Scalable Robust and Adaptive Inventory Routing

Dimitris Bertsimas
Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139
dbertsim@mit.edu

Swati Gupta
H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332
swatig@alum.mit.edu

Joel Tay
Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139
joeltay@mit.edu

We consider the finite horizon inventory routing problem with uncertain demand, where a supplier must deliver a particular commodity to its customers periodically, such that even under uncertain demand the customers do not stock out, e.g. supplying residential heating oil to customers. Current stochastic, robust or adaptive optimization techniques that solve this problem with uncertain demand do not scale to real-world data sizes, with the status quo being only able to perform inventory routing for ~100 customers. We propose a scalable approach to solving a robust and adaptive mixed integer optimization formulation that is made tractable with algorithms for generating worst-case demand vectors, heuristic route selection, warm starts and column generation. We demonstrate experimentally a mean reduction in stockouts of over 94% in our robust and adaptive formulations, translating to a cost savings of over 14%. We also show how to modify our model to achieve further cost savings through fleet size reduction. Our robust and adaptive formulations are tractable for ~6000 customers.

Key words: robust optimization, vehicle routing, inventory routing, stock-out, demand, uncertainty

1. Introduction

We consider the rich problem of inventory routing where a supplier has a contract with individual customers to monitor their inventory of a commodity that diminishes over time, and to resupply that commodity to maintain customer stocks above a certain threshold. Some sizeable industries concerned with inventory routing problems of this type are those supplying commodities such as soft drinks in vending machines, portable water in offices, or heating oil in residential areas. In many of these inventory routing applications, the presence of uncertainty in the customers' demand for the commodity (and other uncertainties in data, e.g. temperature in heating oil usage models) is a critical issue that must be addressed in order to provide solutions that are of practical value in the real world. In this paper, we provide novel scalable and adaptive algorithms to address the inventory routing problem using a robust and adaptive optimization framework.

Given a network of customers spread over a geographic area, the supplier needs to make the following key operational decisions:

- **Fleet size**: ahead of the operational period, the supplier needs to decide the number of vehicles to be maintained and the crew size required. A larger crew size and more vehicles increase the cost of operation, whereas a reduction in these may impact the quality of service negatively, and require a larger emergency fleet to handle stockouts.
- **Routes and Schedules**: the supplier needs to determine which routes to utilize to visit customers, and when to schedule these routes, while minimizing their cost of operation (thus maximizing their profits).
- **Refuelling quantities**: when a customer is visited, the supplier needs to determine how much of the commodity to resupply. Attempting to resupply all customers to their full capacity might not be feasible for the vehicles' capacity, or it might limit the number of customers that a vehicle can resupply.

Having defined our key operational decisions, we now consider the key objectives that a supplier is concerned with, namely: (i) **reducing the frequency of stockouts** and (ii) **minimizing the cost of operations**. Reducing the frequency of stockouts is important, as, besides the obvious damage to brand image that results from customers' stocks being depleted, it is also highly undesirable for suppliers because they have to designate vehicles to make unplanned emergency replenishments of these customers, often at very short notice.

Regarding operational cost, much of the inventory routing literature (e.g. Irnich et al. (2014)) has focused on minimizing the routing cost while maintaining a desired level of service. However, an important reason that many approaches to this problem do not scale well is that they attempt to solve for the optimal routes. As this requires solving the Vehicle Routing Problem as a subproblem, it becomes difficult to use these approaches to solve problems of the sizes required in real-world applications.

In our discussions with a local heating oil company, we learnt that the most important objective for them was to be able to reduce their fleet size, while maintaining schedules that are robust to uncertainties in the rate of customers' demand for the commodity. As commercial routing solutions are already of relatively high quality, it was thus felt that this would allow planners to reduce both expensive stockout resupplies, and the capital, maintenance and labor cost of the vehicle fleet, which is usually particularly high in the peak season and has a greater cost savings potential than fuel cost (e.g. Hall (2016)). We therefore focused our efforts on reducing the vehicle fleet size and reducing stockouts, and used a fast heuristic for the routing component, which ensures feasibility for the routing for a given vehicle fleet size.

Current exact approaches in the literature (Solyali et al. (2012), Aghezzaf (2007)) solve only up to around a hundred customers and do not scale to problem sizes that arise in real life, while heuristic solutions usually decompose the problem into a series of problems with shorter time horizons because of concerns about tractability and uncertain data (e.g. Dror and Ball (1987), Prescott-Gagnon et al. (2014)). Our main application throughout the paper is to companies that provide heating oil in residential areas. For example, a typical company of this nature in New England might have a customer base spanning north central Massachusetts and southern New Hampshire with around 10,000 customers. Our key contribution is a robust and adaptive mixed integer optimization (MIO) formulation that scales to large problem sizes, augmented with a demand uncertainty set that varies with temperature and heuristic route generation. Using data sets generated from real temperature data, we demonstrate both the effectiveness and scalability of our approach.

The rate of demand of the commodity has typically been considered in the literature (e.g. Chepuri and Homem-De-Mello (2005), or the survey of Gendreau et al. (1996)) to be either (dynamically) deterministic or stochastic. A deterministic rate of demand, as with many optimization problems, leads to more tractable but less realistic models. A stochastic rate of demand, however, is less tractable for large instances and often leads to heuristic solutions which are sensitive to the assumptions made about the probability distribution of the demand. In contrast, a robust optimization approach combines the tractability of deterministic models with the realism of stochastic approaches by modeling uncertainty in a deterministic manner, and leads to solutions that are less sensitive to the probabilistic assumptions made about the underlying demand.

Our contributions in this work can be summarized as follows:

1) **Robustness.** We present a robust formulation of the uncertainty set for demand that captures, for the case of resupplying heating oil, the dependence of demand on temperature as well as individual customers' rates of consumption. This results in a novel non-convex uncertainty set which we are able to tractably optimize over, thus generating the critical worst-case demand scenarios.

2) **Adaptability.** For the case where customer demand can be recorded remotely, we present an approach that allows us to adapt our operational decisions according to observed demand. We demonstrate computationally that the adaptive solutions outperform both the deterministic and robust formulations.

3) **Scalability.** By combining:

(a) novel ways to generate the critical worst-case demand scenarios,

(b) automated neighborhood route selection,

(c) generating constraints on the fly for the adaptive formulation, we are able to solve problems with ~6000 customers over a time horizon of 151 days, which is approximately the length of the planning season, within two hours for both the robust and adaptive formulations.

4) **Quality of solutions.**

(a) We demonstrate that the robust solutions of our model materially decrease stockouts and are relatively insensitive to estimation noise in demand and temperature, achieving across a variety of data sets of sizes ranging from 51 to 5915, an average reduction in stockouts of over 94% from a deterministic model.

(b) We show that both the robust and adaptive formulations can be used to reduce vehicle fleet size, while still outperforming the deterministic solution.

(c) We demonstrate that the robust and adaptive solutions lead to a decrease in total operational cost for the supplier, when combining routing cost with vehicle fleet cost and the cost of resupplying customers who experience stockouts.

The remainder of this paper is structured as follows: in Section 2, we survey some of the related literature and discuss why current approaches do not scale well. In Section 3, we introduce deterministic, robust and adaptive models for the capacitated inventory routing problem. We define our uncertainty set, and provide an algorithm that maximizes affine functions of demand over the uncertainty set. In Section 4, we discuss our techniques for route generation and some heuristics to further improve our routes. We detail our experiments and computational results in Section 5. We finally conclude with overall discussion and some future directions in Section 6.

2. Related Work

Vehicle routing problems (VRPs) arise naturally from many problem contexts, and as such have been extensively studied in many flavors. Beginning with "The Truck Dispatching Problem", proposed by Dantzig and Ramser (1959), the difficulty of these problems and their relevance to many industries have generated much research over the past few decades.

One of the best-studied formulations of vehicle routing problems is the capacitated vehicle routing problem (CVRP), which in its most basic form describes the problem of determining a minimum-cost set of routes by which a fleet of delivery vehicles with limited capacity delivers quantities of a product or commodity to customers at various locations. When the costs of potential routes and the customer demands are assumed to be fixed and known, this is a deterministic problem. Early approaches for getting exact solutions of the CVRP were for decades dominated by branch-and-bound algorithms (e.g. Christofides and Eilon (1969), Christofides et al. (1981), Laporte et al. (1986)); in addition, branch-and-cut algorithms were later developed with many different families of cuts (Laporte et al. (1985), Augerat (1995), Ralphs et al. (2003), Lysgaard et al. (2004), Baldacci et al. (2004)), often building on research on the Travelling Salesman Problem. More recently, another popular approach is to solve the problem using column generation alongside cut generation (e.g. Fukasawa et al. (2006), Baldacci et al. (2008), Pecin et al. (2014)). We refer the reader to Cordeau et al. (2006), Golden et al. (2008), Laporte (2009), Baldacci et al. (2010), and Toth and Vigo (2014) for detailed literature surveys about the CVRP and related vehicle routing problems.

However, the solutions to deterministic VRPs can be sensitive to errors or uncertainties in the parameters of the problem, becoming suboptimal or even infeasible for real-world actualizations. This has typically been addressed by taking the uncertain parameters as random variables, and utilizing stochastic programming to formulate the model. Assuming a known probability distribution for the uncertain parameters, probabilistic guarantees can then be made (e.g. a chance-constrained VRP). (More generally, the field of stochastic programming is described in much greater detail in Birge and Louveaux (2011) and Shapiro et al. (2014), just to give two examples.) However, stochastic VRPs are much harder to solve than their deterministic counterparts (Dror et al. (1989)). Developing exact algorithms that solve these problems to optimality has been challenging for problems of any realistic size, and much work has been done on heuristics (for a detailed survey, Toth and Vigo (2001)), and recently, metaheuristics (Toth and Vigo (2014), Archetti and Speranza (2014)) that work well on VRPs.

Often, in addition to finding suitable routes, the planner has to manage levels of inventory between a number of customers or retailers, i.e., solve an inventory routing problem (IRP) (Federgruen and Zipkin (1984)). An important version of the IRP that we address is that of Vendor-Managed Inventory (VMI). VMI is a business practice that was popularized in the 1980s by Walmart and Procter & Gamble, where the suppliers are responsible for monitoring the inventory levels of their customers, and deciding on replenishment schedules and quantities accordingly. This can result in benefits such as lower inventory required (Waller et al. (1999)), cost reductions (Sahin and Robinson (2005)), and a smaller bullwhip effect in supply chains (Disney and Towill (2003)).

In formulating VMI models, it is usually assumed that the planner has real-time telemetry measurements of the customer's inventory, an assumption which we make in the adaptive formulation in Section 3.4. This allows the planner to be responsive to changing conditions and monitor actual consumption. However, often it is the case that either this technology has not yet been implemented, or the planner is operating in one of the several industries where making such telemetry measurements is not cost-effective. In this case, when the model is first solved, we have to make do with estimates of the customers' existing inventory levels, as we do in the robust formulation in Section 3.2.

Due to the increased difficulty of simultaneously solving for schedules and quantities, in practice the planner can simplify the problem of formulating VMI models in a variety of ways. For example, the quantities could be decided by a deterministic order-up-to-level policy (Bertazzi et al. (2002), Archetti et al. (2007)) where the customer is always replenished to maximum capacity. Alternatively, sample-based methods can be used to extend methods for deterministic demand to work with stochastic demand (Hemmelmayr et al. (2010)).

An approach commonly used in practice is to model the resupplying problem as a series of one-day problems, where forecasting models based on the historical consumption of the customers and temperature data indicate which customers are likely to need replenishment within the next day. The planner then optimizes a capacitated routing problem to determine routes that will cover these customers, along with customers who will need replenishment in the following days. Prescott-Gagnon et al. (2014) introduces a tabu search metaheuristic and two large neighborhood search metaheuristics for this problem, while Dror et al. (1985) and Dror and Ball (1987) developed a two-stage approach for propane delivery where the customers are first assigned to specific days, and then routes are constructed daily.

A different approach to modeling IRPs with stochastic demands has been to handle the demands' dependence on uncertain temperatures by using Markov decision process models, (e.g. Dror et al. (1989), Kleywegt et al. (2002) and Adelman (2004), among many more). These approaches tend to have higher computational times, making them less useful for more complex problems of realistic sizes. For more comprehensive reviews of the general inventory routing problem (IRP) and various solution approaches, we refer the reader to Federgruen and Simchi-Levi (1995), Campbell et al. (1998), Kleywegt et al. (2002), Cordeau et al. (2006), Bertazzi et al. (2008), Coelho et al. (2013), and a recent comparison of different IRP formulations (Archetti et al. (2014)).

