Submitted to *Transportation Science*
manuscript 0

Probabilistic Traveling Salesman Problem with Deadlines

Ann M. Campbell, Barrett W. Thomas
Department of Management Sciences, University of Iowa
108 John Pappajohn Business Building, Iowa City, IA 52242-1994, USA
ann-campbell@uiowa.edu, barrett-thomas@uiowa.edu

Time-constrained deliveries are one of the fastest growing segments of the delivery business, and yet, there is surprisingly little literature that addresses time constraints in the context of stochastic customer presence. We begin to fill that void by introducing the probabilistic traveling salesman problem with deadlines (PTSPD). The PTSPD is an extension of the well-known probabilistic traveling salesman problem in which, in addition to stochastic presence, customers must also be visited before a known deadline. We present two recourse models and a chance constrained model for the PTSPD. Special cases are discussed for each model, and computational experiments are used to illustrate under what conditions stochastic and deterministic models lead to different solutions.

Key words: traveling salesman problem; probabilistic; deadlines

1. Introduction

In recent years, there has been tremendous growth in time-definite delivery services for full-truckload, less-than-truckload (LTL), and small package delivery companies (Foster, 1999, U.S. Department of Transportation Federal Highway Administration, 2004, Schulz, 2003, Shanahan, 2003). The most common example of these services is next-day and second-day package delivery featured by United Parcel Service (UPS) and FedEx. Next-day delivery providers usually offer a choice of deadlines such as 10 am, noon, or 3 pm. For the most recent years for which data is available, freight volume associated with time-definite services was growing at 12% per year (Foster, 1999), and in 2003, next-day delivery was the fastest growing segment of the LTL trucking market (Schulz, 2003). The Colography Group, Inc., the industry leader in trend forecasting in transportation and logistics, expects this growth in time-definite services to continue (Scherck, 2003). Group President Theodore Scherck says 95% of shippers in his company's surveys cite "on-time" delivery as the most important factor in carrier selection. Scherck explains, "More than ever, transport in the U.S., and abroad as well, reflects modal neutrality and an emphasis on time-definite services."

Because of limited enabling technology and high cost, many delivery companies employ a pre-planned, or a priori, route which specifies an ordering of all possible customers that a particular driver may need to visit. The driver then skips those customers on the route who do not receive a delivery. In addition to being easily implementable, these a priori routes offer both drivers and customers consistency and help to improve driver efficiency as the driver becomes increasingly familiar with the route. With time-definite services becoming a larger part of the delivery business, companies have been developing ad-hoc measures to accommodate deadlines in their a priori tours. Yet, the explicit consideration of deadlines into a priori route design has the potential to not only reduce transportation costs, but also the penalty costs delivery service providers incur for late deliveries. Consequently, it is important to re-examine the design of a priori delivery routes in light of these time constraints. However, the consideration of time constraints in a priori route construction is strikingly absent in the academic literature.

In this paper, we take the first steps in incorporating time constraints into a priori routes by introducing the probabilistic traveling salesman problem with deadlines (PTSPD). While a number of elements of this problem could be modeled stochastically, such as travel time, we maintain the presence of the customer as the only stochastic element in our model, as in the well-known probabilistic traveling salesman problem (PTSP). As a result, the PTSPD can be considered an extension of the PTSP. The PTSP is the problem of finding a minimum expected cost a priori tour through a set of customers $N = \{i \mid 1, \ldots, n\}$ with probabilities $P = \{p_i \mid 1, \ldots, n\}$ of requiring service on any given day. The travel time between any two customers i and j is given by d_{ij}, where $d_{ij} = d_{ji}$. These travel times also serve as the cost to traverse each arc. In the PTSPD, also associated with each customer $i \in N$ is a known deadline l_i. Service at each customer should begin at or before its deadline. Since most delivery service providers offer time-definite services based on deadlines, such as FedEx's 10 am delivery option, this approach captures the time constraints faced by most companies.

As will be demonstrated in this paper, deadlines present particularly challenging and unique modeling issues when considered in a stochastic context. This paper's primary contribution is a series of models representing different ways in which deadline violations can be measured and addressed in a stochastic environment. These models require non-trivial derivations to calculate the probability of late arrivals at customers. For each model, we also identify special cases. Special cases are types of instances whose structure leads to reduced computational complexity relative to the general problem. The first two models are recourse models, while the third is a chance constrained approach (see Birge and Louveaux (1997) and Charnes and Cooper (1959, 1963) for an overview of the two types of formulations). Both recourse and chance constrained models have been considered for other stochastic routing problems, but not in conjunction with both time constraints and stochastic customer presence. Recourse models account for deadline violations by penalizing any violations in the objective. In contrast, chance constrained models restrict the probability that a deadline violation can occur. We also present the results of computational experiments which identify the circumstances under which PTSPD solutions differ from their deterministic analogs.

This paper is organized as follows. In Section 2, we survey related literature. Sections 3, 4, and 5 each present a general model for the PTSPD as well as a discussion of special cases for each. Sections 3 and 4 are recourse models with different penalty functions, where Section 5 involves a chance constrained model for the PTSPD. Section 6 introduces an example that illustrates the behavior of the proposed models, and Section 7 includes the computational experiments. Finally, Section 8 summarizes our work and offers directions for future research into the PTSPD.

