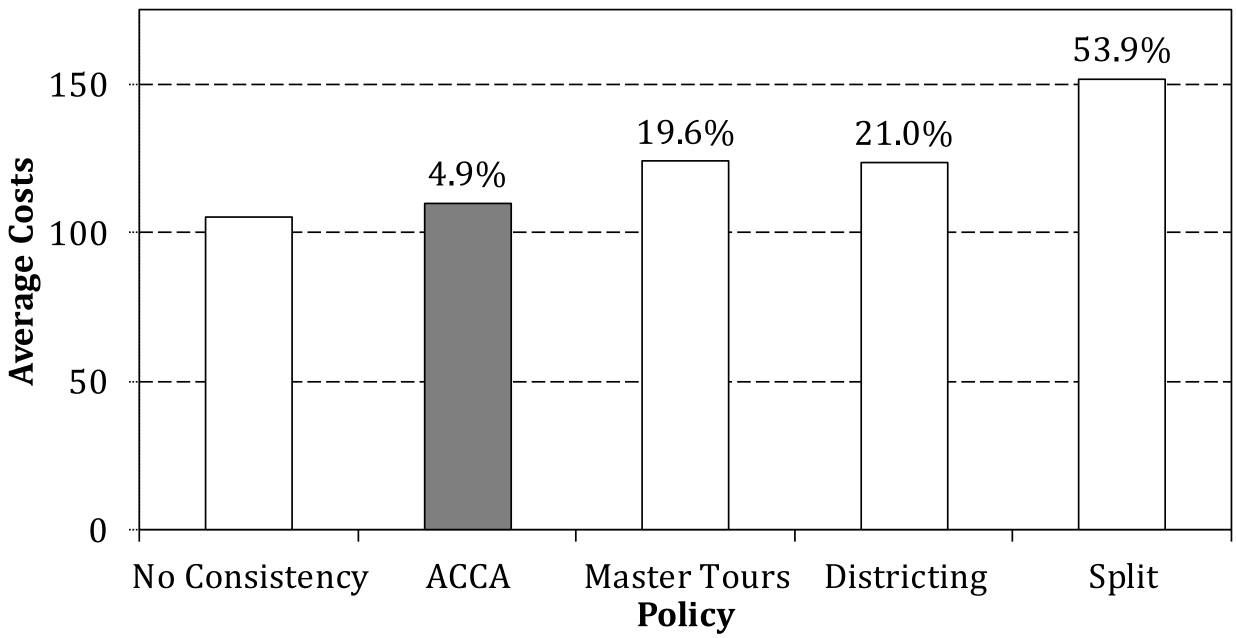

Figure 2: Average Costs and Average Gap to No Consistency over all 300 Instance Settings.

$$\frac{\text{Policy} - \text{No Consistency}}{\text{No Consistency}}. \tag{6}$$

We then calculate the average over all gaps of the individual instances.

Figure 2 shows the average out-of-sample costs over all 300 instances and 300 realization per instance as well as the average gap between a policy and the No Consistency policy. The x-axis shows the five policies, the y-axis depicts the average costs. The values at the bars indicate the average gap.

As expected, No Consistency performs best resulting in the lowest costs on average as well as for all individual instances. This method exploits the substantial flexibility in daily routing without the consistency constraint. It is therefore a lower bound. Notably, ACCA performs only slightly worse than No Consistency and significantly better than the other policies considering consistency.

The gaps are in accordance to the absolute values of Figure 2. The gap for ACCA is relatively narrow with less than 5%. The gap for Master Tours and Districting is large with about 20% higher costs on average. Policies Master Tours and Districting perform similarly with respect to the objective function. However, as we will show later, they provide very different solutions. The gap for Split is tremendous with more than 50% higher costs. This result suggests that splitting the fleet is an inefficient means of achieving consistency. The advantages of ACCA are also reflected in the individual results. For 277 of the 300 individual instances, ACCA provides the best result except of the No Consistency policy. On average, ACCA improves upon policy Master Tours by 14.1%, upon Districting by 15.3%, and upon Split by 46.4%.

The significant gap between ACCA ans Master Tours also gives insight in the performance of scenario analysis (SCA) for stochastic vehicle routing problems. As discussed in Section 4.1, both ACCA and the Master Tour rely on SCA but extract different information to generate the first-stage solution. The significant difference between the performances of the two approaches indicates that

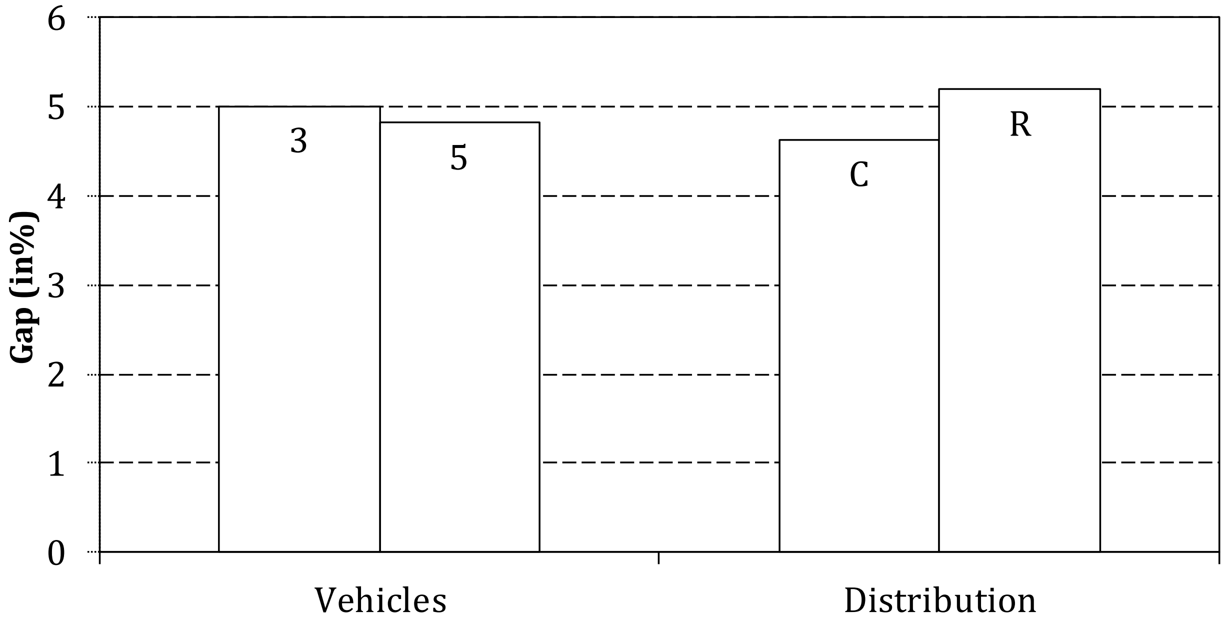

Figure 3: Average Gap to No Consistency for Instances Grouped by Number of Vehicles and Distribution.

transferring only the assignment structure from the scenarios may be more suited for SCA than extracting the routing structures.