A problem with a similar flavor is that of IRPs with transportation procurement, where the planner outsources the deliveries to the customers. In contexts where this is possible, it can lead to more flexibility as the planner does not have a fixed fleet size constraint, and requires optimization of the purchase of transport capacity in each time period, rather than routing each vehicle. We direct interested readers to the recent works of Bertazzi et al. (2015) and Bertazzi et al. (2016).

A paradigm that has proven useful in approaching problems modeling optimization under uncertainty is Robust Optimization (RO) (for instance Bertsimas and Sim (2003), Ben-Tal et al. (2009), Delage and Ye (2010), Bertsimas et al. (2011a)). This approach leads to solutions that are guaranteed to satisfy the constraints for all uncertain parameters in a chosen uncertainty set, and often leads to tractable models requiring weaker assumptions on the uncertain parameters than stochastic formulations. RO formulations have been found in practice to yield solutions that are competitive with the optimal deterministic solution, and perform significantly better in worst-case scenarios. They also tend to be less affected by errors in parameter estimation or structural assumptions (Goldfarb and Iyengar (2003), Bertsimas and Sim (2004)).

While demand uncertainty has long been considered in its stochastic form (Bertsimas (1992), Bertsimas and Simchi-Levi (1996), Gendreau et al. (1996)), recent works have proven the usefulness of RO in formulating certain varieties of VRPs (Ordóñez (2010)). For instance, Sungur et al. (2008) consider a formulation of the single-stage Robust Capacitated VRP (RCVRP) under demand uncertainty that can be solved deterministically, using the budget-of-uncertainty approach first developed in Bertsimas and Sim (2003), and Gounaris et al. (2013) consider the RCVRP with more general demand uncertainty that can be reformulated to yield numerical solutions.

Solyali et al. (2012) and Aghezzaf (2007) have previously addressed the inventory routing problem within a RO framework. Solyali et al. (2012) report solving instances with a branch-and-cut algorithm, solving a Travelling Salesman problem as a subproblem exactly, for up to 30 customers and a time horizon of seven periods. Aghezzaf (2007) uses a heuristic approach to generate routes, proposing a nonlinear MIO problem, and reports solving for cyclic distribution routes for 50 customers. In contrast, our methods allow us to solve problems with the number of customers two orders of magnitude larger than both of these, over a time horizon which is an order of magnitude larger than Solyali et al. (2012), by solving a deterministic MIO to generate robust solutions, and using a cutting-plane algorithm to generate adaptive solutions.

Finally, we consider the problem of formulating an adaptive multistage robust optimization model. While the fully adaptive robust optimization problem is intractable via a dynamic programming approach, affinely-adaptive robust optimization solutions have been found to perform almost as well, while retaining the tractability of single-stage robust optimization problems (Ben-Tal et al. (2004), Bertsimas et al. (2010)). This approach has recently been applied to the unit commitment problem in power generation (Bertsimas et al. (2013), Lorca et al. (2016)). Finite adaptability is a different approach that works well for some multistage robust optimization models (Bertsimas and Caramanis (2010), Bertsimas et al. (2011b)), but we chose affine adaptability due to its stronger scalability characteristics.

3. Problem Formulation

To view the problem we address in a concrete context, consider the following inventory routing problem over a finite horizon: A company has customers who consume a homogeneous commodity over time, and a fleet of vehicles that is used to resupply them. We would like to generate a feasible schedule of routes for the vehicles that satisfies capacity constraints for users and vehicles, and leads to a low likelihood of stockouts for the customers.

A key insight that helps us achieve this is the observation that in practice, customers are often located in small neighborhoods, and that most of the variable cost (i.e., travelling distance) of the routing problem is derived from travel between depots and these small neighborhoods of customers. Within these neighborhoods, then, routes can be optimized sufficiently for industrial purposes by local search algorithms such as 2-opt (Croes (1958)). Therefore, our approach is to think of routes not as a list of customers, but as a neighborhood which a vehicle might travel between in a given time period. Upon selecting a route for a vehicle, a feasible schedule is then one which assigns customers to that vehicle that are on that route, i.e., in the associated selected neighborhood. Correspondingly, we assign costs to routes based on travel between the depot and the customers in the selected neighborhood, bearing in mind that the costs are to be taken as accurate only to the first order. The realized cost will depend on the customers we assign to the vehicle servicing a route.

This has a few key advantages. Firstly, it leverages the current knowledge of the company in the form of extant routes and neighborhoods, driver experience and other geographic and network information. In other words, it allows us to warm start our model with a set of routes that are already known to be feasible, and gradually introduce routes to improve the solution quality of our model as needed. Furthermore, it significantly reduces the solution space of feasible routes, which helps the model to scale to large problem sizes more easily. By varying the sizes and coverage of the set of routes that we optimize over, we can exercise control over the tradeoff between scalability and solution quality, as needed.

Naturally, this observation does not hold true for all problem domains. Where the routing cost is dominated by the cost of individual links of the route, e.g. in a problem where a driver has to visit all customers on a small route, then more sophisticated vehicle routing algorithms will be necessary. Often, though, heuristic algorithms are sufficient for route optimization at a local level, and indeed planners in many industries will be best served to use commercially available routing software within small neighborhoods.

For the vehicle routing problem under consideration, our decision-making has to take into account two sources of uncertainty in the demand for the commodity. The more important of these is the uncertainty associated with changes in temperature, which is correlated across all the customers. To a smaller extent, there is also an uncertainty in demand specific to each customer, which we assume is uncorrelated across customers. Using the well-established RO methodology, we define appropriate uncertainty sets (see (12) below) that capture these phenomena. In Section 3.3, we discuss ways to initialize the parameters of this uncertainty set from observations or simulations of the uncertain data.

3.1. Nominal Formulation

We begin by defining the nominal formulation of the inventory routing problem - in other words, we solve the problem for the case where demand is fixed rather than uncertain. Consider N customers who need to be resupplied over a time horizon T, who we index as customers $i \in [N] = \{1, \ldots, N\}$. The customers are to be resupplied with a fleet of M vehicles, each of capacity S. In a single time period, the vehicles can be assigned to a tour $\theta \in [\Theta]$, each which has associated cost c_θ. Each customer i has a maximum capacity of Q_i, and we suppose that customer i begins the season with Z_i of the commodity remaining. For the nominal formulation, we assume that demand d_i^t is known for all customers and time periods.

We consider the following decision variables:

- $g_{i,\theta}^t$, the amount of fuel that customer i will be resupplied via route θ at time t,
- u_i^t, the total amount of fuel that customer i will be supplied at time t,
- Binary variable v_θ^t which is 1 if and only if tour θ is selected at time t.

Then, the nominal formulation is:

$$\min_{\mathbf{u},\mathbf{v},\mathbf{g}} \quad \sum_{t=1}^{T} \sum_{\theta=1}^{\Theta} c_\theta v_\theta^t \tag{1}$$

$$\text{s.t.} \quad 0 \leq Z_i + \sum_{\tau=1}^{t} u_i^\tau - \sum_{\tau=1}^{t} d_i^\tau, \ \forall i \in [N], \ \forall t \in [T], \tag{2}$$

$$Z_i + \sum_{\tau=1}^{t} u_i^\tau - \sum_{\tau=1}^{t-1} d_i^\tau \leq Q_i, \ \forall i \in [N], \ \forall t \in [T], \tag{3}$$

$$\sum_{\theta=1}^{\Theta} v_\theta^t \leq M, \ \forall t \in [T], \tag{4}$$

$$u_i^t \leq \sum_{\theta=1}^{\Theta} g_{i,\theta}^t, \ \forall i \in [N], \ \forall t \in [T], \tag{5}$$

$$\sum_{i=1}^{N} g_{i,\theta}^t \leq S v_\theta^t, \ \forall \theta \in [\Theta], \ \forall t \in [T], \tag{6}$$

$$g_{i,\theta}^t = 0, \ \forall i \in [N], \ \forall \theta : i \notin \theta, \ \forall t \in [T], \tag{7}$$

$$g_{i,\theta}^t \geq 0, \ \forall i \in [N], \ \forall \theta \in [\Theta], \ \forall t \in [T],$$

$$u_i^t \geq 0, \ \forall i \in [N], \ \forall t \in [T],$$

$$v_\theta^t \in \{0, 1\}, \ \forall \theta \in [\Theta], \ \forall t \in [T].$$

Eq. (1) expresses the cost minimization objective. Eq. (2) guarantees that each customer is resupplied so that their supply of the commodity is never depleted, while Eq. (3) enforces customer capacity constraints. Eq. (4) respects the fleet size. Eq. (5) ensures that the amount of fuel assigned to refuel a customer is also assigned to some route in the same time period. Eq. (6) both allows us to assign fuel to a route only if the route is actually selected, and if so, also enforces vehicle capacity limits. Eq. (7) ensures that assignments are only made for customers that are on a given route.

3.2. Robust Formulation

Now we move to the robust formulation of the inventory routing problem. Here, rather than assume we know what the demand $\mathbf{d}$ is, we assume rather that it lies within an uncertainty set $\mathcal{U}$ which we have constructed beforehand. We discuss the construction of $\mathcal{U}$ in more detail in the next subsection. We also assume that the amounts of fuel that customers start with, Z_i, take values in the interval $[\underline{Z}_i, \bar{Z}_i]$.

As before, we consider the same variables $g_{i,\theta}^t$, u_i^t and v_θ^t. Then the robust formulation is the same as before, except that now the constraints given by Eqs. (2) and (3) become:

$$0 \leq Z_i + \sum_{\tau=1}^{t} u_i^\tau - \sum_{\tau=1}^{t} d_i^\tau, \ \forall i \in [N], \ \forall t \in [T], \ \forall \boldsymbol{d} \in \mathcal{U}, \ \forall Z_i \in [\underline{Z}_i, \bar{Z}_i], \tag{8}$$

$$Z_i + \sum_{\tau=1}^{t} u_i^\tau - \sum_{\tau=1}^{t-1} d_i^\tau \leq Q_i, \ \forall i \in [N], \ \forall t \in [T], \ \forall \boldsymbol{d} \in \mathcal{U}, \ \forall Z_i \in [\underline{Z}_i, \bar{Z}_i], \tag{9}$$

with the same interpretations.

3.3. Constructing $\mathcal{U}$

We describe here one method of constructing $\mathcal{U}$ based on insights from the Central Limit Theorem (see Bandi and Bertsimas (2012)), particularly applicable to the scenario of supplying heating oil to residences during winter. To do this, we assume that for any given customer, expected demand is constant above a certain temperature and increases linearly as the temperature decreases below that point. Specifically, for customer i, we assume that there exists a breakpoint Ψ_i above which expected demand is B_i^0, and that if the temperature decreases below Ψ_i, the expected demand increases with a slope (against temperature) of B_i^1. We operate with the supposition that Ψ_i, B_i^0 and B_i^1 have been estimated for each customer from historical data.

We now assume that for each time period t, the temperature τ_t is subject to i.i.d. variation, and thus construct a CLT-style uncertainty set $\mathcal{U}_\tau$ for the temperature,

$$\mathcal{U}_\tau = \left\{ \boldsymbol{\tau} : \left| \frac{\sum_{t=1}^{T} (\tau_t - \bar{\tau}_t)}{\sigma_\tau \sqrt{T}} \right| \leq \Gamma_\tau, \ \ \bar{\tau}_t - 3\sigma_\tau \leq \tau_t \leq \bar{\tau}_t + 3\sigma_\tau \ \ \forall t \in [T] \right\}. \tag{10}$$

Here $\bar{\tau}_t$ and σ_τ are the mean and standard deviation of the temperatures respectively, and Γ_τ is a robust parameter that we are free to select, which we discuss below. We refer to the value $\sqrt{T}\sigma_\tau\Gamma_\tau$

as the budget of variation in temperature, i.e., the net amount our temperatures are allowed to vary from their means.

We next consider the additional noise in the demand. For simplicity, we assume the demand is subject to additional zero-mean noise that has the same distribution for each time period, but is i.i.d. across customers, and thus construct a CLT-style uncertainty set $\mathcal{U}_\epsilon$ for the noise in demand,

$$\mathcal{U}_\epsilon = \left\{ \boldsymbol{\epsilon} : \left| \frac{\sum_{i=1}^{N} \epsilon_i}{\sigma_\epsilon \sqrt{N}} \right| \leq \Gamma_\epsilon, \quad -3\sigma_\epsilon \leq \epsilon_i \leq 3\sigma_\epsilon \;\; \forall i \in [N] \right\}, \tag{11}$$

where σ_ϵ is the standard deviation of the demand noise.