2. Literature Review

The PTSPD has its roots in both the traveling salesman problem with time windows (TSPTW) and the stochastic vehicle routing literature. The TSPTW is a deterministic analog of the PTSPD and has been widely studied in the literature. Solution approaches for the TSPTW range from exact mathematical programming techniques to various heuristic approaches. Christofides et al. (1981) and Baker (1983) present branch-and-bound algorithms that solve problems with up to 50 vertices, but require "moderately tight" time windows and/or little overlap between them. Langevin et al. (1993) introduce a two-commodity flow formulation well-suited to handling time windows; they solve instances with up to 40 nodes. Dumas et al. (1995) extend earlier dynamic programming approaches by using state space reduction techniques that enable the solution of problems with up to 200 customers. In an alternate approach, Pesant et al. utilize constraint programming to develop an exact method (Pesant et al., 1998) and a heuristic (Pesant et al., 1999) for the TSPTW. Similarly, Focacci et al. (2002) embed optimization techniques within a constraint programming approach.

Because of limitations in solving exact formulations (Savelsbergh (1985) proves that even finding a feasible solution to the TSPTW is an NP-complete problem), many authors focus on heuristic techniques for the TSPTW. Carlton and Barnes (1996) solve the TSPTW with a tabu search approach that considers infeasible solutions in its search neighborhood through the implementation of a static penalty function. Gendreau et al. (1998) offer a construction and post-optimization heuristic based on a near-optimal traveling salesman problem (TSP) heuristic presented by Gendreau et al. (1992). Wolfler Calvo (2000) introduces a heuristic that constructs an initial tour using a unique relaxation to the assignment problem. Ohlmann and Thomas (to appear) apply a variant of simulated annealing, compressed annealing, to the TSPTW and find many new best solutions.

Gendreau et al. (1995b) introduce the traveling salesman problem with deadlines for which the objective is to maximize the number of customers served before a deadline. They characterize a series of special cases and introduce an enumerative approach that solves problems up to 100 customers.

In contrast to these deterministic problems, the research presented in this paper focuses on a problem in which customer presence on the tour is random. In an early treatment of stochastic routing in this context, Bartholdi et al. (1983) introduce a space-filling curve heuristic for constructing a priori tours for meals-on-wheels routing. Jaillet (1985, 1988) introduces an analytical framework for the PTSP. In addition, Jaillet demonstrates that an optimal solution to the deterministic traveling salesman problem (TSP) may not be the best solution when customer presence is stochastic. Laporte et al. (1994) provide an exact algorithm for the PTSP. However, the exact approach is limited to small problem sizes. Consequently, much of the PTSP literature focuses on heuristic approaches. Bertsimas et al. (1990) discuss space-filling curve and iterative heuristics. Bertsimas and Howell (1993) and Chervi (1988) introduce equations for efficiently evaluating the cost of local-search moves for the PTSP. Bianchi et al. (2005) and Bianchi and Campbell (2007) provide corrections for the equations in Bertsimas and Howell and in Chervi, respectively. Recent work by Campbell (2006) and Tang and Miller-Hooks (2004) focuses on approximations for the PTSP. Overviews of the research in this area can be found in Powell et al. (1995), Bertsimas and Simchi-Levi (1996), and Gendreau et al. (1996).

While the literature contains research into many constrained versions of the TSP, we are aware of only one constrained variant of the PTSP, the stochastic vehicle routing problem (SVRP). The SVRP requires the consideration of vehicle capacity in the formation of the tours, and rather than customer presence, customer demand is usually the stochastic element of the problem. The first mention of this problem can be found in Tillman (1969). Bertsimas (1988, 1992) introduces an analytical framework and bounds for the SVRP. Other work can be divided into consideration of chance constrained and recourse model formulations. Stewart and Golden (1983), Laporte et al. (1989), and Bastian and Rinnooy Kan (1992) provide chance constrained formulations and demonstrate how they can be transformed into deterministic problems. Dror et al. (1989), Dror (1993), and Bastian and Rinnooy Kan (1992) present stochastic programming solutions to various recourse models for the SVRP. Many offer heuristics for the SVRP, including Dror and Trudeau (1986), Bramel et al. (1992), Bertsimas et al. (1995), Savelsbergh and Goetschalckx (1995), and Yang et al. (2000). Gendreau et al. (1995a) offer a stochastic programming approach for an SVRP variant in which both customer presence and customer demand is stochastic.

The authors know of only a few papers that address routing under uncertainty with time constraints. These papers consider time constraints in the context of stochastic travel times. Teng et al. (2004) apply the L-shaped algorithm to the time-constrained traveling salesman problem (TCTSP) with stochastic travel and service times. In the TCTSP, the time constraint is on the length of the tour, which contrasts with this paper where the time constraints control when individual customers can be visited. Wong et al. (2003) introduce a 2-stage stochastic integer program with recourse for a problem where customers have time windows and travel times are stochastic.

3. Recourse I

In next two sections, we present recourse models for the PTSPD. A recourse model is characterized by two stages. In stage one, an a priori solution is determined. Then, in the second stage, the random variables are realized, and a recourse action is applied to the first stage solution. In essence, the recourse actions are corrective measures taken as the result of any infeasibility that has arisen because of particular realizations of the random variables. These corrective measures result in added costs that are included in evaluating the expected cost of a proposed solution. Thus, the expected cost of a PTSPD tour when posed as a recourse model is the sum of the expected travel costs between customers and the expected costs associated with the given infeasibility penalties.