On average, with ACCA, consistency can be achieved with an increase in costs of less than 5%. As the individual results in the Appendix show, these values vary for the individual instances. However, they do not change substantially if we group instances with respect to distributions or number of vehicles as we show in Figure 3. On the left, the average gaps for the 150 instances with 3 vehicles and the 150 instances with 5 vehicles are shown. On the right, average gaps for the 150 instances with clustered customer locations and for the 150 instances with uniformly distribution customer locations are shown. The gaps differ only marginally between 4.6% and 5.2%. Thus, the customer distribution and the fleet size do not impact the costs of consistency substantially. The cost of consistency changes with the percentage of subscription customers though, as we will show later in the paper.

6.2 Balancing Services and Routing Costs

For our problem, we experience a tradeoff in routing costs and on-demand customers served. We address this tradeoff in our model by incorporating a reward generated by the customers in the objective function. However, solutions with similar objective values may differ significantly in routing costs and customers served. This difference is important because operational decisions on routing and customer services may have longer term impacts on the business model. Routing costs may impact the pricing of goods and services while good customer service may increase our customer base, for example, by turning on-demand customers into subscription customers. Thus, it is important to analyze how the policies perform with respect to the individual components of the model's objective. In the following, we analyze the average routing costs and services for our

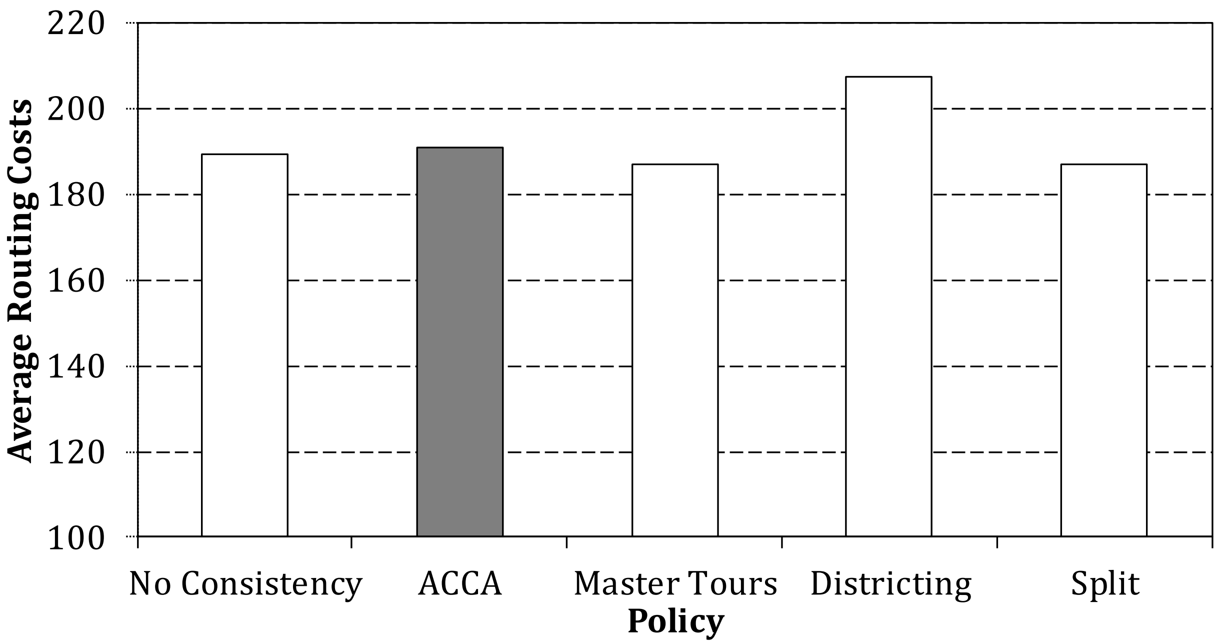

Figure 4: Average Routing Costs.

policies. To do so, we take the solutions obtained from our policies and analyze each component of the objective function individually.

We start with the average routing costs. These are depicted in Figure 4. We observe that policies No Consistency and ACCA show similar routing costs. Policies Master Tours and Split have slightly lower costs. The routing costs for the Districting policy are significantly higher. The equal distribution of customers to vehicles comes at the costs of longer routes and higher costs.

The second components of the objective is to serve many on-demand customers. To analyze the results for our policies, we calculate the average on-demand services, that is the percentage of on-demand customers served on average over all instances. The results are shown in Figure 5.

We observe that polices No Consistency, ACCA, and Districting show similar percentages of services. Policies Master Tours and Split serve significantly fewer customers. Both policies allow only limited flexibility due to the predefined sequence in customer visits and the exclusive use of vehicles, respectively. The Districting policy allows many services. For this policy, the assignments are made in a way that a better coverage of the service area is enforced. This coverage allows the efficient integration of on-demand customers and a high service level.

The comparison between Master Tours and Districting provides some interesting insight in how planning should be conducted if one objective may be more relevant than the other. If the main objective is saving money by routing the subscription customers at low costs, creating master tours may be a reasonable approach. If the main objective is to grow market share by serving many on-demand customers, a suitable coverage of the service area as provided by Districting may be the way to go. However, ACCA is able to balance the tradeoff between costs and services. It allows both a relatively high service level as Districting and relatively low routing costs as Master Tours. It exploits the strengths of Master Tours and Districting by avoiding their weaknesses.