This gives us our uncertainty set for demand, $\mathcal{U}$, which, as described above, consists of all demand vectors for which the corresponding temperature and demand noise simultaneously lie within the uncertainty sets $\mathcal{U}_\tau$ and $\mathcal{U}_\epsilon$, respectively.

$$\mathcal{U} = \left\{ \boldsymbol{d} : d_i^t = B_i^0 + B_i^1 \max(0, \Psi_i - \tau_t) + \epsilon_i, \;\; \boldsymbol{\tau} \in \mathcal{U}_\tau, \;\; \boldsymbol{\epsilon} \in \mathcal{U}_\epsilon \right\}, \tag{12}$$

where B_i^0, B_i^1 and and Ψ_i are all parameters estimated from data.

Selecting robust parameters The uncertainty sets $\mathcal{U}_\tau$ and $\mathcal{U}_\epsilon$ involve the parameters Γ_τ and Γ_ϵ that represent the planner's desired balance between optimality and robustness. We next outline our approach for selecting these parameters. Assuming that temperatures τ_t are independent for each time period t, with mean $\bar{\tau}_t$ and variance σ_τ^2 from an otherwise unknown distribution, we select Γ_τ such that $\mathcal{U}_\tau$ contains the realized temperature with probability 99% for large T. Specifically, from the Central Limit Theorem,

$$\lim_{T \leftarrow \infty} \mathbb{P}\left(\left| \sum_{t=1}^{T} (\tau_t - \bar{\tau}_t) \right| \leq \Phi^{-1}(0.99)\, \sigma_\tau \sqrt{T} \right) = 0.99, \tag{13}$$

where Φ is the cdf of the standard normal distribution, and so we select $\Gamma_\tau = \Phi^{-1}(0.99)$. A similar approach is used for selecting Γ_ϵ. For other possible approaches to selecting the robust parameters, see Ben-Tal and Nemirovski (1999), Bertsimas and Sim (2004), Ben-Tal et al. (2009), Bertsimas et al. (2018).

For a given planning horizon T and N customers, let the demands $\boldsymbol{d} \in \mathbb{R}^{N \times T}$ lie in the uncertainty set $\mathcal{U}$ given in Eq. (12), which is non-convex, necessitating a novel approach to generate critical worst-case scenarios.

Note that the only robust constraint in our formulation is Eq. (8), which requires us to protect against the maximum and minimum values of $\sum_{\tau=1}^{t} d_i^\tau$ over $\mathcal{U}$ for each customer i in $[N]$

and each day t in $[T]$. We next give algorithm Opt-Temp that allows us to optimize over $\mathcal{U}$ an affine combination of convex non-increasing functions of temperature. Note that demand without customer-specific noise is a convex non-increasing function of temperature in our model. In addition, as each robust constraint only involves one customer, the worst-case ϵ_i can always be taken to be $3\sigma_\epsilon$ for maxima, and $-3\sigma_\epsilon$ for minima. Given a customer i and day t, we can use these to construct a demand vector $d_i \in \mathbb{R}^T$ that maximizes the sum $\sum_{\tau=1}^{t} d_i^\tau$. This enables us to solve the robust formulation as a deterministic problem, vastly improving computational performance. For notational convenience, we refer to the natural projection of $\mathcal{U}$ onto the set of demand vectors for customer i as $\mathcal{U}_{[i]}$.

Summary of algorithm: To maximize the sum $\sum_{t=1}^{T} a_t d_t(\tau_t)$ for $\tau \in \mathcal{U}_\tau$ as defined by Eq. 10, where $d_t(\tau_t)$, for each t, is a convex non-increasing function of τ, we let the set of days with non-negative affine coefficients, i.e., $a_t \geq 0$, be T_1, and those with negative affine coefficients, i.e., $a_t < 0$, be T_2. In algorithm Opt-Temp, we consider two cases: (i) $|T_1| \geq |T_2|$ and (ii) $|T_1| < |T_2|$. For the first case, we set all temperatures to be at their upper bounds, i.e., $\tau_t = \bar{\tau}_t + 3\sigma_\tau$. We then greedily choose the days in T_1 and for each such t decrease its corresponding temperature as far as possible. In the second case, we set all the temperatures to be at their lower bounds, i.e., $\tau_t = \bar{\tau}_t - 3\sigma_\tau$. We then optimize the restricted objective function over the days T_2 using standard convex optimization techniques. In both cases, we ensure that the temperatures selected respect the bound $|\sum_{t=1}^{T} \tau_t - \bar{\tau}_t| \leq \Gamma_\tau \sqrt{T} \sigma_\tau$, where Γ_τ is a robust parameter. To prove optimality, we show that there exists an optimal solution with at most one temperature not attaining one of its bounds, and that our algorithm finds such a solution.

Formally, we present in Algorithm 1 an algorithm Opt-Temp for maximizing an affine combination of convex non-increasing functions over $\mathcal{U}_\tau$. The algorithm finds, for convex non-increasing functions $d_t(\tau)$ and coefficients a_t, a temperature vector yielding $\max_{\boldsymbol{\tau} \in U_\tau} \sum_{t=1}^{T} a_t d(\tau_t)$. In our presentation of the algorithm we use a sorting function Sort(R), which sorts the set of days R in descending order of the difference in the objective function when the temperature is changed from $\bar{\tau}_t + 3\sigma_\tau$ to $\bar{\tau}_t - 3\sigma_\tau$, i.e., Sort$(R) = \{t_1, t_2, \ldots, t_{|R|}\}$ such that $\Delta(t_x) \geq \Delta(t_y)$ whenever $x < y$, where:

$$\Delta(t) = a_t(d_t(\bar{\tau}_t - 3\sigma_\tau) - d_t(\bar{\tau}_t + 3\sigma_\tau)).$$

$D(q, k, F)$ calculates the increase in objective value that we could get, for fixed k and F, of allowing q to be the single time period that does not achieve either of its temperature bounds. Figure 1 explains the logic of the algorithm graphically.

Theorem 1. *The temperature vector $\boldsymbol{\tau}^* \in \mathbb{R}^T$ output by the Algorithm 1 maximizes $\sum_{t=1}^{T} a_t d_t(\tau_t)$ over $\mathcal{U}_\tau$.*

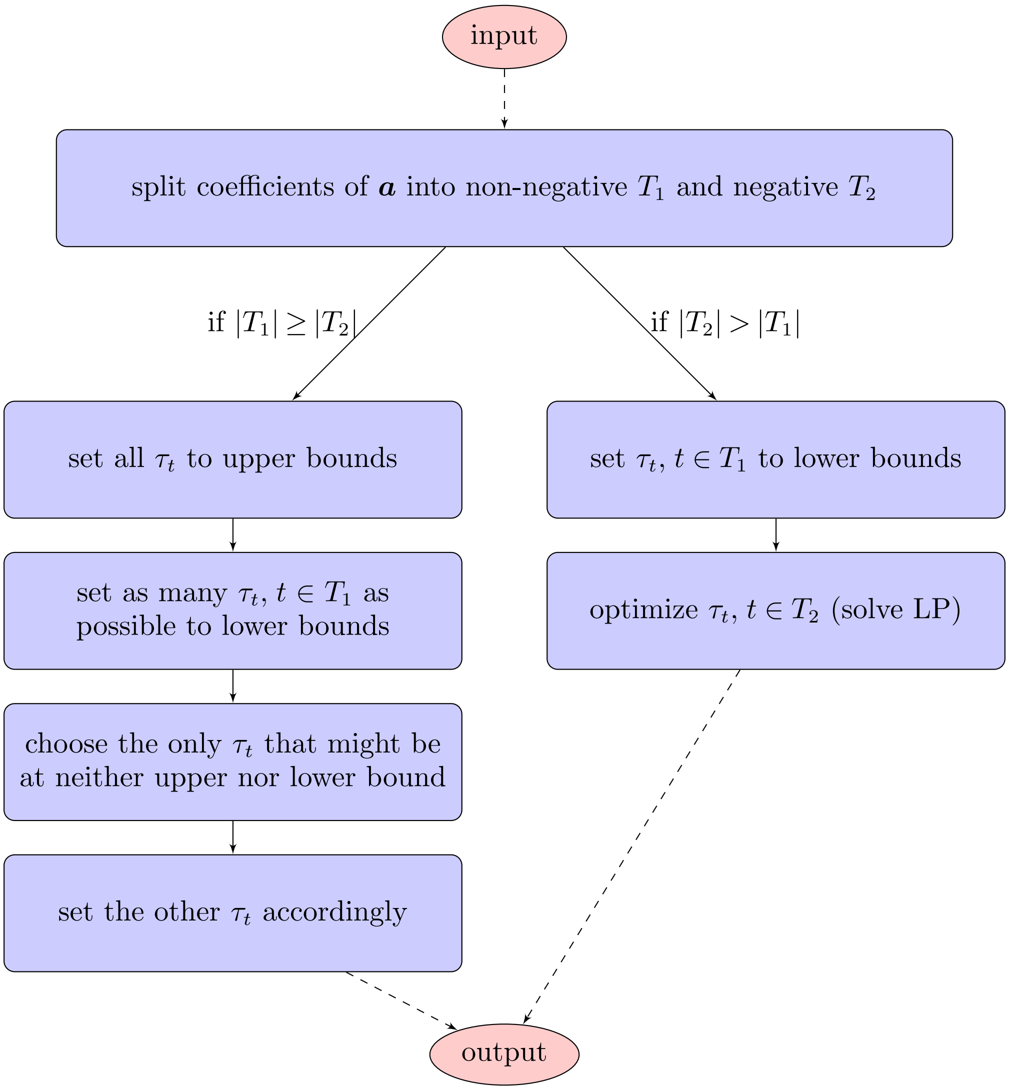

Figure 1 The logic of Algorithm 1.

Proof: We first show that $\boldsymbol{\tau}^*$ is feasible. For the case where $|T_1| < |T_2|$, the temperatures are guaranteed to be feasible by definition of the optimization subproblem that we solve (Note that as we only optimize for T_2, this is a convex optimization problem and so tractable). To show feasibility for the case where $|T_1| \geq |T_2|$, we consider the bookkeeping variable F, which tracks the value of $\sum_{t=1}^{T}(\tau_t - \bar{\tau}_t)$. Before we update a temperature, we check that F will not exceed the CLT-type bounds $-\Gamma\sqrt{T}\sigma \leq F \leq \Gamma\sqrt{T}\sigma$, and limit the magnitude of our updates accordingly. Similarly, the temperatures are initialized at their upper bounds and never decreased by more than 6σ, the width of the feasible interval for a single temperature. Also, note that we are assured of the existence of a feasible solution (e.g. setting the temperatures to their mean values). Thus, $\boldsymbol{\tau}^*$ is feasible.

Next we prove that $\boldsymbol{\tau}^*$ is optimal. Suppose we had a feasible temperature vector where for some $r \in T_1$, $\tau_r > \bar{\tau}_r - 3\sigma$, and for some $s \in T_2$, $\tau_s < \bar{\tau}_s + 3\sigma$. Then we could decrease τ_r and increase τ_s

Algorithm 1: OPT-TEMP

Input: $\Gamma > 0$, σ, $\bar{\boldsymbol{\tau}} \in \mathbb{R}^T$, $\boldsymbol{a} \in \mathbb{R}^T, d_t : \mathbb{R} \to \mathbb{R}\ \forall t \in [T]$

Output: $\tau \in \arg\max_{\tau \in \mathcal{U}_\tau} \sum_{t=1}^T a_t d_t(\tau_t)$

$T_1 = \{t \in T : a_t \geq 0\},\ T_2 = T \setminus T_1, k = l = m = 1;$

if $|T_1| \geq |T_2|$ **then**

  $\tau_t = \bar{\tau}_t + 3\sigma\ \forall t \in [T],\ F = 3T\sigma;$

  $\{t_1, t_2, \ldots, t_{|T_1|}\} = \textsc{Sort}(T_1);$

  while $F \geq 6\sigma - \Gamma\sqrt{T}\sigma$ **and** $k < |T_1|$ **do**

    $(\tau_{t_k}, F) \leftarrow (\tau_{t_k} - 6\sigma, F - 6\sigma);$

    $k \leftarrow k + 1;$

  end

  if $F > -\Gamma\sqrt{T}\sigma$ **and** $k < |T_1|$ **then**

    $q^* = \arg\max_{q \in T_1} D(q, k, F);$

    if$(q^* \leq k-1)$ **then** $(\tau_{t_k}, \tau_{t_{q^*}}, F) \leftarrow (\tau_{t_k} - 6\sigma, \tau_{t_{q^*}} + 6\sigma - F - \Gamma\sqrt{T}\sigma, -\Gamma\sqrt{T}\sigma);$

    else $(\tau_{t_{q^*}}, F) \leftarrow (\tau_{t_{q^*}} - F - \Gamma\sqrt{T}\sigma, -\Gamma\sqrt{T}\sigma);$

  end

end

else $\tau = \arg\max_{\tau \in \mathcal{U}'} \sum_{t=1}^T a_t d_t(\tau)$ for $\mathcal{U}' = \mathcal{U}_\tau \cap \{\tau : \tau_t = \bar{\tau}_t - 3\sigma_\tau\ \forall t \in T_1\};$

where SORT is the sorting function defined earlier, and

where $D(q, k, F) =$

$$\begin{cases} a_q d_q(\bar{\tau}_q + 3\sigma - F) - a_q d_q(\bar{\tau}_k + 3\sigma) & \text{if } q > k, \\ a_q d_q(\bar{\tau}_q + 3\sigma - F) + a_{k+1} d_{k+1}(\bar{\tau}_{k+1} - 3\sigma) - a_q d_q(\bar{\tau}_q - 3\sigma) - a_{k+1} d_{k+1}(\bar{\tau}_{k+1} + 3\sigma) & \text{if } q \leq k. \end{cases}$$

by some small ϵ, while not decreasing the objective function. This means that we can limit ourself to optimal solutions where either the temperatures in T_1 all attain their lower bounds $\bar{\tau}_r - 3\sigma_\tau$, or the temperatures in T_2 all attain their upper bounds $\bar{\tau}_s + 3\sigma_\tau$. (We will show that the smaller set attains its bounds.)