For the PTSPD, we present models for two possible recourse actions. For each, we present expressions that allow for efficient computation of the expected cost of a given tour (in the manner of (Jaillet, 1988)). In Recourse I, the vehicle is allowed to visit a customer after the delivery deadline has passed, but incurs a penalty for doing so. We consider both a per unit time of violation charge and a fixed charge for violating the deadline. The penalty for the per-unit-time charge is represented by a customer-dependent λ_i, and the fixed charge penalty by a customer-dependent ϕ_i. The per-unit-time charge represents cases where the delivery company is charged per unit time of lateness. For instance, FedEx Custom Critical refunds varying percentages of the cost of a shipment based on how late the shipment is delivered (FedEx, 2003). For additional examples, see (Charnsirisakskul et al., 2004) and (Slotnick and Sobel, 2005). The fixed-charge recourse represents the case where the delivery company reimburses the customer for the cost of the delivery in the event that the deadline is not met. Well-known examples of such penalties are FedEx's and UPS' money-back guarantees (FedEx, 2004, United Parcel Service, 2004).

In our discussion of recourse models, we let τ define an order, or tour, in which the customers $i = 1, \ldots, n$ are to be visited. We assume that the customers are indexed according to their position in τ. We assume that we always begin a tour at a fixed depot, and the depot is indexed as $i = 0$. Unless otherwise indicated, we assume integer travel times and deadlines and that all tours start at time $t = 0$. Our recourse models provide equations for evaluating the cost of a given tour τ.

3.1. General Model

As discussed previously, the Recourse I model visits each realized customer regardless of whether or not a deadline violation occurs. As a result, the expected travel costs can be calculated as they are for the well-known PTSP (see (Jaillet, 1988) for further reference) with a straightforward modification for the fixed depot:

$$\sum_{j=1}^{n} p_j d_{0,j} \prod_{k=1}^{j-1}(1-p_k) + \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} p_i p_j d_{i,j} \prod_{k=i+1}^{j-1} (1-p_k) + \sum_{i=1}^{n} p_i d_{i,0} \prod_{k=i+1}^{n} (1-p_k). \tag{1}$$

This equation calculates the probability and the resulting expected cost of any arc that may appear in the tour. The expected cost of an arc (i, j) depends on customers i and j being realized and no customers k, $k = i+1, \ldots, j-1$, being realized.

To account for deadline violations, we add a penalty term to equation 1. For both the per-unit-time and fixed-charge penalties, we determine the probability that a deadline violation occurs. To begin, let the random variable

$$X_i = \begin{cases} 0 \text{ if customer } i \text{ is not realized} \\ 1 \text{ if customer } i \text{ is realized.} \end{cases}$$

Also, let A_i be a random variable representing the time of arrival to customer i. We assume that the deadline at customer i cannot be violated if customer i is not realized. Hence, we are left to compute $P(A_i = t \mid X_i = 1)$, where t represents the time of arrival to customer i. For notational convenience, let $g(i,t) = P(A_i = t \mid X_i = 1)$. We recognize that, given customer i is realized, arrival to customer i is at time t only if arrival to the last visited customer, say h, was at time $t - d_{h,i}$, and all customers k between h and i are not realized. Then, for $t = d_{0i}$:

$$g(i,t) = \prod_{k=1}^{i-1}(1-p_k) + \sum_{h=1, t>d_{h,i}}^{i-1} p_h g(h, t-d_{h,i}) \prod_{k=h+1}^{i-1} (1-p_k),$$

otherwise

$$g(i,t) = \sum_{h=1, t>d_{h,i}}^{i-1} p_h g(h, t-d_{h,i}) \prod_{k=h+1}^{i-1} (1-p_k).$$

We note that $g(i,t)$ can be computed recursively for all $t > 0$.

In the case of the per-unit-time penalty, we add the following term to equation 1:

$$\sum_{i=1}^{n} p_i \sum_{t=l_i+1}^{T_i} \lambda_i g(i,t)(t-l_i), \tag{2}$$

where T_i is latest time that we could possibly arrive to customer i. For each customer $i, \ldots, n$, we can compute $\hat{T}_i = \sum_{j=0}^{i-1} d_{jj+1}$. That is, the latest that we can possibly arrive to i is the result of all customers from the depot to i being realized.

In the case of a fixed-charge penalty, we need to calculate $P(A_i \le t \mid X_i = 1)$, which is the probability that arrival to customer i occurs at or before time t given that customer i is realized. For convenience, we let $G(i,t) = P(A_i \le t \mid X_i = 1)$. Then, the probability that arrival to a realized customer i is before or at l_i, $P(A_i \le l_i \mid X_i = 1)$, is given by $G(i, l_i)$. The calculation of $G(i,t)$ is straightforward:

$$G(i,t) = \sum_{k=0}^{t} g(i,k) = G(i, t-1) + g(i,t). \tag{3}$$

For the fixed-charge penalty, we add the following term to equation 1:

$$\sum_{i=1}^{n} p_i \phi_i \bar{G}(i, l_i), \tag{4}$$

where $\bar{G}(i, l_i) = 1 - G(i, l_i)$.

Computation of the expected cost equations for Recourse I is more computationally complex than for the PTSP. When using a per-unit penalty, we must consider all arrival times that would violate a customer's deadline. Given the model described in equation 2, computing all of the necessary g values requires $O(n^2 \max_i\{T_i\})$ time. Without restrictions on distances, this is not necessarily polynomial in n. Once the g values are known, the per-unit penalty portion of the objective can be computed in $O(n \max_i\{T_i - l_i\})$ time where the distance portion remains $O(n^2)$. Thus, the complexity of the function evaluation is dominated by the computation of the g values and is $O(n^2 \max_i\{T_i\})$. With a fixed charge penalty, we only need to consider arrival times up to the largest l_i value, making the complexity $O(n^2 \max_i\{l_i\})$.