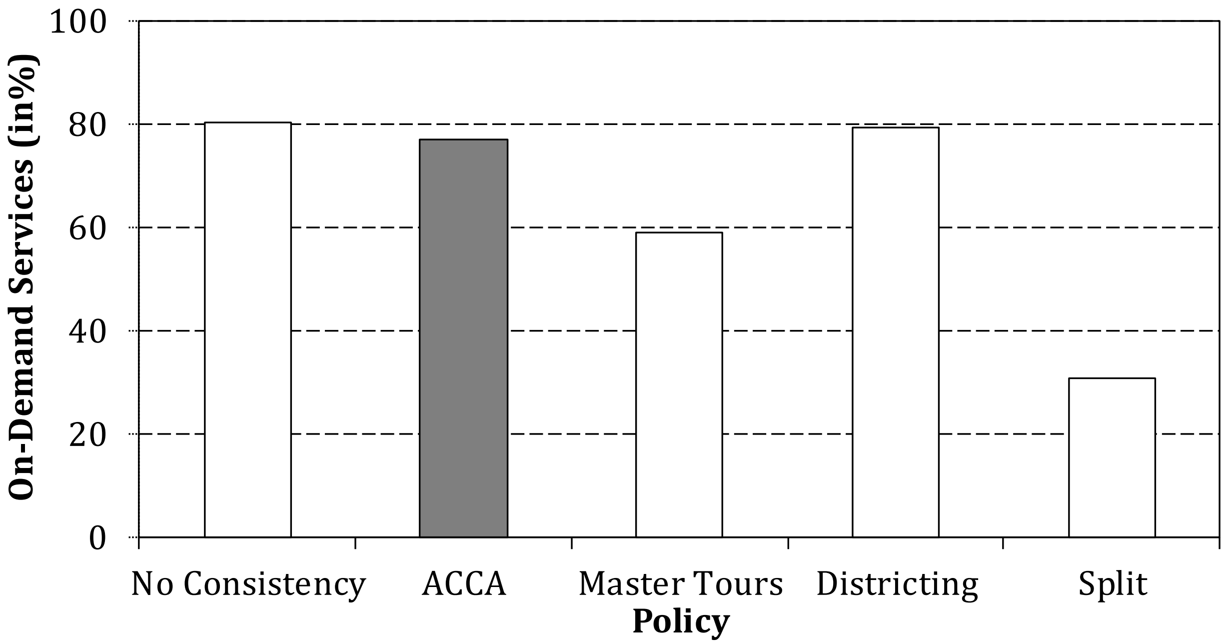

Figure 5: Average Percentage of Served On-Demand Customers.

6.3 Growing the Subscriber Base

The longer term goal of the service provider is to increase the subscriber base. One opportunity for increasing the subscriber base is to convert satisfied on-demand customers to subscribers. In the previous section, we have shown that ACCA enables service for nearly 80% of the on-demand customers. We now analyze, how growth changes both the costs of consistency as well as the service of on-demand customers.

To this end, we group the instances with respect to the ratio of subscription customers to on-demand customers. This leads to three sets, each with 100 instances. The first set comprises instances with 10 subscription customers and up to 20 on-demand customers, the second comprises instances with 15 subscription customers and up to 15 on-demand customer, and the third set comprises instances with 20 subscription customers and up to 10 on-demand customer. Thus, the sets have an increasing percentage of subscription customers which reflects an increase in subscriber base by "conversion" of on-demand customers. For these three sets, we calculate the average gap of ACCA compared to the No Consistency solution. These values are shown in Figure 6. On the x-axis, the number of subscription customers is shown. On the left y-axis, the gap compared to the No Consistency solution is depicted. The corresponding values are indicated by the bars. We observe a decreasing gap with a growing subscriber base. With only 10 subscription customers, the gap is relatively large with over 7%. With 15 subscription customers, the gap is reduced to 4.6% and with 20 subscription customers, the gap is small with 2.9%. That indicates that - compared to the optimal day to day routing without consistency - consistency provided by ACCA becomes even less expensive when the subscription customer base grows.

To show that reduction in costs does not come on the expense of the remaining on-demand customers, we calculate the percentage of on-demand customer services. The values for the different numbers of subscription customers are depicted on the right y-axis of Figure 6 and indicated by the

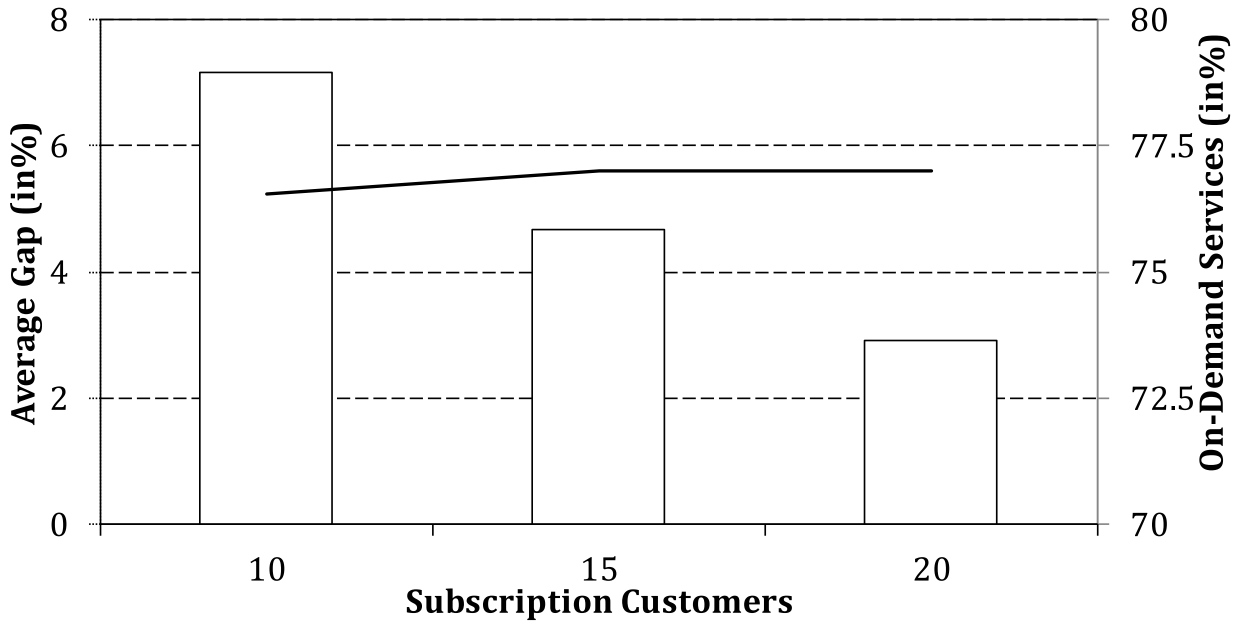

Figure 6: Gap by Increasing Percentage of Subscription Customers.

line. We observe that the value remains nearly constant around 77% regardless the number of subscription customers. Thus, if we are able to convince on-demand customers to subscribe, the costs of consistency decrease and we are still able to serve the vast majority of remaining on-demand customers.