Case (i) $|T_1| < |T_2|$: We show that in this case, there exists at least one optimal temperature vector τ^* such that $\tau^*_t = \bar{\tau}_t - 3\sigma$ for all days in T_1. (Note that such an optimal temperature vector is easy to find: for days in T_1, all the temperatures are at their lower bounds, and temperatures for days in T_2 can be found using linear optimization). Consider any optimal temperature vector τ^{opt} that maximizes $\sum_{t=1}^T a_t d_t(\tau_t)$ such that all the temperatures in T_2 attain their upper bounds, i.e., $\tau^{opt}_t = \bar{\tau}_t + 3\sigma$ for $t \in T_2$ (if not, as argued above, all the temperatures in T_1 must be at their lower bounds, thus proving our claim). Let $F^{opt} = \sum_{t=1}^T (\tau^{opt}_t - \bar{\tau}_t) = \sum_{t \in T_1} (\tau^{opt}_t - \bar{\tau}_t) + 3\sigma|T_2|$. Note that $F^{opt} \leq \Gamma\sqrt{T}\sigma$ since τ^{opt} is feasible. Now, consider a temperature vector τ' such that $\tau'_t = \bar{\tau}_t - 3\sigma$ for $t \in T_1$ and $\tau'_t = \bar{\tau}_t + 3\sigma$ for $t \in T_2$. Let $F' = \sum_{t \in T_1} (\tau'_t - \bar{\tau}_t) + \sum_{t \in T_2} (\tau'_t - \bar{\tau}_t) = 3(|T_2| - |T_1|) \geq 0$. Also, note that $F' \leq F^{opt} \leq \Gamma\sqrt{T}\sigma$. Thus, τ' is feasible and its function value is no worse than τ^{opt}.

Hence, we have proved that there exists an optimal temperature vector which attains the lower bounds for temperatures in T_1.

In this case, optimality follows from the definition of the optimization subproblem that we solve, restricted to T_2.

Case (ii) $|T_1| \geq |T_2|$: Similar to the previous case, we can assume that the temperatures in T_2 all attain their upper bounds, i.e., for all $t \in T_2$, we have $\tau_t = \bar{\tau}_t + 3\sigma$. We next show that there exists such an optimal solution where at most one temperature τ_t for a $t \in T_1$ is neither at $\bar{\tau}_t - 3\sigma$ nor $\bar{\tau}_t + 3\sigma$.

Suppose we had some feasible solution with $r, s \in T_1$, $\tau_r \neq \bar{\tau}_r \pm 3\sigma$, $\tau_s \neq \bar{\tau}_s \pm 3\sigma$. We want to adjust these temperatures so that one attains its bound, without decreasing the objective function. Let $a = \min(\tau_r - (\bar{\tau}_r - 3\sigma), \bar{\tau}_s + 3\sigma - \tau_s)$, $b = \min(\bar{\tau}_r + 3\sigma - \tau_r, \tau_s - (\bar{\tau}_s - 3\sigma))$. By the convexity of d_r and d_s, we use Jensen's inequality to get:

$$\frac{b}{a+b} d_r(\tau_r - a) + \frac{a}{a+b} d_r(\tau_r + b) \geq d_r(\tau_r), \tag{14}$$

$$\frac{a}{a+b} d_s(\tau_s - b) + \frac{b}{a+b} d_s(\tau_s + a) \geq d_s(\tau_s). \tag{15}$$

Adding these inequalities implies that either $d_r(\tau_r - a) + d_s(\tau_s + a)$ or $d_r(\tau_r + b) + d_s(\tau_s - b)$ must be at least $d_r(\tau_r) + d_s(\tau_s)$, and so we can adjust τ_r and τ_s as desired. We thus can limit ourselves to considering temperature vectors with at most one temperature not attaining either of its 3-sigma bounds.

Finally, suppose we knew that τ_t was the temperature not attaining its bounds. Then, a simple greedy algorithm for the temperature values at lower and upper bounds would give the optimal temperature vector.

In our algorithm, we iterate over all the choices for the day with the temperature not attaining its bounds, and select the one with the best objective value. The remaining temperatures are set to their upper or lower bounds, sorted so that they have the same output a greedy algorithm would have. Therefore, we obtain a temperature vector that maximizes the objective function over both sets of days, T_1 and T_2.

□

We can now explicitly find the minima and maxima over $\mathcal{U}$ for the sums of demand seen in the robust constraints. For the maximum demand, we construct a worst-case temperature vector for $\sum_{\tau=1}^{t} d_i^t$ using the above algorithm. As mentioned above, the robust constraints each involve just a single customer and so ϵ_i can be taken to be $3\sigma_\epsilon$.

For the minimum demand, the algorithm requires us to solve a convex optimization subproblem. In fact, for each $s \in [T]$ and $i \in [N]$, we can compute the minimum value of $\sum_{t=1}^{s} d_i^t$ by solving the following linear optimization:

$$\min \sum_{t=1}^{s} d_i^t \tag{16}$$

$$\text{s.t.} \quad -\Gamma_\tau \sqrt{T}\sigma_\tau \leq \sum_{t=1}^{T} (\bar{\tau}_t - \tau_t) \leq \Gamma_\tau \sqrt{T}\sigma_\tau, \tag{17}$$

$$\bar{\tau}_t - 3\sigma_\tau \leq \tau_t \leq \bar{\tau}_t + 3\sigma_\tau, \quad \forall t \in [T], \tag{18}$$

$$d_i^t \geq B_i^0 + B_i^1 x_t - 3\sigma_\epsilon, \quad \forall i \in [N], \quad \forall t \in [T], \tag{19}$$

$$x_t \geq \Psi_i - \tau_t, \quad \forall t \in [T], \tag{20}$$

$$\boldsymbol{x} \geq \mathbf{0}, \ \boldsymbol{d}_i \geq \mathbf{0}. \tag{21}$$

Similar to before, ϵ_i can be taken to be $-3\sigma_\epsilon$. This allows us to replace our robust constraints with $2NT$ deterministic constraints, in each case picking the appropriate endpoint of the interval $[\underline{Z}_i, \bar{Z}_i]$ to robustify against (i.e., $\underline{Z}_i$ for lower bounds and $\bar{Z}_i$ for upper bounds).

In our computational experiments, we observed that as the robust constraints for time t do not involve customer demands for time periods beyond that, it improved the performance of our algorithm to project $\mathcal{U}$ onto the first t time periods and find the worst-case vector corresponding to $\Gamma_\tau \sqrt{t/T}$. This weakens the theoretical probabilistic guarantees that we can make, because the Central Limit Theorem might not apply in small cases. However, in our experiments this adaptation did not result in a significant increase in stockouts, but it did produce a significant decrease in the cost (and conservativeness) of the models. Note that the protection against stockouts is weakest against the earlier time periods at the very start of the heating season, when a customer is less likely to stockout anyway.

3.4. Affine Adaptive Robust Formulation

As technology develops, it is becoming increasingly feasible for companies to install sensors in customers' buildings. This might allow them, for instance, to track the daily consumption of their customers, improving the solution quality of their planning models. While it may be impractical to alter the fleet and crew schedule on short notice, we adapt our formulation so that the quantity of fuel resupplied will now be partially responsive to the actual demand observed. Without this new information from sensors, a company is limited to observations made during scheduled deliveries, i.e., the aggregated demand between refuelling decisions, which is much less informative.

We now define an affine adaptive robust formulation that applies to the scenario where we have additional real-time information about customers' demands. Instead of having the model decide on exact amounts to refuel each customer daily, we set the quantities refuelled to be affine functions of the demand in the previous days, and solve for the coefficients of these affine functions.

To make the formulation adaptive, we substitute each u_i^t with an affine function of previous days' demands: $u_i^t = b_i^{0,t} + \sum_{j=1}^{t-1} b_i^{j,t} d_i^j$ (remember that consumption for a day occurs after any refuelling on that day), where the various $b_i^{j,t}$ are now variables we are solving for. Similarly, we substitute each $g_{i,\theta}^t$ with $g_{i,\theta}^t = a_{i,\theta}^{0,t} + \sum_{j=1}^{t-1} a_{i,\theta}^{j,t} d_i^j$, where $a_{i,\theta}^{j,t}$ are variables.

This leads to the following formulation:

$$\min_{\mathbf{a},\mathbf{b},\mathbf{v},\mathbf{g}} \quad \sum_{t=1}^{T} \sum_{\theta=1}^{\Theta} c_\theta v_\theta^t \tag{22}$$

$$\text{s.t.} \quad 0 \leq Z_i + \sum_{\tau=1}^{t} (b_i^{0,\tau} + \sum_{j=1}^{\tau-1} b_i^{j,\tau} d_i^j) - \sum_{\tau=1}^{t} d_i^\tau, \ \forall i \in [N], \ \forall t \in [T], \ \forall \mathbf{d} \in \mathcal{U}, \tag{23}$$

$$Z_i + \sum_{\tau=1}^{t} (b_i^{0,\tau} + \sum_{j=1}^{\tau-1} b_i^{j,\tau} d_i^j) - \sum_{\tau=1}^{t-1} d_i^\tau \leq Q_i, \ \forall i \in [N], \ \forall t \in [T], \ \forall \mathbf{d} \in \mathcal{U}, \tag{24}$$

$$\sum_{\theta=1}^{\Theta} v_\theta^t \leq M, \ \forall t \in [T], \tag{25}$$

$$b_i^{0,t} + \sum_{j=1}^{t-1} b_i^{j,t} d_i^j \leq \sum_{\theta=1}^{\Theta} (a_{i,\theta}^{0,t} + \sum_{j=1}^{t-1} a_{i,\theta}^{j,t} d_i^j), \ \forall i \in [N], \ \forall t \in [T], \ \forall \mathbf{d} \in \mathcal{U}, \tag{26}$$

$$\sum_{i=1}^{N} (a_{i,\theta}^{0,t} + \sum_{j=1}^{t-1} a_{i,\theta}^{j,t} d_i^j) \leq S v_\theta^t, \ \forall \theta \in [\Theta], \ \forall t \in [T], \tag{27}$$

$$a_{i,\theta}^{j,t} = 0, \ \forall i \in [N], \ \forall \theta : i \notin \theta, \ \forall t \in [T], \ \forall j \in \{0, \ldots, t-1\}, \tag{28}$$

$$a_{i,\theta}^{j,t} \geq 0, \ \forall i \in [N], \ \forall \theta \in [\Theta], \ \forall t \in [T], \ \forall j \in \{0, \ldots, t-1\},$$

$$b_i^{j,t} \geq 0, \ \forall i \in [N], \ \forall t \in [T], \ \forall j \in \{0, \ldots, t-1\},$$

$$v_\theta^t \in \{0, 1\}, \ \forall \theta \in [\Theta], \ \forall t \in [T].$$

Note that all the constraints in the adaptive robust formulation have the same interpretation as their counterparts in the robust formulation, although fuel supplied is now adaptive in that it is an affine function of demand. Furthermore, the starting quantities, Z_i, are no longer taken to be uncertain, as we would expect real-time measurements of demand to also yield exact information about the customers' remaining fuel.