3.2. Special Cases for Recourse I

Because of the computational complexity of evaluating the expected cost of a tour, there is merit to exploring special cases and approximations that have reduced complexity. In future work, these special cases may form the basis for heuristic solutions to Recourse I. The remainder of this section introduces special cases of the general Recourse I model.

The set of special cases for Recourse I are based on assumptions regarding $p_i, \lambda_i, \phi_i, l_i$, and $d_{i,j}$. The first two special cases assume $d_{i,j} = 1$ for every i and j. The assumption that $d_{i,j} = 1$ makes all function evaluations polynomial in n. That is, $\max_i\{T_i\}$ becomes $O(n)$, and the expected cost evaluation becomes $O(n^3)$. The final special case makes that assumption that all customers are located on a line. Gendreau et al. (1995b) note that automated guided vehicle systems often exhibit this structure. For all three special cases, we characterize the optimal tour.

3.2.1. Unit Distance, Homogeneous p, Homogeneous l, Homogeneous λ (or ϕ) For the first special case, we assume that $p_i = p$ and $\lambda_i = \lambda$ for all i (or $\phi_i = \phi$), $d_{i,j} = 1$ for every i and j, and $l_i = l$ for every i. Under these assumptions, each customer is identical and hence every tour τ has the same expected cost.

3.2.2. Unit Distance, Homogeneous p, Homogeneous l For the second special case, we allow the penalty value λ_i (or ϕ_i) to be customer dependent, but maintain our assumptions $p_i = p$ for every i, $d_{i,j} = 1$ for every i and j, and $l_i = l$ for every i. Under these assumptions, we recognize that, if the delivery is late to one customer, it is going to be late to every subsequent customer on the tour. At the same time, because of the unit distance assumption, the computation of $G(i,t)$ depends only on how many customers precede customer i on the tour. Which customers precede i and in what order are now irrelevant. As i increases, the probability of violating the deadline at i increases and hence the probability of incurring the penalty λ_i (or ϕ_i) increases. Since the travel time between customers is irrelevant, the optimal tour is then a tour that orders the customers in decreasing order of λ_i (or ϕ_i).

3.2.3. Straight Line Distance For the third special case, we assume that we can order the customers $1,\ldots,n$ such that $d_{i,k} = d_{i,j} + d_{j,k}$ for $i < j < k$ and $d_{i,j} = d_{ji}$ for $i \neq j$. Given the assumptions on distance, the identity order, $1,\ldots,n$, is an optimal tour. Consider any non-identity tour τ. As a non-identity tour, τ contains at least one customer who is not in topological order. Let i be the first customer not in topological order in τ. Let τ' be the tour such that i is removed from its position in τ and inserted back into the tour so that i is in topological order.

Consider any realization of customer presence ω. If i is not present in the realization ω, then the cost of τ and τ' are the same for the realization ω. Suppose i is present in ω. Let m be the last realized customer in τ before i, and n be the first realized customer in τ after i. Let h be the last realized customer in τ' before i and j be the first realized customer in τ' after i.

Given these definitions, the change in cost for this switch from τ to τ' for realization ω is $d_{h,i} + d_{i,j} + d_{mn} - (d_{mi} + d_{in} + d_{hj})$. Given the construction of τ', $h < i < j$ and thus $d_{hj} = d_{h,i} + d_{i,j}$. The triangle equality assumption implies $d_{mn} \leq d_{mi} + d_{in}$. Consequently, for any realization of the customers, the cost of τ' is less than or equal to the travel cost of τ. Likewise, we arrive to each customer no later in τ' than in τ. Hence, the penalty cost is also not increased by the transformation to τ' from τ. Continuing the proposed tour improvement scheme, we can transform a tour τ into a lower expected cost tour τ' eventually resulting in the identity tour.

4. Recourse Model II

In this section, we introduce another reasonable approach to handling deadline violations. We assume that any customer whose deadline would be violated in the realization of the tour is skipped, with the delivery company incurring a fixed charge. This penalty represents the cost to service the violated customer with a separate vehicle. This recourse action represents the situation in which the company must serve all customers by their specified deadline. For each customer i, we represent this customer-dependent fixed charge by ϕ_i.

4.1. General Model

In the Recourse II model, a customer on the tour is skipped if:

1. the customer is not realized
2. the customer is realized, but arrival to the customer would occur after the customer's deadline.

Because of the two options, calculating the probability of any particular arc appearing on a tour is more complicated to derive than Recourse I. Notably, the probability of violating one customer's deadline is dependent on the order of previous customers in the tour, whether or not those customers are realized, and whether or not any of the previous customers are skipped. To understand the interactions among the potential scenarios, it is instructive to present the complete derivation of the expected cost equation for a given tour.

To begin, again let the random variable

$$X_i = \begin{cases} 0 \text{ if customer } i \text{ is not realized} \\ 1 \text{ if customer } i \text{ is realized.} \end{cases}$$

Also, define the random variable

$$Y_i = \begin{cases} 0 \text{ if the time of arrival to customer } i \text{ would be later than } l_i \\ 1 \text{ if the time of arrival to customer } i \text{ would be earlier than or equal to } l_i. \end{cases}$$

While our assumptions imply that the depot is always realized and has no deadline to violate, for convenience, our derivation employs X_0 and Y_0 with $P(X_0 = 1) = 1$ and $P(Y_0 = 1 \mid X_0 = 1) = 1$.