7 Conclusion

In this paper, we have presented new and effective measures to ensure customer consistency for subscription customers, when additional on-demand customers request service on a day to day basis. To this end, we have modeled the problem as a two-stage stochastic program. On the first stage, the assignment decisions for vehicles and subscription customers are determined. On the second stage, a team orienteering problem with time windows and mandatory visits by predefined drivers is solved.

To solve the problem, we have introduced the anticipatory consistent customer assignment policy (ACCA). ACCA samples a set of scenarios and solves the team orienteering problem for each scenario by means of branch-and-price. ACCA then determines the most similar solution of all scenarios based on a newly introduced consensus functions. In our comprehensive computational study, we have shown how ACCA outperforms methods from the literature significantly by exploiting their strengths and alleviating their weaknesses. Furthermore, ACCA allows consistency with a very small increase in costs of less than 5%.

Future research may focus on both method and model. The method presented in this paper relies on an MSA framework that falls into what would be called scenario analysis (SCA) in stochastic programming. It is well established that as a general approach, SCA is not at all guaranteed to deliver good much less optimal solutions, see for example Wallace (2000) and King and Wallace

(2012) for general discussions. This result is also reflected in the development of scenario aggregation by Rockafellar and Wets (1991) and in approaches for solving stochastic integer programs using scenario aggregation as a heuristic Crainic et al. (2011). The point is that properties that are shared by the scenario solutions can be exactly what is *not* needed as they reflect the fact that the individual scenario problems are deterministic, and hence free of any kind of options (if the options come at a cost, which they normally do). On the other hand, several papers on stochastic network design, for example Wang et al. (2018), show that although the deterministic solution as such is rather bad in the stochastic setting, it has properties that are extremely useful to find good solutions to the stochastic problem. For network design it turns out that in many cases the deterministic problem (where all demands are set at their expected values) produce good network structures, but very bad arc capacities, i.e., they determine the 0/1 variables well, but not the continuous ones. This paper, while primarily comparing different heuristics, also sheds light on the qualities of SCA. We saw that SCA approaches depend heavily on how information is extracted. This always needs to be tested. Further, the No Consistency approach produces a lower bound on the true objective function value in the stochastic two-stage model. We do not know how "optimistic" this lower bound is, but we see that ACCA produces results on average only 5% above this lower bound. Hence, we see that for this problem setting, the SCA approach is not too bad when used in our specific way. A direction for future work is to try to better assess the quality of the lower bound and thereby better understand the quality of the ACCA approach.

In future research, the model may be extended to address additional business questions. In our paper, we assumed that the delivery frequency of subscription customers is the same (for example, daily, weekly, etc.). However, it may be the case that frequencies differ. One customer expects delivery every second day while another customer expects delivery once a week. This adds multi-periodicity to the problem. In this context, it may be reasonable to extend consistency to not only single drivers but small groups of drivers that can serve a customer. In may be further valuable to combine our problem with the revenue management problems addressed in the attended home delivery literature. Thus, not all on-demand customers become known at once but request delivery subsequently over an order phase. Thus, the second stage of our problem becomes a dynamic decision problem where customers need to be offered time windows for delivery.

Another interesting avenue is to analyze consistency for changing workforces and/or customer bases. On the one hand side, new customers may start ordering and on-demand customers may turn into subscription customers. Furthermore, drivers may quit and new drivers may be hired. For these longer term problem, suitable methodology should consider the flexibility of assignment decisions with respect to potential future changes in workforce and customer base.

Acknowledgment

The authors like to thank the VeRoLog and the TSL community for supporting this research with the TSL/VeRoLog-Cross Region Grant.

References

Niels Agatz, Ann Campbell, Moritz Fleischmann, and Martin Savelsbergh. Time slot management in attended home delivery. *Transportation Science*, 45(3):435–449, 2011.

Enrico Angelelli, Claudia Archetti, Carlo Filippi, and Michele Vindigni. The probabilistic orienteering problem. *Computers & Operations Research*, 81:269–281, 2017.

Nabila Azi, Michel Gendreau, and Jean-Yves Potvin. A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199(1):103–112, 2012.

Russell W Bent and Pascal Van Hentenryck. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52(6):977–987, 2004.

Ann M Campbell and Martin Savelsbergh. Incentive schemes for attended home delivery services. *Transportation Science*, 40(3):327–341, 2006.

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

John Gunnar Carlsson. Dividing a territory among several vehicles. *INFORMS Journal on Computing*, 24(4):565–577, 2012.

John Gunnar Carlsson and Erick Delage. Robust partitioning for stochastic multivehicle routing. *Operations research*, 61(3):727–744, 2013.

Teodor Gabriel Crainic, Xiaorui Fu, Michel Gendreau, Walter Rei, and Stein W Wallace. Progressive hedging-based metaheuristics for stochastic network design. *Networks*, 58(2):114–124, 2011.

Teodor Gabriel Crainic, Fausto Errico, Federico Malucelli, and Maddalena Nonato. Designing the master schedule for demand-adaptive transit systems. *Annals of Operations Research*, 194(1):151–166, 2012.

Kevin Dalmeijer and Remy Spliet. A branch-and-cut algorithm for the time window assignment vehicle routing problem. *Computers & Operations Research*, 89:140–152, 2018.

Jan F Ehmke and Ann M Campbell. Customer acceptance mechanisms for home deliveries in metropolitan areas. *European Journal of Operational Research*, 233(1):193–207, 2014.

Gianpaolo Ghiani, Emanuele Manni, and Barrett W Thomas. A comparison of anticipatory algorithms for the dynamic and stochastic traveling salesman problem. *Transportation Science*, 46(3):374–387, 2012.

Aldy Gunawan, Hoong Chuin Lau, and Pieter Vansteenwegen. Orienteering problem: A survey of recent variants, solution approaches and applications. *European Journal of Operational Research*, 255(2):315–332, 2016.