The number of variables in the adaptive formulation is an order of magnitude greater than the nominal or robust case. Thus it is impractical to solve it using a deterministic linear MIP, as we did for the nominal formulation. In addition, the constraints (23), (26) and (27) involve products of our decision variables and the uncertain demand. This means that to separate over these constraints one would need to solve a quadratic optimization problem over a non-convex set.

We instead use a cutting-plane algorithm that exploits the structure of the uncertainty set $\mathcal{U}$, to tractably solve the adaptive formulation. Given a candidate solution, we can, for each of the constraints (23) or (26), use Opt-Temp to give us the worst-case demand corresponding to that particular constraint and candidate solution, i.e., if the constraint is violated, we can find a demand vector in $\mathcal{U}$ that shows the violation, giving us a feasible cutting plane. Specifically, as noise for each customer is constant across time periods, the noise ϵ_i for a worst-case demand vector for that constraint-candidate pair is given by a greedy algorithm sorting on its coefficient, $\sum_{i=1}^{N} a_{i,\theta}^{0,t} \sum_{j=1}^{t-1} a_{i,\theta}^{j,t}$. If $\mathbf{d}^*$ is such a demand vector that causes a violation, we can add new deterministic constraints that check the violated constraints against $\mathbf{d}^*$. We then reoptimize the model, each time enforcing a check against all the previously-violated constraints with their associated demand vectors, and generate a new candidate solution. We repeat the process until the candidate solution we have does not violate any of the constraints.

4. Route Generation

Route generation is a widely studied problem, especially given its importance in various vehicle and inventory routing problems (for example, see Laporte (1992), Francis et al. (2008), Golden et al. (2008)), applied to a plethora of real-world applications such as routing for bakery companies (Pacheco et al. (2012)), blood product distribution (Hemmelmayr et al. (2009a)), grocery industry (Semet and Taillard (1993)), ship-routing (Gunnarsson et al. (2006)). The literature is ripe with a number of exact (Baldacci et al. (2011), Laporte (1992), Laporte et al. (1986)) and heuristic (Hemmelmayr et al. (2009b), Vidal et al. (2012), Liu (1997), Semet and Taillard (1993)) approaches for route generation. Since tractability is a major concern with exact approaches, we employ heuristic methods to generate a feasible set of routes. We would however like to emphasize that exploring route generation techniques is not the main focus of this work.

In this work, we consider a route to be not just a feasible tour, but a neighborhood of customers that can potentially be served on a given trip. The schedule that the solution of the model outputs will specify which subset of customers of the chosen route is actually served in a given time period. From discussions with the company that we collaborated with, we learnt that even without considering the labor and maintenance cost of a vehicle, most of the variable cost that arises from the choice of a route is given in our problem domain by the fuel and time taken to travel to the neighborhood of customers to be serviced. As this led to our primary objectives of fleet size reduction and reducing stockouts, it was natural for us to make the simplifying assumption that there is a fixed cost to "visit" a neighborhood, irrespective of how many houses on that route are actually resupplied in any given time period.

We generate an initial set of feasible neighborhoods for our datasets in two ways: (i.) using a user-operated GUI where the supplier can manually select neighborhoods that are typically served together, (ii.) using an automated sweep of the customer locations. The first approach is preferred when the supplier would like to utilize prior knowledge and accumulated expertise about different neighborhoods. Our second approach is an automated sweep of the geographic area under consideration. Our algorithm creates a *cover* of the entire space with 'neighborhoods' or boxes so that the number of customers in each box lies in an interval. This interval is determined from the expected demand and the vehicle capacity in such a way that we would expect a vehicle to have the capacity to resupply about half the customers in the neighborhood to full capacity. Experimentally, we found that these sweeps generate a good first set of feasible routes that make the problem scalable. As emphasized above, we do not claim any special properties from this heuristic selection, and in particular we would not necessarily expect it to outperform user-selected routes.

We further improve the quality of the routes using a set-cover formulation inspired by the work of (Cacchiani et al. (2014)) augmented with well-studied tabu-search heuristics for improving routes (for e.g. in Cordeau et al. (1997), Gendreau et al. (1994)). We consider a set of possible schedules for the customers, covered using feasible routes such that on any day at most M vehicles are used. However, we deviate from Cacchiani et al. (2014) by assuming that the cost, c_θ, of a route θ is given by the Euclidean distance of the tour suggested by the 2-opt (TSP) heuristic from the depot that a customer is served from (which is usually good enough in practice). We construct the following input from a pre-computed solution of the nominal problem.

- T, the total number of days in the planning horizon,
- M, the maximum number of vehicles in the fleet,
- N, the number of customers,
- S_i, the set of valid schedules that a customer i could be visited at. In order to construct this set, we consider the service schedule suggested by the nominal solution, i.e., which days customers are scheduled to be resupplied. We then allow the schedule to shift such that each customer is visited only up to three days before or after a delivery in the nominal model's service schedule, and call the resulting set of potential service schedules S_i.
- Ω, the set of feasible routes. We initialize Ω with the set of routes obtained by either neighborhood selection or automatic sweep. We will now describe how we use the following formulation to improve the quality of the set of routes.

We next select routes, using the following set-covering-like formulation. Let a^i_θ be a constant equal to 1 if customer i is on route θ, and 0, otherwise, for all $i \in [N], \theta \in \Omega$. Let b^t_p be equal to 1 if time period t is selected on service schedule p, and 0, otherwise. We use two sets of binary variables: x^t_θ and y^i_p. Here x^t_θ is 1 if and only if route θ is selected on day t and 0, otherwise. Finally,

y_p^i is 1 if and only if schedule p is selected for customer i and 0, otherwise. We now formulate an binary linear program as follows:

$$\min \sum_{\theta\in\Omega}\sum_{t\in[T]} x_\theta^t c_\theta \tag{29}$$

$$\text{s.t.} \sum_{p\in S_i} y_p^i = 1, \quad i\in[N], \tag{30}$$

$$\sum_{\theta\in\Omega} x_\theta^t a_\theta^i - \sum_{p\in S_i} y_p^i b_p^t \geq 0, \quad i\in[N],\ t\in[T], \tag{31}$$

$$\sum_{\theta\in\Omega} x_\theta^t \leq M, \quad \forall t\in[T], \tag{32}$$

$$x_\theta^t \in\{0,1\}, \quad r\in\Omega,\ t\in[T], \tag{33}$$

$$y_p^i \in\{0,1\}, \quad p\in S_i,\ i\in[N]. \tag{34}$$

The objective function aims at minimizing the cost of the routes selected. Constraints (30) guarantee that exactly one feasible service schedule is selected for a customer. Constraints (31) guarantee that if the selected service schedule for customer i requires service on day t, then there must be a route selected on day t with customer i on the route. Constraints (32) ensure that at most M vehicles are used on any day, thereby respecting the fleet size.

We relax the above integer optimization problem and do column generation on the resulting optimization relaxation. We generate a set of candidate routes using the following heuristic operations on all the existing routes:

- **Insert** a customer into an existing neighbouring route,
- **Swap** two customers from neighboring routes,
- **Remove** customers from an existing route,
- **Construct** new routes for each day t by considering customers i such that their dual variables $p_{i,t}$ take large values.

For each candidate route θ for each day t, we compute its reduced cost as follows: $\bar{c}_\theta = c_\theta - \sum_{i\in\theta} p_{i,t} - p_M$ where $p_{i,t}$ is the optimal dual variable corresponding to constaints (31) and p_M is the dual variable corresponding to the constraint (32). We add the route that has the most negative reduced cost out of our set of candidate solutions, until a set with the required number of routes covering each customer is generated.

5. Computational Experiments

To test the scalability of our problem formulations and the quality of our solutions, we generated a number of datasets based on real-world problems. We imported customer locations from a few instances in the TSPLIB, the standard test bed of the Traveling Salesman Problem, with the size of these instances (i.e., the number of customers) ranging from 51 to 5915 (data instances `eil51,` `rat99, kroB200, rat575, pcb1173, d2103, rl5915`). For simplicity, we assumed a homogenous fleet of vehicles (in particular, with identical maximum capacity), operating from a single depot located at the centroid of the users.

We assumed all the customers to have homogenous heating oil tanks with identical maximum capacity. For each customer, we generated a base temperature above which their expected demand was near-zero and constant, and below which it increased linearly as temperature decreased. We randomly generated family sizes for each customer ranging from 2 to 4, and scaled the mean demand accordingly, adding noise in both the temperature and for each user's demand as described in Section 3.2. To tune our uncertainty parameters for temperatures, we used actual data for Boston for the months of November 2013 to March 2014, representing a full season of heating oil consumption (Weather Underground (2014)).

We further generated estimated initial amounts for each customer, and subjected these to further zero-mean uncertainty proportional to the difference of these amounts to the full customer capacity, encapsulating the principle that a customer with higher usage or a customer who was serviced a longer time ago should have more uncertainty in their starting amounts. To be precise, if the estimated initial amount for a customer was z_i^{est}, the noisy initial amount was

$$z_i = z_i^{est} + (Q - z_i^{est}) \times U_i, \;\; U_i \sim U(-1/2, 1/2).$$

(These are all decisions appropriate for actual companies (e.g. as described in Dror (2005) for propane delivery), and in practice companies already use estimates based on similar parameters.)

Family sizes and uncertain initial amounts were randomized separately to get *training* and *testing* datasets. We used the training set to tune our robust parameters Γ_τ and Γ_ϵ, and set these parameters correspondingly in the testing set to test the performance of our approach in terms of running time and effectiveness. We assumed in our experiments that a centralized depot serves all the customers, although the formulation generalizes easily to applications with multiple depots, each with their own fleet of vehicles. We give exact details of the parameters in Appendix A.

For our computational experiments, we let the nominal, robust and adaptive formulations solve for two hours each, using the nominal solution as a warm start (though infeasible) to the robust model, and the robust solution as a warm start to the adaptive model.

While the adaptive model we presented in Section 3.4 schedules a customer with refuelling quantities that are affine in all of that customer's observed demand, we improved the tractability of our implementation of the adaptive model by relaxing the number of terms of the adaptability. Specifically, we limited the refuelling quantity for a customer for time period t to be a base amount $b_i^{0,t}$, plus a term linear in that customer's demand during time period $t-1$ (i.e., d_i^{t-1}), a term linear in the total demand of the customer during time periods $t-3$ and $t-2$ (i.e., $d_i^{t-3}+d_i^{t-2}$), and a term linear in the total demand of the customer during time periods $t-7$ to $t-4$ (i.e., $d_i^{t-7}+\cdots+d_i^{t-4}$). We also solved for non-adaptive quantities across different vehicles, i.e., for $g_{i,\theta}^t$. We used the robust solution as a warm start to the base amount, and initialized the other affine coefficients as zero. Note that this is, by design, already a feasible solution to the adaptive model, ensuring that the adaptive model always gave feasible output.

Each dataset (both training and testing) contained fifty generated scenarios for any computational experiment. All our instances were solved with Gurobi 6.0.0 on a Intel Xeon E5687W (3.1 GHz) processor with 16 cores and 128 GB of RAM. Our models were coded in JuMP (Dunning et al. (2017)).

To investigate the quality of the solutions resulting from our models, we ask the questions of a) whether our robust and adaptive models lead to fewer stockouts, b) what effect this has on service cost, and whether this suggests a possible reduction of the vehicular fleet size.

5.1. Stockouts

We first investigate whether our robust and adaptive inventory routing models lead to a reduction in stockouts. As mentioned above, the three formulations were solved sequentially (nominal-robust-adaptive) on problems of a fixed vehicle fleet size and a decision horizon of 151 days. Figures 2, 3 and 4 show the average percentage of customers who experienced stockouts for the nominal, robust and adaptive formulations respectively. The standard variation in the (temperature and demand) noise was scaled appropriately for each dataset so that the models were trained on a base value of 5.

We observe that across data sets, while the nominal formulation had between 160%-225% of customers stocking out (some customers experienced stockouts multiple times), the robust formulation decreased this to below 9% of all customers, and in most cases half of that or even less. Stockouts decreased further by 0.5%-1% of all customers from the robust to the adaptive formulation for all the data sets, i.e., a decrease of over 5% in stockouts from the robust formulation. We also notice that the robust and adaptive formulations were less sensitive than the nominal formulation to increasing variance in the noise (i.e., errors in tuning the robust parameters).