We let $W_i = j$ if the arc (i,j) appears on the tour. We note that arc (i,j) appears in the tour if customers i and j are realized, i and j's deadlines are met, and customers k, $k = i+1, \dots, j-1$, are either not realized or are skipped if they are realized.

If Z_i is the random variable representing the cost incurred traveling from customer i to the next customer that appears on the tour after i, the expected cost of a tour is:

$$E\left[\sum_{i=0}^{n} Z_i\right] = \sum_{i=0}^{n} E[Z_i].$$

We assume that $E[Z_i \mid X_i = 0] = 0$. That is, no cost is associated with customer i if i is not realized. Then, using conditioning arguments, we have:

$$\begin{aligned} E[Z_i] &= p_i E\left[Z_i \mid X_i = 1\right] = p_i E\left[E[Z_i \mid X_i = 1, Y_i]\right] \\ &= p_i \phi_i P(Y_i = 0 \mid X_i = 1) + p_i E[Z_i \mid X_i = 1, Y_i = 1] P(Y_i = 1 \mid X_i = 1), \end{aligned} \tag{5}$$

where ϕ_i in equation 5 follows from the fact that, if i is realized, but its deadline is violated, then customer i is skipped on the tour and the penalty ϕ_i is incurred.

We now consider the expectation in the second term in equation 5. Conditioning on which arc (i,j) appears next on the tour, we have:

$$\begin{aligned} E[Z_i \mid X_i = 1, Y_i = 1] &= \sum_{j=i+1}^{n} E[Z_i \mid X_i = 1, Y_i = 1, W_i = j] P(W_i = j \mid X_i = 1, Y_i = 1) \\ &\quad + E[Z_i \mid X_i = 1, Y_i = 1, W_i = 0] P(W_i = 0 \mid X_i = 1, Y_i = 1) \qquad (6) \\ &= \sum_{j=i+1}^{n} d_{i,j} P(W_i = j \mid X_i = 1, Y_i = 1) + d_{i,0} P(W_i = 0 \mid X_i = 1, Y_i = 1). \qquad (7) \end{aligned}$$

We note that the second term in equation 6 captures the cost of the return to the depot. The substitution of $E[Z_i \mid X_i = 1, Y_i = 1, W_i = j]$ with $d_{i,j}$ in equation 7 follows from the fact that, if arc (i,j) is present, then a cost of $d_{i,j}$ is incurred.

In order to evaluate the expected cost of a tour, we are then left to calculate the probabilities $P(Y_i = 0 \mid X_i = 1)$, $P(Y_i = 1 \mid X_i = 1)$, and $P(W_i = j \mid X_i = 1, Y_i = 1)$. We will again make use of the probability $G(i,t)$, which we defined in Section 3.1 to be the probability that the arrival time to customer i is at or before time t, given customer i is realized. With this definition, we have $P(Y_i = 1 \mid X_i = 1) = G(i, l_i)$ and $\bar{G}(i, l_i) = 1 - G(i, l_i) = P(Y_i = 0 \mid X_i = 1)$. As in equation 3, we can recursively compute $G(i,t)$ once we compute the probabilities $g(i,t) = P(A_i = t \mid X_i = 1)$, where $t = 0, \dots, l_i$. As before, A_i is a random variable representing the time of arrival to i. It will be useful to introduce the indicator function

$$\mathcal{I}(i,t) = \begin{cases} 0 \text{ if } l_i < t \\ 1 \text{ otherwise.} \end{cases}$$

We also note that $g(0,0) = 1$ and $g(0,t) = 0$ for $t > 0$. Then,

$$g(j,t) = \sum_{i=0}^{j-1} P(A_j = t \mid X_j = 1, X_i = 1, Y_i = 1, A_i = t - d_{i,j}) P(X_i = 1, Y_i = 1, A_i = t - d_{i,j} \mid X_j = 1) \tag{8}$$

$$
\begin{aligned}
&= \sum_{i=0}^{j-1} P(A_j = t \mid X_j = 1, X_i = 1, Y_i = 1, A_i = t - d_{i,j}) P(X_i = 1) P(A_i \le l_i, A_i = t - d_{i,j} \mid X_i = 1) && (9)\\
&= P(A_j = t \mid X_j = 1, X_0 = 1, Y_0 = 1, A_0 = t - d_{0,j}) P(A_0 = t - d_{0,j} \mid X_0 = 1)\\
&+ \sum_{i=1}^{j-1} P(A_j = t \mid X_j = 1, X_i = 1, Y_i = 1, A_i = t - d_{i,j}) p_i \mathcal{I}(i, t - d_{i,j}) P(A_i = t - d_{i,j} \mid X_i = 1) && (10)\\
&= P(A_j = t \mid X_j = 1, X_0 = 1, Y_0 = 1, A_0 = t - d_{0,j}) g(0, t - d_{0,j})\\
&+ \sum_{i=1}^{j-1} P(A_j = t \mid X_j = 1, X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_k + Y_k \le 1, k = i+1, \ldots, j-1) p_i \mathcal{I}(i, t - d_{i,j})\\
&\times g(i, t - d_{i,j}) P(X_k + Y_k \le 1, k = i+1, \ldots, j-1 \mid X_i = 1, Y_i = 1, X_j = 1, A_i = t - d_{i,j}), && (11)
\end{aligned}
$$