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

Dag Haugland, Sin C Ho, and Gilbert Laporte. Designing delivery districts for the vehicle routing problem with stochastic demands. *European Journal of Operational Research*, 180(3):997–1010, 2007.

Lars M Hvattum, Arne Løkketangen, and Gilbert Laporte. Solving a dynamic and stochastic vehicle routing problem with a sample scenario hedging heuristic. *Transportation Science*, 40(4):421–438, 2006.

Lars Magnus Hvattum, Arne Løkketangen, and Gilbert Laporte. A branch-and-regret heuristic for stochastic and dynamic vehicle routing problems. *Networks*, 49(4):330–340, 2007.

Ola Jabali, Roel Leus, Tom Van Woensel, and Ton de Kok. Self-imposed time windows in vehicle routing problems. *OR Spectrum*, 37(2): 331–352, 2013.

Alan J King and Stein W Wallace. *Modeling with stochastic programming*. Springer Science & Business Media, 2012.

Robert Klein, Michael Neugebauer, Dimitri Ratkovitch, and Claudius Steinhardt. Differentiated time slot pricing under routing considerations in attended home delivery. *Transportation Science*, 2017.

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

Kunlei Lian. *Service Consistency in Vehicle Routing*. PhD thesis, University of Arkansas, Fayetteville, Arkansas, May 2017.

Shih-Wei Lin and F Yu Vincent. Solving the team orienteering problem with time windows and mandatory visits by multi-start simulated annealing. *Computers & Industrial Engineering*, 114:195–205, 2017.

Yongliang Lu, Una Benlic, and Qinghua Wu. A memetic algorithm for the orienteering problem with mandatory visits and exclusionary constraints. *European Journal of Operational Research*, 268(1):54–69, 2018.

Kris Oldland. The brilliant field services pro: Creating great customer experiences and driving customer success, August 2017. URL `http://fieldservicenews.com/brilliant-field-services-pro-creating-great-customer-experiences-driving-customer-success/`. [Online; accessed 05-April-2018].

Pamela J Palomo-Martínez, M Angélica Salazar-Aguilar, and Víctor M Albornoz. Formulations for the orienteering problem with additional constraints. *Annals of Operations Research*, 258(2):503–545, 2017a.

Pamela J Palomo-Martínez, M Angélica Salazar-Aguilar, Gilbert Laporte, and André Langevin. A hybrid variable neighborhood search for the orienteering problem with mandatory visits and exclusionary constraints. *Computers & Operations Research*, 78:408–419, 2017b.

Julie Paquette, François Bellavance, Jean-François Cordeau, and Gilbert Laporte. Measuring quality of service in dial-a-ride operations: the case of a Canadian city. *Transportation*, 39(3):539–564, 2012.

R Tyrrell Rockafellar and Roger J-B Wets. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of operations research*, 16(1):119–147, 1991.

Inmaculada Rodríguez-Martín, Juan-José Salazar-González, and Hande Yaman. The periodic vehicle routing problem with driver consistency. *European Journal of Operational Research*, to appear.

M Angélica Salazar-Aguilar, André Langevin, and Gilbert Laporte. The multi-district team orienteering problem. *Computers & Operations Research*, 41:76–82, 2014.

Marius M Solomon. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35 (2):254–265, 1987.

Remy Spliet and Rommert Dekker. The driver assignment vehicle routing problem. *Networks*, 68(3):212–223, 2016. ISSN 1097-0037. doi: 10.1002/net.21694. URL `http://dx.doi.org/10.1002/net.21694`.

Remy Spliet and Guy Desaulniers. The discrete time window assignment vehicle routing problem. *European Journal of Operational Research*, 244(2):379–391, 2015.

Remy Spliet and Adriana F. Gabor. The time window assignment vehicle routing problem. *Transportation Science*, 49(4):721–731, 2015.

Remy Spliet, Said Dabia, and Tom van Woensel. The time window assignment vehicle routing problem with time-dependent travel times. *Transportation Science*, 2017. doi: 10.1287/trsc.2016.0705. URL `https://doi.org/10.1287/trsc.2016.0705`.

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

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

Anastasios D. Vareias, Panagiotis P. Repoussis, and Christos D. Tarantilis. Assessing customer service reliability in route planning with self-imposed time windows and stochastic travel times. *Transportation Science*, 2017. doi: 10.1287/trsc.2017.0748. URL `https://doi.org/10.1287/trsc.2017.0748`.

Stacy A Voccia, Ann M Campbell, and Barrett W Thomas. The same-day delivery problem for online purchases. *Transportation Science*, to appear.

Stein W Wallace. Decision making under uncertainty: Is sensitivity analysis of any use? *Operations Research*, 48(1):20–25, 2000.

Xin Wang, Teodor Gabriel Crainic, and Stein W Wallace. Stochastic scheduled service network design: The value of deterministic solutions. *INFORMS Journal on Computing*, 2018.

Richard T Wong. Vehicle routing for small package delivery and pickup services. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, pages 475–485. Springer, New York, 2008.

Zefeng Xu and Yanguang Cai. Variable neighborhood search for consistent vehicle routing problem. *Expert Systems with Applications*, 113:66–76, 2018.

Xinan Yang and Arne K Strauss. An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3):935–945, 2017.

Mengying Zhang, Jin Qin, Yugang Yu, and Liang Liang. Traveling salesman problems with profits and stochastic customers. *International Transactions in Operational Research*, 25(4):1297–1313, June 2016.

Mengying Zhang, John Wang, and Hongwei Liu. The probabilistic profitable tour problem. *International Journal of Enterprise Information Systems (IJEIS)*, 13(3):51–64, 2017.