We explore how the reduction in stockouts was distributed across the fifty scenarios for each data point in the above experiment. Figure 5 shows the standard box plot for the reduction in robust model stockouts as a fraction of nominal model stockouts, for the uncertainty regime of the base level of variance in the noise. We observe that every scenario generated had at least 94% relative reduction in stockouts from the nominal to robust models, with an average of over 96% relative reduction for each dataset.

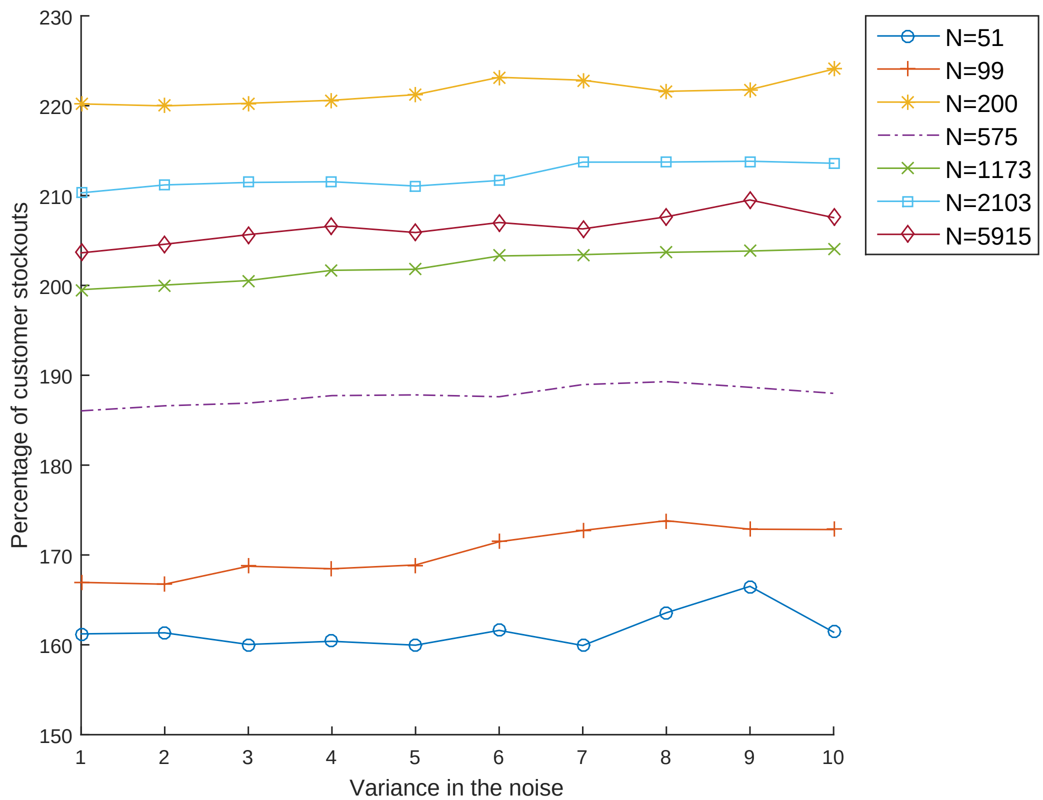

Figure 2 **Average stockout percentages for the nominal solutions for data sets of different sizes.**

Finally, we ran experiments to determine whether we had chosen sufficiently many routes or neighborhoods to cover each customer. We found that as long as customers were included in 2-3 neighborhoods, any additional coverage was superfluous and did not result in further cost reductions. This makes sense because it is unlikely that it improves our situation for a customer to be assigned to a vehicle that is not going to resupply other customers in their immediate neighborhood.

5.2. Service cost

We next consider the effect on the cost of servicing the customers with the different formulations. As before, we solve problems of a fixed vehicle fleet size. To get the combined cost of the problem, we consider costs from two sources, namely: 1) the variable cost from the routes, which is the objective function of the optimization model, and 2) the cost of refuelling a customer who experiences a stockout. Because these stockouts occur randomly throughout the course of a time period, and must be addressed urgently, the planner must send an emergency refuelling vehicle out each time

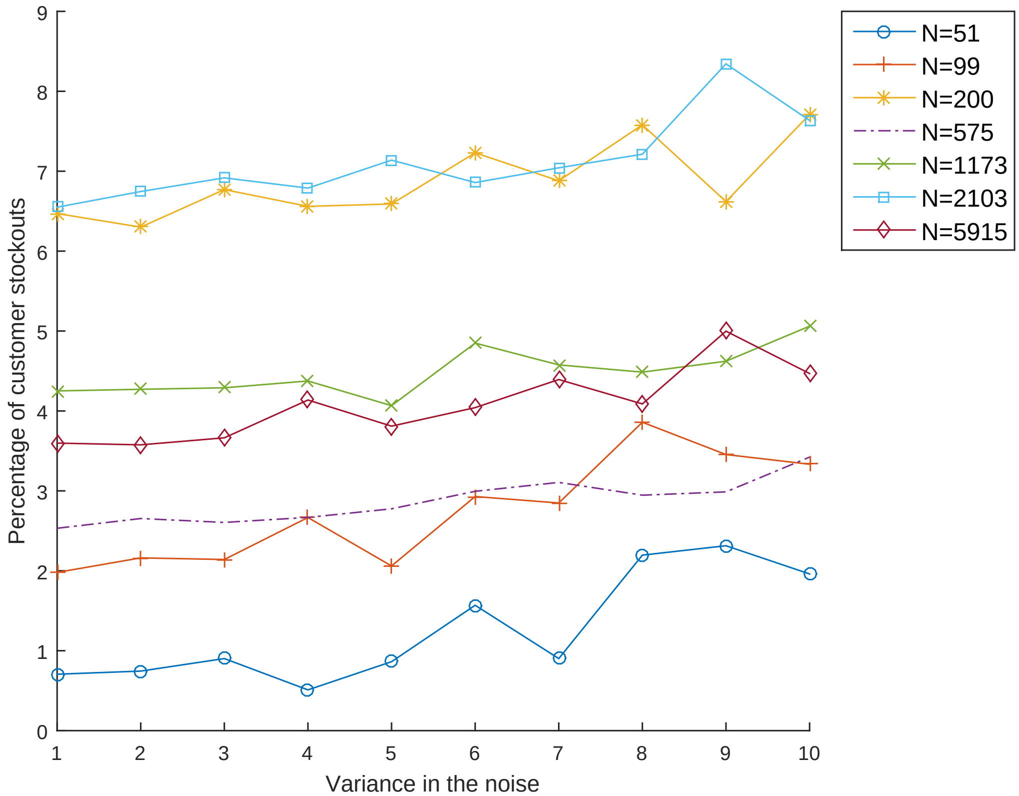

Figure 3 **Average stockout percentages for the robust solutions for data sets of different sizes.**

a customer stocks out. We assume that due to the reduced efficiency of the smaller emergency refuelling vehicle, its cost per unit distance is twice that of the usual refuelling vehicle fleet.

Table 1 compares the combined cost for the respective models, along with the percentage gap compared to the best lower bound the solver could find within the time limits we set. C_N, C_R and C_A are the combined costs of the nominal, robust and adaptive models respectively, while G_N, G_R and G_A are the respective provable duality gaps output by the Gurobi solver.

In all cases, the robust model had a combined cost no higher than 86% of the nominal model's. With larger data sets of over a hundred customers, the cost savings were 44% or more of the combined cost of the nominal model. The adaptive model had a combined cost that was a further 0.2%-0.3% lower than that of the robust model, i.e., an additional 0.1%-0.3% decrease to the combined cost of the nominal model beyond the improvements from going to the robust model.

To ensure that the solver gaps were not indicative of any major problems, we allowed the two smallest nominal models to run for 24 hours. At that point, the optimality gaps decreased to below 5%, with no further change to the solutions, showing that the solution found by the solver was

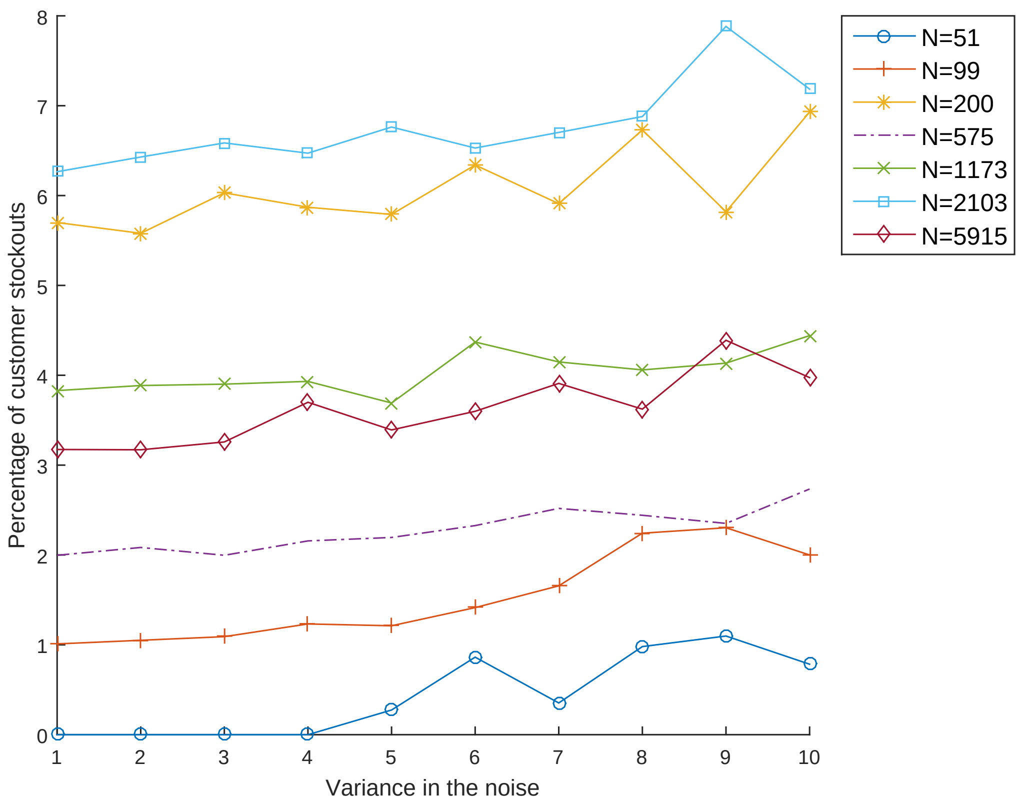

Figure 4 **Average stockout percentages for the adaptive solutions for data sets of different sizes.**

indeed close to the true optimal solution, and also provides us with improved bounds for the robust and adaptive models. Similarly, no improvements were made to the larger cases after 24 hours of running time, though in these cases the optimality gaps did not decrease sufficiently to draw the same conclusion.

Customers	C_N	G_N	C_R	G_R	C_A	G_A	C_R/C_N	C_A/C_N
51	16032	26.9	13791	37.2	13768	56.6	0.860	0.859
99	83527	45.7	66894	65.1	66684	66.5	0.801	0.798
200	5.295e6	77.3	2.966e6	74.3	2.959e6	75.5	0.560	0.559
575	1.001e6	1.61	459040	16.4	457547	18.9	0.459	0.457
1173	1.474e7	36.2	8.003e6	46.1	7.986e6	47.9	0.543	0.542
2103	3.190e7	31.8	1.373e7	38.0	1.370e7	39.7	0.431	0.429
5915	3.946e8	35.6	1.676e8	43.1	1.670e8	44.7	0.425	0.423

Table 1 **Costs and solver gaps for data sets of different sizes**

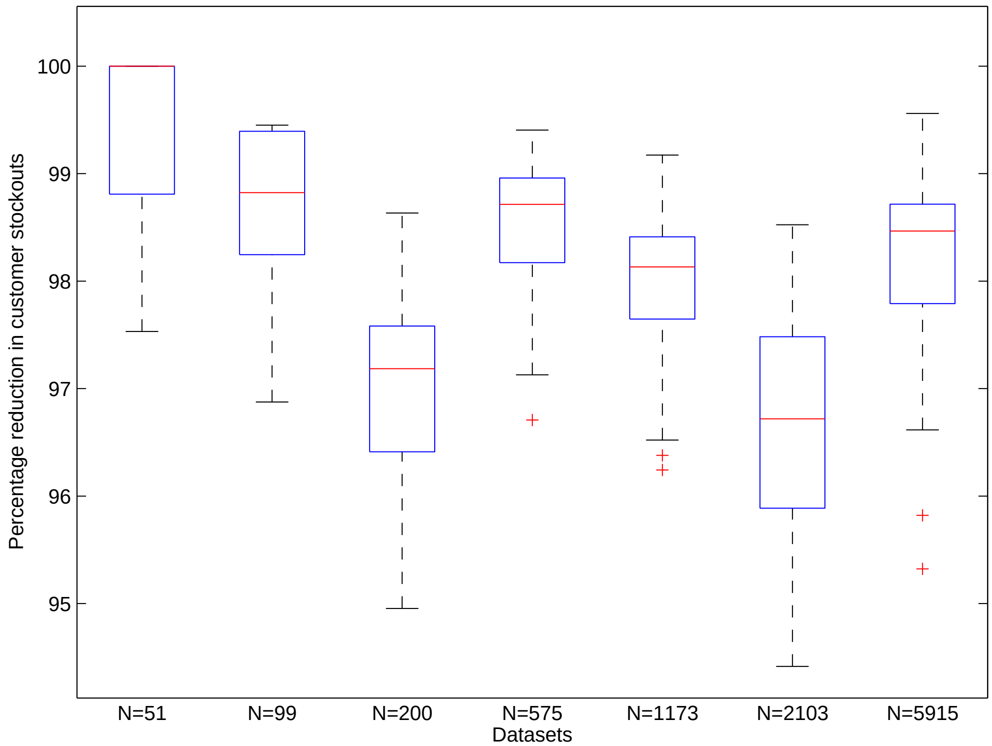

Figure 5 **25%-75% quantiles and extreme values for the reduction in robust model stockouts as a fraction of nominal model stockouts.**

5.3. Fleet reduction

Finally, whereas in the previous subsections, the vehicle fleet size was constant for each data set, here we investigate the tradeoffs of reducing the vehicle fleet size. We focus on a single data set with $N = 575$, for which our previous experiments used a fleet of 11 vehicles.