where equation 8 follows from conditioning, and equation 9 is based on the chain rule, the definition of Y_i, and customer i's presence being independent of customer j's presence. Equation 10 results from explicit treatment of the depot, the definition of X_i and the recognition that, for customer i's deadline to be met and arrival to customer i to be at time $t - d_{i,j}$, then $t - d_{i,j}$ must be less than or equal to l_i. Equation 11 results from conditioning on the fact that arc (i, j) appears in the tour only if customers k, $k = i+1, \ldots, j-1$, are either not realized or are skipped if they are realized. The substitution of $g(i, t - d_{i,j})$ follows from its definition. Finally, we note that, given the conditions,

$$
P(A_j = t \mid X_j = 1, X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_k + Y_k \le 1, k = i+1, \ldots, j-1) = 1. \tag{12}
$$

We now demonstrate how the probability $P(X_k + Y_k \le 1, k = i+1, \ldots, j-1 \mid X_i = 1, Y_i = 1, X_j = 1, A_i = t - d_{i,j})$ can be computed recursively. We note that this is the probability that customers k, $k = i+1, \ldots, j-1$, do not appear on the tour given that the last customer visited on the tour was customer i and arrival to i occurred at time $t - d_{i,j} \le l_i$. The key is to recognize customer k, $k = i+1, \ldots, j-1$, does not appear on the tour if either it is not realized, or if it is realized, it is skipped. A realized customer k is skipped if arrival would occur at some time $t > l_k$. Importantly, whether or not customer k is skipped is independent of whether or not customer $q > k$ appears on the tour. Hence, we can drop the condition on X_j. For notational convenience, we let $H(i, j-1, t - d_{i,j}) = P(X_k + Y_k \le 1, k = i+1, \ldots, j-1 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j})$. Then, we begin by applying the chain rule and conditioning on whether or not customer $i+1$ is realized. We get

$$
\begin{aligned}
H(i, j-1, t - d_{i,j}) &= [P(Y_{i+1} = 0 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} = 1) p_{i+1}\\
&+ P(Y_{i+1} \le 1 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} = 0)(1 - p_{i+1})]\\
&\times P(X_k + Y_k \le 1, k = i+2, \ldots, j-1 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} + Y_{i+1} \le 1)\\
&= [(1 - \mathcal{I}(i+1, t - d_{i,j} + d_{i,i+1})) p_{i+1} + (1 - p_{i+1})]\\
&\times P(X_k + Y_k \le 1, k = i+2, \ldots, j-1 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} + Y_{i+1} \le 1), && (13)
\end{aligned}
$$

where equation 13 follows from finding that, if customer $i+1$ is not realized, its deadline cannot be violated by assumption and

$$
P(Y_{i+1} = 0 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} = 1) = \begin{cases} 0 & \text{if } t - d_{i,j} + d_{i,i+1} \le l_{i+1} \\ 1 & \text{otherwise.} \end{cases}
$$

That is, we cannot violate the deadline at customer $i+1$ if, traveling from customer i at time $t - d_{i,j}$, arrival to $i+1$ is at or before time l_{i+1}. Letting $(1 - \mathcal{I}(i+1, t - d_{i,j} + d_{i,i+1})) p_{i+1} + (1 - p_{i+1}) = H(i, i+1, t - d_{i,j})$, again applying the chain rule, and this time conditioning on X_{i+2}, equation 13 becomes

$$
\begin{aligned}
&H(i, i+1, t - d_{i,j}) \, [(1 - \mathcal{I}(i+2, t - d_{i,j} + d_{ii+2})) p_{i+2} + (1 - p_{i+2})]\\
&\times P(X_k + Y_k \le 1, k = i+3, \ldots, j-1 \mid X_i = 1, Y_i = 1, A_i = t - d_{i,j}, X_{i+1} + Y_{i+1} \le 1, X_{i+2} + Y_{i+2} \le 1). && (14)
\end{aligned}
$$

Continuing to condition on k, $k = i+3, \ldots, j-1$, in the manner of equation 14, we derive a recursion in the general form $H(i, m, t) = H(i, m-1, t)[(1 - \mathcal{I}(m, t + d_{im})) p_m + (1 - p_m)]$. $H(i, m, t)$ is the probability that, given customer i was visited at time $t \le l_i$, customers k, $k = i+1, \ldots, m$, are either not realized or are skipped. With this recursion and equation 12, we can rewrite equation 11 as

$$
g(0, t - d_{0,j}) H(0, j-1, t - d_{0,j}) + \sum_{i=1}^{j-1} p_i \mathcal{I}(i, t - d_{i,j}) g(i, t - d_{i,j}) H(i, j-1, t - d_{i,j}).
$$