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

A Individual Results

Table A1: Results: 3 Vehicles, 10 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	72.8	75.4	93.2	84.6	139.0
C101(2)	70.6	73.2	92.2	82.5	138.5
C101(3)	72.5	74.7	93.9	83.7	139.2
C101(4)	70.5	73.2	91.8	83.0	138.5
C101(5)	73.6	76.2	93.7	84.6	139.4
C102(1)	79.2	81.8	97.6	90.4	138.8
C102(2)	79.5	81.8	92.3	90.9	138.8
C102(3)	79.2	88.5	98.0	90.7	138.7
C102(4)	79.4	82.0	97.5	91.0	138.6
C102(5)	78.6	81.3	92.0	90.3	138.4
C103(1)	37.8	40.4	52.1	49.3	101.9
C103(2)	39.3	42.3	53.8	51.0	102.4
C103(3)	38.9	42.2	53.7	50.4	102.8
C103(4)	39.0	41.1	52.9	50.9	102.8
C103(5)	39.1	42.2	53.2	50.5	103.0
C104(1)	81.7	84.2	112.0	89.2	147.3
C104(2)	81.6	84.1	113.5	89.7	147.1
C104(3)	82.0	84.8	115.7	90.1	147.5
C104(4)	82.8	85.2	114.5	90.0	147.7
C104(5)	82.7	85.7	114.4	91.1	147.6
C105(1)	58.2	64.2	65.7	64.3	131.0
C105(2)	59.1	61.7	66.9	65.5	130.8
C105(3)	58.7	60.8	72.7	64.5	130.7
C105(4)	57.5	75.1	66.2	63.4	130.8
C105(5)	57.7	59.9	66.1	63.3	130.2
R101(1)	63.5	69.9	78.9	67.8	103.6
R101(2)	61.9	66.2	76.6	66.2	102.9
R101(3)	61.8	71.0	77.0	66.2	103.3
R101(4)	63.4	69.6	78.0	67.6	103.5
R101(5)	61.8	66.4	77.0	66.4	103.3
R102(1)	55.8	63.1	78.2	70.9	94.9
R102(2)	56.1	60.9	68.9	71.2	95.3
R102(3)	55.6	62.7	74.0	71.0	94.9
R102(4)	55.6	62.2	69.9	70.2	95.0
R102(5)	56.4	63.1	79.1	71.5	95.4
R103(1)	42.8	49.5	56.5	53.7	94.2
R103(2)	41.9	44.7	55.1	52.5	94.0
R103(3)	43.5	53.4	56.7	53.6	94.3
R103(4)	42.5	45.2	55.8	53.2	93.8
R103(5)	42.2	44.8	55.0	53.4	93.7
R104(1)	75.0	76.4	88.2	80.1	125.6
R104(2)	76.3	77.4	88.8	81.1	126.4
R104(3)	75.5	80.5	86.9	80.5	125.9
R104(4)	74.1	75.3	86.6	78.8	125.3
R104(5)	74.7	75.6	87.0	79.5	125.1
R105(1)	65.2	72.7	75.4	77.0	100.5
R105(2)	65.4	72.9	75.6	76.4	99.9
R105(3)	65.5	68.2	75.6	77.1	100.3
R105(4)	67.3	70.5	76.6	78.7	101.3
R105(5)	65.6	72.8	75.2	77.5	101.0

Table A2: Results: 3 Vehicles, 15 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	135.3	138.4	151.8	144.7	204.0
C101(2)	135.6	148.0	151.6	144.3	204.0
C101(3)	133.2	143.1	150.3	141.7	204.0
C101(4)	134.9	138.2	151.8	142.8	204.0
C101(5)	133.8	136.9	151.3	142.5	204.0
C102(1)	99.6	103.8	119.2	108.5	155.0
C102(2)	98.3	102.3	118.0	107.5	155.0
C102(3)	100.7	104.5	119.3	109.0	155.0
C102(4)	99.2	103.7	118.4	107.8	155.0
C102(5)	99.4	103.5	119.5	108.4	155.0
C103(1)	105.8	110.2	114.7	118.9	150.5
C103(2)	106.0	110.5	116.3	119.1	150.5
C103(3)	105.4	107.6	115.1	118.9	150.5
C103(4)	104.6	109.0	113.9	118.0	150.5
C103(5)	106.4	108.4	115.7	119.5	150.5
C104(1)	129.5	141.3	149.0	144.1	182.7
C104(2)	129.4	135.1	154.0	143.2	183.0
C104(3)	128.5	139.9	147.1	142.6	183.0
C104(4)	130.7	136.7	150.6	143.7	183.4
C104(5)	128.0	139.5	148.5	141.6	182.6
C105(1)	109.5	111.5	147.5	118.4	173.5
C105(2)	110.3	122.5	146.9	118.8	173.5
C105(3)	109.6	120.4	147.3	118.0	173.5
C105(4)	108.6	114.3	147.5	116.9	173.5
C105(5)	108.5	113.6	147.3	116.0	173.5
R101(1)	113.6	119.6	146.7	127.1	165.5
R101(2)	112.4	115.5	145.3	126.0	165.5
R101(3)	113.7	117.0	145.8	127.4	165.5
R101(4)	112.2	115.5	145.7	125.8	165.5
R101(5)	113.6	116.7	146.3	127.1	165.5
R102(1)	112.7	115.6	130.9	119.4	153.5
R102(2)	112.0	114.9	130.3	118.5	153.5
R102(3)	113.0	115.6	130.0	119.3	153.5
R102(4)	112.8	117.5	132.9	119.2	153.5
R102(5)	113.6	117.7	132.9	120.3	153.5
R103(1)	87.2	96.5	105.2	95.8	134.4
R103(2)	86.6	87.4	103.8	95.2	134.2
R103(3)	87.1	88.5	103.8	95.8	134.1
R103(4)	86.4	95.7	103.6	95.5	133.9
R103(5)	88.1	93.2	105.1	96.9	134.4
R104(1)	78.6	80.2	95.9	89.1	116.7
R104(2)	78.9	80.3	97.0	89.6	117.4
R104(3)	79.8	88.0	97.5	90.3	117.6
R104(4)	80.4	82.0	97.8	91.3	118.0
R104(5)	81.5	83.2	98.7	91.8	118.4
R105(1)	85.3	87.3	104.3	87.7	131.5
R105(2)	85.6	87.6	105.3	88.0	131.5
R105(3)	87.0	89.1	105.4	89.6	131.5
R105(4)	86.3	88.1	105.3	88.5	131.5
R105(5)	85.3	93.3	104.8	87.8	131.5