To allow the models to output a solution even with an infeasibly small fleet size, we introduce slack variables into our model that allow the demand constraints to be relaxed for a steep penalty (we took this to be 10^7 times the amount of violation). Taking the combined cost introduced in Section 5.2, we now further add to this the fixed cost of a vehicle fleet of a given size, taken to be 10,000 per vehicle. Table 2 compares the new combined cost for the best solutions for vehicle fleets of different sizes. C_N, C_R and C_A are the combined costs of the nominal, robust and adaptive models respectively, while S_N, S_R and S_A are the average number of customers who stock out.

We observe that the robust and adaptive solutions allow us to decrease the fleet size to 8 without increasing stockouts, and so decrease the combined cost. Decreasing the fleet size below 8 leads to an increase in combined cost for the robust and adaptive models, as demand is shifted from scheduled refuellings to emergency refuellings. On the other hand, with the nominal model, removing even

Vehicles	C_N	S_N	C_R	S_R	C_A	S_A
11	959490	852.9	529377	15.12	527747	12.38
10	1013552	932.54	519377	15.12	517747	12.38
9	1003552	932.54	509377	15.12	507747	12.38
8	993552	932.54	499377	15.12	497747	12.38
7	983552	932.54	556791	110.00	548587	95.04
6	991370	976.52	546791	110.00	538587	95.04
5	981370	976.52	1493560	1888.88	1493049	1860.24
4	1458104	1823.82	1483560	1888.88	1482509	1885.62
3	2157939	3233.10	2040526	2998.22	2038394	2994.58

Table 2 Stockout percentages for $N = 575$ with different fleet sizes

one vehicle leads to an increased combined cost, as the savings from the smaller fleet size are lost to increased refuelling costs from the increased numbers of customers who stock out.

With five vehicles or fewer, the robust and adaptive solutions are unable to find high-quality solutions; because our penalty is applied to the total unmet demand, these models minimize this by spreading the shortfall out over a large number of customers, suggesting that increasing the fleet size is crucial to reduce the emergency refuellings - in the worst case, we are experiencing over five times the number of stockouts as we have customers. At this point, our fleet sizes are highly infeasible for the models, and a significant part of the "cost" is from the penalty from the slack variables.

However, with at least six vehicles, we get not only a significant cost decrease, but we also observe that the adaptive solution has 81-86% of the number of stockouts that the robust solution has.

5.4. Summary of computational findings

We first found that across a variety of data sets ranging from 51 to 5915 customers, using the robust formulation led to an average reduction in stockouts of over 96% from the nominal formulation, and a further 5% decrease was attained by the adaptive formulation.

Next, we fixed the fleet size and considered the costs of the different formulations, namely regular routing costs and emergency refuelling. We found that the robust and adaptive models saved at least 14% in all cases, and at least 56% for larger cases (over 100 customers).

Finally, we considered a single data set, and looked more closely at whether we were able to reduce the fleet size without increasing stockouts. We found that while the nominal formulation was unable to do so, the robust and adaptive formulations were both able to maintain the same level of stockouts while reducing the fleet by three vehicles.

We also ran experiments to check the simple route generation heuristics that we used, and found that for our data sets, a small number of routes covering each customer was sufficient for our model to be able to find high-quality solutions.

6. Conclusion

We have presented robust and adaptive formulations for the finite horizon inventory routing problem that are tractable for ~6000 customers. For an uncertainty set where customers' demands demonstrate limited dependence, where the usual methods of robust optimization are insufficient, we have constructed an algorithm that allows us to find worst-case scenarios deterministically (robust formulation) or relative to a candidate solution (adaptive formulation). We have shown a significant decrease in stockouts (over 94% in all test cases) for our models, translating to a 14% decrease in cost for the supplier. In addition, we have shown that our models, with slack variables, are capable of providing further cost savings through a reduction in the vehicle fleet size.

While our work here has been in the context of a heating oil problem, it is applicable more broadly to other problems where the customer demand satisfied by a Vendor-Managed Inventory paradigm can be modeled by a tractable uncertainty set. Such problems might include beverages in vending machines, or more recently, bike-sharing in cities, where demand is dependent on temperature.

We would also like to explore possible improvements in the provable lower bounds on our solutions, for example along the lines of Bertsimas and de Ruiter (2016). Other improvements include more sophisticated ways of modelling emergency refuelling routing decisions, and smarter ways of managing fuel quantities dynamically.

Acknowledgments

We wish to thank Lorden Oil Co for much help in inspiring our interest in this problem, and providing us with data from which our experiment parameters were derived. We also wish to thank the Associate Editor and referees for their thorough and extensive comments on an earlier version of this paper, which were constructive and of great value to us.

Appendix A: Data

Let the number of customers in the data set be N and let each customer have a capacity of $Q = 20$. We consider the length of the planning horizon, T, to be 151 time units. However, to account for end-of-horizon effects, in our experiments we solve the model for $T = 151$, but only calculate costs for the first 141 time units. We assume a homogenous fleet of vehicles, each with capacity $S = 200$.

1. *Estimated initial amounts*: We generate the *estimated* initial level of oil for each customer $i \in N$ using the following formula:

$$z_i^{est} = Q \times (1 - \min(0.9, |X_i|/3)),$$

   where X_i are i.i.d. standard normal random variables sampled once for each customer. We generate z_i^{est} once for each customer for all the training scenarios, and once for each customer for all the testing scenarios.

2. *Realized initial amounts*: Once the estimated initial amounts are fixed, we generate *actual* customer levels at the start of the horizon, called z_i. These are generated with randomness proportional to the estimated amount already consumed. More precisely, for each scenario, we generate the initial customer level using:

$$z_i = z_i^{est} + (Q - z_i^{est}) \times U_i,$$

where U_i are i.i.d. uniform random variables distributed as $U_i \sim U(-1/2, 1/2)$. We finally clip the z_i within the interval $[0.5, Q]$.

3. *Estimated Temperature*: We set the base temperature, T_{base}, to $70°F$. Estimated temperature for day $t \in [T]$ is computed as:

$$T_t^{est} = T_{base} - 5 - t * 0.2.$$

4. *Realized Temperature*: We consider different scenarios with the noise in temperature, δ_t, varying in the set $\{0.02, 0.04, 0.06, \ldots, 1.0\}$. For each value for δ_t, we create instances with temperature generated using the following relation:

$$T_t = T_t^{est} + \max(-3 * \delta_t, \min(3 * \delta_t, \delta_t * X_t)).$$

5. *Fleet Size*: For datasets with less than 1000 customers, we assume a fleet with approximately $\sqrt{N}/2$ vehicles. Otherwise, we assume approximately $3D/ST$ vehicles, where D is the total mean demand across all the customers over the entire planning period, S is the vehicle capacity and T is the planning period. For data sets of size 51, 99, 200, 575, 1173, 2103 and 5915, this means a fleet size of 3, 4, 7, 11, 24, 42 and 117 vehicles respectively. These numbers were chosen based on the total average demand and vehicle capacity.

6. *Routes*: We start with an automatically set of generated routes such that for datasets with 51, 99 and 200 customers we have 4 routes covering each customer and for the larger datasets, where clusters are more stable, we have only 1 route covering each customer. The characteristics of the routes used in our experiments are included in Table 3, showing the number of routes (`Num Routes`), Minimum cost of the routes (`Min cost`), Maximum cost of the routes (`Max cost`), Average cost of the routes (`Avg cost`), number of routes covering each customer (`Per cust`), Minimum route size (`Min size`), Maximum route size (`Max size`) and Average route size (`Avg size`).

Dataset	Num routes	Min cost	Max cost	Avg cost	Per cust	Min size	Max size	Avg size
51	15	113.36	234.08	153.57	4	10	20	13.6
99	27	198.1877	521.12	297.89	4	11	23	14.67
200	55	2454.26	10200.675	5216.824	4	11	24	14.54
575	18	406.06	903.31	709.18	1	27	36	31.944
1173	36	212.18	628.137	392.403	1	12	37	32.58
2103	58	196.8	764.52	422.35	1	8	46	36.25
5915	174	3585.996	40454.5	17088.686	1	16	42	33.99

Table 3 Characteristics and coverage of the routes used in the experiments for various datasets.

References

Adelman, Daniel. 2004. A price-directed approach to stochastic inventory/routing. *Operations Research* **52**(4) 499–514.

Aghezzaf, El-Houssaine. 2007. Robust distribution planning for supplier-managed inventory agreements when demand rates and travel times are stationary. *Journal of the Operational Research Society* **59**(8) 1055–1065.

Archetti, Claudia, Luca Bertazzi, Gilbert Laporte, Maria Grazia Speranza. 2007. A branch-and-cut algorithm for a vendor-managed inventory-routing problem. *Transportation Science* **41**(3) 382–391.

Archetti, Claudia, Nicola Bianchessi, Stefan Irnich, M Grazia Speranza. 2014. Formulations for an inventory routing problem. *International Transactions in Operational Research* **21**(3) 353–374.

Archetti, Claudia, M Grazia Speranza. 2014. A survey on matheuristics for routing problems. *EURO Journal on Computational Optimization* **2**(4) 223–246.

Augerat, Philippe. 1995. Polyhedral study of the capacitated vehicle routing. Ph.D. thesis, Universite Joseph Fourrier, Grenoble.

Baldacci, Roberto, Nicos Christofides, Aristide Mingozzi. 2008. An exact algorithm for the vehicle routing problem based on the set partitioning formulation with additional cuts. *Mathematical Programming* **115**(2) 351–385.

Baldacci, Roberto, Eleni Hadjiconstantinou, Aristide Mingozzi. 2004. An exact algorithm for the capacitated vehicle routing problem based on a two-commodity network flow formulation. *Operations Research* **52**(5) 723–738.

Baldacci, Roberto, Aristide Mingozzi, Roberto Roberti. 2011. New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research* **59**(5) 1269–1283.

Baldacci, Roberto, Paolo Toth, Daniele Vigo. 2010. Exact algorithms for routing problems under vehicle capacity constraints. *Annals of Operations Research* **175**(1) 213–245.

Bandi, Chaithanya, Dimitris Bertsimas. 2012. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming* **134**(1) 23–70.

Ben-Tal, Aharon, Laurent El Ghaoui, Arkadi Nemirovski. 2009. *Robust Optimization*. Princeton University Press.

Ben-Tal, Aharon, Alexander Goryashko, Elana Guslitzer, Arkadi Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* **99**(2) 351–376.

Ben-Tal, Aharon, Arkadi Nemirovski. 1999. Robust solutions of uncertain linear programs. *Operations Research Letters* **25**(1) 1–13.

Bertazzi, Luca, Adamo Bosco, Demetrio Laganà. 2015. Managing stochastic demand in an inventory routing problem with transportation procurement. *Omega* **56** 112–121.