Next, we derive $P(W_i = j \mid X_i = 1, Y_i = 1)$ for $0 < i < j$. Recall that the realization $W_i = j$ can only occur if customers i and j are realized, their deadlines met, and all customers k, $k = i+1, \ldots, j-1$, are either not realized or are skipped if they are realized. Hence,

$$P(W_i = j \mid X_i = 1, Y_i = 1) = \sum_{t=0}^{\min\{l_i, l_j - d_{i,j}\}} P(W_i = j \mid X_i = 1, Y_i = 1, A_i = t) P(A_i = t \mid X_i = 1, Y_i = 1) \tag{15}$$

$$= \sum_{t=0}^{\min\{l_i, l_j - d_{i,j}\}} \frac{P(A_j \leq l_j, X_j = 1, X_i = 1, Y_i = 1, A_i = t, X_k + Y_k \leq 1, k = i+1, ..., j-1)}{P(X_i = 1, Y_i = 1, A_i = t)}$$
$$\times P(A_i = t \mid X_i = 1, Y_i = 1) \tag{16}$$

$$= \sum_{t=0}^{\min\{l_i, l_j - d_{i,j}\}} p_j P(A_i = t \mid X_i = 1, X_j = 1) P(X_k + Y_k \leq 1, k = i+1, \ldots, j-1 \mid A_i = t, X_i = 1, Y_i = 1, X_j = 1)$$
$$\times P(A_j \leq l_j \mid X_i = 1, Y_i = 1, A_i = t, X_j = 1, X_k + Y_k \leq 1, k = i+1, \ldots, j-1) \frac{1}{P(Y_i = 1 \mid X_i = 1)} \tag{17}$$

$$= \sum_{t=0}^{\min\{l_i, l_j - d_{i,j}\}} p_j \frac{g(i,t)}{G(i, l_i)} H(i, j-1, t), \tag{18}$$

where equation 15 follows from conditioning on the arrival time to customer i. We need only sum from $t = 0$ to $t = \min\{l_i, l_j - d_{i,j}\}$ since customer i's deadline must be met and the fact that we cannot possibly meet the deadline at customer j if arrival to customer i occurred after $l_j - d_{i,j}$. Equation 16 follows from the definition of conditional probability. Equation 17 follows from the chain rule, the fact that $P(Y_i = 1 \mid X_i = 1, A_i = t, X_j = 1) = 1$ given the available values of t, and recognizing that customer demand is independent. Finally, Equation 18 results from the definition of $G(i, l_i)$ and from replacing $P(X_k + Y_k \leq 1, k = i+1, \ldots, j-1 \mid A_i = t, X_i = 1, X_j = 1)$ with $H(i, j-1, t)$. As before, given the conditions, $P(A_j \leq l_j \mid X_i = 1, Y_i = 1, A_i = t, X_j = 1, X_k + Y_k \leq 1, k = i+1, \ldots, j-1) = 1$.

In the case of travel to and from the depot, we must modify our computation of $P(W_i = j \mid X_i = 1, Y_i = 1)$. When traveling from the depot, we get:

$$P(W_0 = j \mid X_0 = 1, Y_0 = 1) = p_j \mathcal{I}(j, d_{0,j}) H(0, j-1, 0).$$

The probability of returning to the depot from customer i is calculated as:

$$P(W_i = 0 \mid X_i = 1, Y_i = 1) = \sum_{t=0}^{l_i} \frac{g(i,t)}{G(i, l_i)} H(i, n, t). \tag{19}$$

The probability in 19 follows from the fact that traveling from i to the depot occurs if all customers following i on the tour are either skipped or not realized.

We can now write the expected cost of a given tour as

$$\begin{aligned} E\left[\sum_{i=1}^{n} Z_i\right] &= \sum_{j=1}^{n} d_{0,j} p_j \mathcal{I}(j, d_{0,j}) H(0, j-1, 0) \\ &+ \sum_{i=1}^{n} p_i \left[\phi_i \bar{G}(i, l_i) + \left[\sum_{j=i+1}^{n} d_{i,j} p_j \sum_{t=0}^{\min\{l_i, l_j - d_{i,j}\}} \frac{g(i,t)}{G(i,l_i)} H(i, j-1, t)\right.\right. \\ &+ \left.\left. d_{i,0} \sum_{t=0}^{l_i} \frac{g(i,t)}{G(i,l_i)} H(i, n, t)\right]\right]. \end{aligned} \tag{20}$$

Clearly, direct computation of the objective function value is computationally expensive. Since a customer is never served after its deadline, it is $O(n^2 \max_i\{l_i\})$ to compute all of the g, H, and $P(W_i = j \mid X_i = 1, Y_i = 1)$ values. Once these values are computed, it is $O(n \max_i\{l_i\})$ to compute all of the G values and $O(n^2)$ to compute the expected cost of the tour. Thus, the complexity of the Recourse II cost expression is dominated by the time to compute the g, H, and $P(W_i = j \mid X_i = 1, Y_i = 1)$ values and is $O(n^2 \max_i\{l_i\})$.

4.2. Special Cases for Recourse II

As we did for Recourse I, we present special cases for Recourse II. Again, these special cases are important because they offered a reduced complexity for the evaluation of expected cost. The set of special cases for Recourse II are based on assumptions regarding p_i, ϕ_i, l_i, and $d_{i,j}$. The first three special cases assume $d_{i,j} = 1$ for every i and j. As with Recourse I, unit distances make all function evaluations polynomial in n with $\max_i\{l_i\}$ being $O(n)$ and the expected cost evaluation now $O(n^3)$. The fourth special case assumes that all customers are located on a line.

4.2.1. Unit Distance, Homogeneous p, Homogeneous l, Homogeneous ϕ In this special case, we assume that $p_i = p$ and $\phi_i = \phi$ for all i, $d_{i,j} = 1$ for every i and j, and $l_i = l$ for every i. As with Recourse I, under these assumptions, each customer is identical and hence every tour τ has the same expected cost.