Table A3: Results: 3 Vehicles, 20 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	127.9	133.0	153.6	137.3	164.0
C101(2)	128.0	133.5	151.6	136.9	164.0
C101(3)	127.1	132.9	153.2	136.3	164.0
C101(4)	127.6	131.4	152.8	137.2	164.0
C101(5)	128.1	133.5	153.6	137.1	164.0
C102(1)	137.9	139.4	161.0	146.1	175.5
C102(2)	137.9	140.0	160.7	145.6	175.5
C102(3)	137.0	139.0	160.6	145.3	175.5
C102(4)	136.8	138.4	160.1	145.3	175.6
C102(5)	137.3	139.2	161.2	145.0	175.5
C103(1)	172.3	176.4	189.7	185.3	201.5
C103(2)	172.9	183.6	198.6	185.3	201.5
C103(3)	172.9	177.3	190.4	185.5	201.5
C103(4)	171.6	176.0	189.9	184.1	201.5
C103(5)	172.1	182.4	191.1	185.2	201.5
C104(1)	168.5	177.4	191.3	176.8	212.0
C104(2)	169.0	178.0	192.0	177.3	212.0
C104(3)	168.5	177.6	190.6	177.0	212.0
C104(4)	167.1	176.4	190.4	175.8	212.0
C104(5)	168.2	177.2	191.0	176.1	212.0
C105(1)	182.2	184.2	198.0	186.4	225.0
C105(2)	180.9	182.8	197.2	184.9	225.0
C105(3)	181.3	183.3	197.1	185.3	225.0
C105(4)	180.3	182.1	197.4	184.5	225.0
C105(5)	181.2	182.6	197.0	184.9	225.0
R101(1)	136.5	148.2	147.7	148.7	172.5
R101(2)	136.1	138.8	153.3	147.7	172.5
R101(3)	136.3	138.9	148.3	148.2	172.6
R101(4)	136.3	151.4	153.9	148.6	172.7
R101(5)	134.9	137.6	147.2	147.4	172.3
R102(1)	155.6	165.4	181.0	165.7	198.5
R102(2)	155.4	158.8	181.0	165.8	198.5
R102(3)	155.6	158.5	181.1	165.6	198.5
R102(4)	156.2	159.5	181.7	166.3	198.5
R102(5)	155.8	158.5	181.0	165.3	198.5
R103(1)	127.4	128.7	139.1	128.7	166.6
R103(2)	127.1	128.3	140.4	128.3	167.3
R103(3)	127.7	128.8	140.0	128.8	167.5
R103(4)	126.6	127.8	138.7	127.8	166.7
R103(5)	128.0	129.1	140.7	129.1	167.3
R104(1)	164.5	170.7	182.1	178.9	194.5
R104(2)	164.8	169.3	182.0	178.7	194.5
R104(3)	163.8	168.7	185.7	177.5	194.5
R104(4)	164.9	169.5	182.2	177.8	194.5
R104(5)	163.9	168.1	194.2	177.5	194.5
R105(1)	123.8	132.2	137.0	128.3	162.0
R105(2)	124.3	131.1	135.4	128.6	162.2
R105(3)	125.3	130.2	139.3	129.8	163.3
R105(4)	123.6	124.6	136.3	127.8	162.6
R105(5)	124.0	125.2	137.3	128.6	162.4

Table A4: Results: 5 Vehicles, 10 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	70.1	71.9	89.3	102.4	138.9
C101(2)	67.9	69.7	88.6	100.3	138.4
C101(3)	69.9	71.7	90.7	102.4	139.1
C101(4)	68.0	69.8	88.0	100.5	138.4
C101(5)	71.0	72.8	90.2	103.3	139.3
C102(1)	78.9	81.8	92.8	111.2	137.9
C102(2)	79.2	81.8	92.6	110.2	138.0
C102(3)	78.8	88.5	98.3	110.4	137.9
C102(4)	79.1	93.0	97.7	110.8	137.8
C102(5)	78.2	81.2	92.1	109.6	137.5
C103(1)	37.2	40.3	51.9	65.6	97.9
C103(2)	38.8	42.0	52.9	67.6	98.6
C103(3)	38.4	41.2	53.5	67.1	98.9
C103(4)	38.4	41.0	52.9	67.3	98.7
C103(5)	38.5	41.4	46.0	67.2	98.9
C104(1)	81.0	84.2	114.0	103.3	147.3
C104(2)	80.9	84.1	116.3	103.1	147.1
C104(3)	81.3	84.8	118.0	103.2	147.5
C104(4)	82.3	85.2	117.1	104.7	147.7
C104(5)	82.1	85.7	116.7	104.9	147.6
C105(1)	57.5	63.9	64.9	74.4	111.9
C105(2)	58.3	61.8	66.1	75.4	112.0
C105(3)	57.8	63.9	71.1	74.9	112.0
C105(4)	56.7	62.8	65.1	73.9	111.7
C105(5)	56.8	59.7	65.4	73.6	111.2
R101(1)	63.3	69.9	78.9	102.2	102.9
R101(2)	61.7	66.2	76.6	100.5	102.2
R101(3)	61.6	66.2	77.0	100.6	102.5
R101(4)	63.2	69.5	78.0	103.3	102.9
R101(5)	61.6	66.4	77.0	100.5	102.9
R102(1)	55.8	60.0	78.2	87.8	94.4
R102(2)	56.1	60.9	68.9	88.7	94.7
R102(3)	55.6	63.0	73.8	88.1	94.1
R102(4)	55.6	62.2	69.8	87.7	94.4
R102(5)	56.4	63.1	79.1	88.6	94.7
R103(1)	42.4	49.5	56.3	67.6	91.9
R103(2)	41.5	44.6	54.9	66.6	91.5
R103(3)	43.1	46.0	56.6	67.9	92.1
R103(4)	42.2	45.1	55.8	67.4	91.7
R103(5)	41.8	44.8	55.1	67.0	91.4
R104(1)	75.0	76.3	88.2	106.6	123.9
R104(2)	76.2	77.4	88.7	108.8	124.8
R104(3)	75.4	80.5	86.8	107.5	124.1
R104(4)	74.0	75.2	86.6	106.0	123.3
R104(5)	74.6	75.6	87.2	106.0	123.2
R105(1)	65.0	72.1	75.4	101.3	99.6
R105(2)	65.3	71.9	75.6	101.2	99.0
R105(3)	65.3	67.9	75.6	101.9	99.6
R105(4)	67.2	70.0	76.2	103.8	100.6
R105(5)	65.5	72.4	74.7	102.6	100.2