Bertazzi, Luca, Adamo Bosco, Demetrio Laganà. 2016. Min-max exact and heuristic policies for a two-echelon supply chain with inventory and transportation procurement decisions. *Transportation Research Part E: Logistics and Transportation Review* **93**(Supplement C) 57 – 70. doi:https://doi.org/10.1016/j.tre.2016.05.008. URL `http://www.sciencedirect.com/science/article/pii/S1366554515302933`.

Bertazzi, Luca, Giuseppe Paletta, M Grazia Speranza. 2002. Deterministic order-up-to level policies in an inventory routing problem. *Transportation Science* **36**(1) 119–132.

Bertazzi, Luca, Martin Savelsbergh, Maria Grazia Speranza. 2008. Inventory routing. *The vehicle routing problem: latest advances and new challenges*. Springer, 49–72.

Bertsimas, Dimitris, David B Brown, Constantine Caramanis. 2011a. Theory and applications of robust optimization. *SIAM Review* **53**(3) 464–501.

Bertsimas, Dimitris, Constantine Caramanis. 2010. Finite adaptability in multistage linear optimization. *Automatic Control, IEEE Transactions on* **55**(12) 2751–2766.

Bertsimas, Dimitris, Frans JCT de Ruiter. 2016. Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS Journal on Computing* **28**(3) 500–511.

Bertsimas, Dimitris, Vineet Goyal, Xu Andy Sun. 2011b. A geometric characterization of the power of finite adaptability in multistage stochastic and adaptive optimization. *Mathematics of Operations Research* **36**(1) 24–54.

Bertsimas, Dimitris, Vishal Gupta, Nathan Kallus. 2018. Data-driven robust optimization. *Mathematical Programming* **167**(2) 235–292.

Bertsimas, Dimitris, Dan A Iancu, Pablo A Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research* **35**(2) 363–394.

Bertsimas, Dimitris, Eugene Litvinov, Xu Andy Sun, Jinye Zhao, Tongxin Zheng. 2013. Adaptive robust optimization for the security constrained unit commitment problem. *Power Systems, IEEE Transactions on* **28**(1) 52–63.

Bertsimas, Dimitris, Melvyn Sim. 2003. Robust discrete optimization and network flows. *Mathematical Programming* **98**(1) 49–71.

Bertsimas, Dimitris, Melvyn Sim. 2004. The price of robustness. *Operations Research* **52**(1) 35–53.

Bertsimas, Dimitris J. 1992. A vehicle routing problem with stochastic demand. *Operations Research* **40**(3) 574–585.

Bertsimas, Dimitris J, David Simchi-Levi. 1996. A new generation of vehicle routing research: robust algorithms, addressing uncertainty. *Operations Research* **44**(2) 286–304.

Birge, John R, Francois Louveaux. 2011. *Introduction to stochastic programming*. Springer Science & Business Media.

Cacchiani, Valentina, VC Hemmelmayr, Fabien Tricoire. 2014. A set-covering based heuristic algorithm for the periodic vehicle routing problem. *Discrete Applied Mathematics* **163** 53–64.

Campbell, Ann, Lloyd Clarke, Anton Kleywegt, Martin Savelsbergh. 1998. The inventory routing problem. *Fleet management and logistics*. Springer, 95–113.

Chepuri, Krishna, Tito Homem-De-Mello. 2005. Solving the vehicle routing problem with stochastic demands using the cross-entropy method. *Annals of Operations Research* **134**(1) 153–181.

Christofides, Nicos, Samuel Eilon. 1969. An algorithm for the vehicle-dispatching problem. *OR* 309–318.

Christofides, Nicos, Aristide Mingozzi, Paolo Toth. 1981. Exact algorithms for the vehicle routing problem, based on spanning tree and shortest path relaxations. *Mathematical Programming* **20**(1) 255–282.

Coelho, Leandro C, Jean-François Cordeau, Gilbert Laporte. 2013. Thirty years of inventory routing. *Transportation Science* **48**(1) 1–19.

Cordeau, Jean-François, Michel Gendreau, Gilbert Laporte. 1997. A tabu search heuristic for periodic and multi-depot vehicle routing problems. *Networks* **30**(2) 105–119.

Cordeau, Jean-François, Gilbert Laporte, Martin WP Savelsbergh, Daniele Vigo. 2006. Vehicle routing. *Transportation, handbooks in operations research and management science* **14** 367–428.

Croes, Georges A. 1958. A method for solving traveling-salesman problems. *Operations Research* **6**(6) 791–812.

Dantzig, George B, John H Ramser. 1959. The truck dispatching problem. *Management Science* **6**(1) 80–91.

Delage, Erick, Yinyu Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* **58**(3) 595–612.

Disney, Stephen M, Denis R Towill. 2003. The effect of vendor managed inventory (vmi) dynamics on the bullwhip effect in supply chains. *International journal of production economics* **85**(2) 199–215.

Dror, Moshe. 2005. Routing propane deliveries. *Logistics Systems: Design and Optimization* 299–322.

Dror, Moshe, Michael Ball. 1987. Inventory/routing: Reduction from an annual to a short-period problem. *Naval Research Logistics (NRL)* **34**(6) 891–905.

Dror, Moshe, Michael Ball, Bruce Golden. 1985. A computational comparison of algorithms for the inventory routing problem. *Annals of Operations Research* **4**(1) 1–23.

Dror, Moshe, Gilbert Laporte, Pierre Trudeau. 1989. Vehicle routing with stochastic demands: Properties and solution frameworks. *Transportation science* **23**(3) 166–176.

Dunning, Iain, Joey Huchette, Miles Lubin. 2017. Jump: A modeling language for mathematical optimization. *SIAM Review* **59**(2) 295–320. doi:10.1137/15M1020575.

Federgruen, Awi, David Simchi-Levi. 1995. Analysis of vehicle routing and inventory-routing problems. *Handbooks in operations research and management science* **8** 297–373.

Federgruen, Awi, Paul Zipkin. 1984. A combined vehicle routing and inventory allocation problem. *Operations Research* **32**(5) 1019–1037.

Francis, Peter M, Karen R Smilowitz, Michal Tzur. 2008. The period vehicle routing problem and its extensions. *The vehicle routing problem: latest advances and new challenges*. Springer, 73–102.

Fukasawa, Ricardo, Humberto Longo, Jens Lysgaard, Marcus Poggi de Aragão, Marcelo Reis, Eduardo Uchoa, Renato F Werneck. 2006. Robust branch-and-cut-and-price for the capacitated vehicle routing problem. *Mathematical Programming* **106**(3) 491–511.

Gendreau, Michel, Alain Hertz, Gilbert Laporte. 1994. A tabu search heuristic for the vehicle routing problem. *Management Science* **40**(10) 1276–1290.

Gendreau, Michel, Gilbert Laporte, René Séguin. 1996. Stochastic vehicle routing. *European Journal of Operational Research* **88**(1) 3–12.

Golden, Bruce L, Subramanian Raghavan, Edward A Wasil. 2008. *The Vehicle Routing Problem: Latest Advances and New Challenges: latest advances and new challenges*, vol. 43. Springer Science & Business Media.

Goldfarb, Donald, Garud Iyengar. 2003. Robust portfolio selection problems. *Mathematics of Operations Research* **28**(1) 1–38.

Gounaris, Chrysanthos E, Wolfram Wiesemann, Christodoulos A Floudas. 2013. The robust capacitated vehicle routing problem under demand uncertainty. *Operations Research* **61**(3) 677–693.

Gunnarsson, Helene, Mikael Rönnqvist, Dick Carlsson. 2006. A combined terminal location and ship routing problem. *Journal of the Operational Research Society* **57**(8) 928–938.

Hall, Justin. 2016. How the cost of crude oil affects truck freight rates. URL `http://shiplps.com/how-the-cost-of-crude-oil-affects-truck-freight-rates/`.

Hemmelmayr, Vera, Karl F Doerner, Richard F Hartl, Martin WP Savelsbergh. 2009a. Delivery strategies for blood products supplies. *OR spectrum* **31**(4) 707–725.

Hemmelmayr, Vera, Karl F Doerner, Richard F Hartl, Martin WP Savelsbergh. 2010. Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research* **202**(3) 686–695.

Hemmelmayr, Vera C, Karl F Doerner, Richard F Hartl. 2009b. A variable neighborhood search heuristic for periodic routing problems. *European Journal of Operational Research* **195**(3) 791–802.

Irnich, Stefan, Paolo Toth, Daniele Vigo. 2014. The family of vehicle routing problems. *Vehicle Routing: Problems, Methods, and Applications* **18** 1.

Kleywegt, Anton J, Vijay S Nori, Martin WP Savelsbergh. 2002. The stochastic inventory routing problem with direct deliveries. *Transportation Science* **36**(1) 94–118.

Laporte, Gilbert. 1992. The vehicle routing problem: An overview of exact and approximate algorithms. *European Journal of Operational Research* **59**(3) 345–358.

Laporte, Gilbert. 2009. Fifty years of vehicle routing. *Transportation Science* **43**(4) 408–416.

Laporte, Gilbert, Hélène Mercure, Yves Nobert. 1986. An exact algorithm for the asymmetrical capacitated vehicle routing problem. *Networks* **16**(1) 33–46.

Laporte, Gilbert, Yves Nobert, Martin Desrochers. 1985. Optimal routing under capacity and distance restrictions. *Operations Research* **33**(5) 1050–1073.

Liu, Bing. 1997. Route finding by using knowledge about the road network. *Systems, Man and Cybernetics, Part A: Systems and Humans* **27**(4) 436–448.

Lorca, Alvaro, X Andy Sun, Eugene Litvinov, Tongxin Zheng. 2016. Multistage adaptive robust optimization for the unit commitment problem. *Operations Research* **64**(1) 32–51.

Lysgaard, Jens, Adam N Letchford, Richard W Eglese. 2004. A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming* **100**(2) 423–445.

Ordóñez, Fernando. 2010. Robust vehicle routing. *TUTORIALS in Operations Research* 153–178.

Pacheco, Joaquín, Ada Alvarez, Irma García, Francisco Angel-Bello. 2012. Optimizing vehicle routes in a bakery company allowing flexibility in delivery dates. *Journal of the Operational Research Society* **63**(5) 569–581.

Pecin, Diego, Artur Pessoa, Marcus Poggi, Eduardo Uchoa. 2014. Improved branch-cut-and-price for capacitated vehicle routing. *Integer programming and combinatorial optimization*. Springer, 393–403.

Prescott-Gagnon, Eric, Guy Desaulniers, Louis-Martin Rousseau. 2014. Heuristics for an oil delivery vehicle routing problem. *Flexible Services and Manufacturing Journal* **26**(4) 516–539.

Ralphs, Ted K, Leonid Kopman, William R Pulleyblank, Leslie E Trotter. 2003. On the capacitated vehicle routing problem. *Mathematical Programming* **94**(2-3) 343–359.

Sahin, Funda, E. Powell Robinson. 2005. Information sharing and coordination in make-to-order supply chains. *Journal of Operations Management* **23**(6) 579 – 598. doi:https://doi.org/10.1016/j.jom.2004.08.007. URL `http://www.sciencedirect.com/science/article/pii/S0272696304001159`.

Semet, Frédéric, Eric Taillard. 1993. Solving real-life vehicle routing problems efficiently using tabu search. *Annals of Operations Research* **41**(4) 469–488.

Shapiro, Alexander, Darinka Dentcheva, et al. 2014. *Lectures on stochastic programming: modeling and theory*, vol. 16. SIAM.

Solyali, Oguz, Jean-François Cordeau, Gilbert Laporte. 2012. Robust inventory routing under demand uncertainty. *Transportation Science* **46**(3) 327–340.

Sungur, Ilgaz, Fernando Ordónez, Maged Dessouky. 2008. A robust optimization approach for the capacitated vehicle routing problem with demand uncertainty. *IIE Transactions* **40**(5) 509–523.

Toth, Paolo, Daniele Vigo. 2001. *The vehicle routing problem*. Society for Industrial and Applied Mathematics.

Toth, Paolo, Daniele Vigo. 2014. *Vehicle Routing: Problems, Methods, and Applications*, vol. 18. SIAM.

Vidal, Thibaut, Teodor Gabriel Crainic, Michel Gendreau, Nadia Lahrichi, Walter Rei. 2012. A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research* **60**(3) 611–624.

Waller, Matt, M Eric Johnson, Tom Davis. 1999. Vendor-managed inventory in the retail supply chain. *Journal of business logistics* **20**(1) 183.

Weather Underground. 2014. Weather history for KBOS. URL `http://www.wunderground.com/history/airport/KBOS/2013/11/1/CustomHistory.html?dayend=31&monthend=3&yearend=2014&req_city=&req_state=&req_statename=&reqdb.zip=&reqdb.magic=&reqdb.wmo=`.