Table A5: Results: 5 Vehicles, 15 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	131.8	136.6	140.8	156.6	190.1
C101(2)	132.1	134.4	141.9	156.1	190.5
C101(3)	129.8	132.2	139.3	154.4	190.0
C101(4)	131.7	135.8	141.8	156.4	190.6
C101(5)	130.3	144.2	140.0	154.8	190.1
C102(1)	98.1	102.1	114.8	116.3	150.4
C102(2)	97.1	100.4	112.9	115.8	150.5
C102(3)	99.2	102.6	113.9	118.0	151.6
C102(4)	97.9	106.7	118.0	116.6	150.9
C102(5)	97.9	101.8	114.5	116.5	151.2
C103(1)	104.7	106.3	111.1	168.4	137.0
C103(2)	104.4	106.3	110.6	168.3	135.9
C103(3)	104.3	106.2	113.1	167.9	136.3
C103(4)	103.4	108.4	112.2	167.3	135.7
C103(5)	104.9	106.8	113.9	168.4	136.3
C104(1)	128.8	140.7	149.2	155.7	182.3
C104(2)	128.9	143.5	154.3	155.5	182.8
C104(3)	128.0	135.0	159.7	155.6	182.6
C104(4)	130.0	140.5	161.2	156.8	183.1
C104(5)	127.4	139.6	148.7	154.5	182.2
C105(1)	107.4	112.8	147.1	127.6	168.1
C105(2)	108.2	113.9	146.8	128.9	168.0
C105(3)	107.6	113.0	147.0	127.8	168.0
C105(4)	106.5	112.1	146.9	126.4	167.8
C105(5)	105.9	113.6	146.8	125.8	167.3
R101(1)	112.2	118.9	144.8	130.6	162.6
R101(2)	111.1	114.5	143.9	129.7	162.7
R101(3)	112.4	115.8	144.2	130.9	162.7
R101(4)	110.9	114.2	143.9	129.1	162.5
R101(5)	112.4	115.7	144.7	131.0	163.0
R102(1)	110.2	120.6	130.5	137.2	145.5
R102(2)	109.0	114.2	134.6	136.9	145.0
R102(3)	110.2	115.3	135.0	138.2	145.5
R102(4)	110.3	120.1	135.2	137.7	145.3
R102(5)	111.0	115.8	130.8	138.5	145.8
R103(1)	87.1	88.0	103.8	117.4	134.0
R103(2)	86.5	87.3	103.5	117.1	133.8
R103(3)	87.1	87.7	103.3	117.9	133.7
R103(4)	86.3	87.2	103.0	117.0	133.7
R103(5)	88.0	89.0	104.5	118.9	134.2
R104(1)	77.8	79.6	94.8	112.8	112.2
R104(2)	78.1	83.2	95.8	113.7	112.8
R104(3)	79.0	87.7	96.5	114.5	113.1
R104(4)	79.7	84.7	96.6	115.1	113.6
R104(5)	80.7	82.8	97.7	116.3	114.0
R105(1)	85.1	87.3	104.6	105.7	126.0
R105(2)	85.4	87.6	105.4	106.5	126.2
R105(3)	86.8	101.6	105.6	107.4	126.4
R105(4)	86.1	88.1	112.3	107.0	126.2
R105(5)	85.1	86.8	105.0	105.6	125.9

Table A6: Results: 5 Vehicles, 20 Subscription Customers

Distribution	No Consistency	ACCA	Master Tours	Districting	Split
C101(1)	127.4	133.0	153.7	158.0	161.0
C101(2)	127.7	133.5	152.9	158.3	161.0
C101(3)	126.7	132.7	153.3	157.5	160.8
C101(4)	127.2	131.4	147.7	158.0	160.8
C101(5)	127.6	133.5	153.6	158.1	160.9
C102(1)	137.5	139.4	161.1	174.3	174.4
C102(2)	137.5	140.0	160.8	174.4	174.4
C102(3)	136.6	139.0	162.2	173.5	174.4
C102(4)	136.5	138.4	160.2	173.7	174.1
C102(5)	136.9	139.2	161.2	173.9	174.3
C103(1)	166.8	168.6	185.6	184.7	196.5
C103(2)	167.4	169.3	180.6	185.4	196.5
C103(3)	167.6	169.4	181.0	185.5	196.5
C103(4)	165.7	167.5	180.4	183.8	196.5
C103(5)	166.6	168.5	180.7	184.0	196.5
C104(1)	164.8	168.4	175.5	175.5	202.8
C104(2)	164.8	169.6	178.2	175.3	202.7
C104(3)	164.2	168.1	177.7	175.5	202.7
C104(4)	162.8	166.3	173.8	173.8	202.4
C104(5)	163.9	168.0	177.0	175.5	203.2
C105(1)	179.7	183.5	196.1	194.4	218.2
C105(2)	178.3	181.2	195.9	193.4	217.8
C105(3)	179.1	181.8	193.1	194.3	218.0
C105(4)	178.0	180.6	196.4	193.0	217.6
C105(5)	178.3	181.1	195.9	193.0	217.6
R101(1)	136.4	148.7	152.8	164.9	172.5
R101(2)	135.9	138.8	153.3	164.9	172.5
R101(3)	136.2	138.9	148.8	164.8	172.6
R101(4)	136.3	145.7	153.9	165.3	172.7
R101(5)	134.8	137.6	147.8	163.1	172.3
R102(1)	153.6	158.3	178.8	163.7	194.5
R102(2)	153.6	158.3	181.4	164.1	194.5
R102(3)	153.6	157.6	178.5	163.5	194.4
R102(4)	154.1	158.9	179.1	164.4	194.4
R102(5)	153.6	158.1	178.7	163.6	194.3
R103(1)	127.4	128.6	139.5	166.4	162.1
R103(2)	127.1	128.3	140.7	166.3	162.8
R103(3)	127.7	128.8	140.0	166.4	162.8
R103(4)	126.5	127.7	139.1	165.8	162.2
R103(5)	128.0	129.0	140.7	167.1	162.8
R104(1)	161.5	169.7	173.4	176.1	193.6
R104(2)	161.6	166.8	179.3	176.5	193.7
R104(3)	160.7	165.9	172.6	175.4	193.6
R104(4)	161.2	167.1	178.0	176.4	193.6
R104(5)	160.7	168.8	177.3	175.2	193.6
R105(1)	123.6	132.0	135.8	153.9	160.5
R105(2)	124.0	131.0	136.8	153.8	160.7
R105(3)	125.1	130.0	137.5	156.2	161.8
R105(4)	123.3	124.4	135.3	153.9	161.0
R105(5)	123.9	125.1	135.9	154.7	160.9