4.2.2. Unit Distance, Homogeneous p, Homogeneous l For this second special case, we allow the penalty value ϕ_i to be customer dependent, but maintain our assumptions $p_i = p$ for every i, $d_{i,j} = 1$ for every i and j, and $l_i = l$ for every i. The later a customer i appears in the tour, the probability of violating the deadline at i increases and hence the probability of incurring the penalty ϕ_i increases. Since the travel time between customers is irrelevant, the optimal tour is then a tour that orders the customers in decreasing order of ϕ_i.

4.2.3. Unit Distance, Homogeneous l For the third special case, we relax any assumptions on p_i and ϕ_i for every i, while maintaining our restrictions that $l_i = l$ for every i and $d_{i,j} = 1$ for every i and j. Given these assumptions, it is true that, if we are going to be late to one customer, we are going to be late to every subsequent customer on the tour. We recognize that arrival to the q^{th} realized customer occurs at time q. To be late to customer i, arrival to customer i must be at time $l+1$ or later. Thus, to be in violation of the delivery deadlines, i must be the $(l+1)^{st}$ or greater realized customer in the tour.

Let $p(k,q)$ be the probability that exactly k of the first q customers are realized. We can compute this probability using the following recursion:

$$p(k,q) = p_q p(k-1,q-1) + (1-p_q)p(k,q-1).$$

To complete our calculations, we use the following initial conditions:

$$p(0,q) = \prod_{i=1}^{q}(1-p_i)$$

and

$$p(q,q) = \prod_{i=1}^{q} p_i.$$

Then, $G(q,t) = \sum_{k=0}^{t-1} p(k,q-1)$, noting $G(q,t) = 1$ when $q \le t$ and $G(0,t) = 1$ for every t. Further, $P(W_i = j \mid X_i = 1, Y_i = 1)$ reduces to

$$p_j \frac{G(i,l-1)}{G(i,l)} \prod_{k=i+1}^{j-1}(1-p_k)$$

for $0 < i < j$ because we will only travel to j from i if customers $k = i+1,\ldots,j-1$ are not realized. Violating a deadline at any customer k, $k = i+1,\ldots,j-1$, means that we would also violate the deadline at customer j. For travel from the depot, $P(W_0 = j \mid X_0 = 1, Y_0 = 1)$ becomes

$$\sum_{j=1}^{n} p_j \mathcal{I}(j,1) \prod_{k=1}^{j-1}(1-p_k).$$

For the return to the depot, $P(W_i = 0 \mid X_i = 1, Y_i = 1)$ reduces to

$$\prod_{m=i+1}^{n}(1-p_m) + \frac{G(i,l) - \bar{G}(i,l-1)}{G(i,l)} - \frac{G(i,l) - \bar{G}(i,l-1)}{G(i,l)} \prod_{m=i+1}^{n}(1-p_m).$$

This term follows from the fact that the depot is visited after i only if either customers $i+1,\ldots,n$ are not realized or if arrival to customer i is after time $l-1$.

We can then compute the expected cost of a tour as follows:

$$d_{0j}\sum_{j=1}^{n} p_j \mathcal{I}(j,1)\prod_{k=1}^{j-1}(1-p_k) + \sum_{i=1}^{n}\Bigg[p_i\phi_i\bar{G}(i,l) + p_i G(i,l)\Bigg[\sum_{j=i+1}^{n} d_{i,j}p_j \frac{G(i,l-1)}{G(i,l)} \prod_{k=i+1}^{j-1}(1-p_k)$$
$$+d_{i,0}\left(\prod_{m=i+1}^{n}(1-p_m) + \frac{G(i,l) - \bar{G}(i,l-1)}{G(i,l)} - \frac{G(i,l) - \bar{G}(i,l-1)}{G(i,l)} \prod_{m=i+1}^{n}(1-p_m)\right)\Bigg]\Bigg].$$

With this special case of Recourse II, the expected cost evaluation becomes $O(n^2)$ rather than $O(n^3)$ as in the general model with unit distances.

4.2.4. Straight Line Distance As we did with Recourse I, for this special case, we assume that we can order the customers $1,\ldots,n$ such that $d_{i,k} = d_{i,j} + d_{j,k}$ for $i < j < k$ and $d_{i,j} = d_{ji}$ for $i \neq j$. Under these assumptions, the identity permutation is also an optimal solution for Recourse II. The result follows in an analogous fashion to the discussion presented in Section 3.2.3. Essentially, if a customer on the identity tour is skipped instead of being served late, it does not allow arrival to the succeeding customer to occur any earlier.

5. Chance Constrained Models

In a chance constrained model, certain deterministic constraints in a mathematical program are replaced by a set of chance constraints. These new constraints restrict the probability that certain events will occur. In many real world applications, finding an a priori tour that visits all customers by their deadlines may be impossible if customers are treated deterministically. A chance constrained approach for the PTSPD allows the possibility of late arrivals at realized customers, but restricts the likelihood in conjunction with minimizing expected costs. This type of constraint is analogous to the cycle-service level which is used extensively in inventory management [see (Hopp and Spearman, 2000) and (Nahmias, 2001)]. In a chance constrained model, there are no penalties associated with being late, unlike recourse models.

In this chance constrained model, each customer i has a value α_i representing the allowed probability it will not be served before its deadline l_i. Assigning each customer the same α value treats all customers identically, where choosing different α values allows a company to prioritize service among customers when resources are tight.

The chance constrained model for the PTSPD is very different than the one developed for the stochastic vehicle routing problem (SVRP). In Dror et al. (1989) and other papers focused on the SVRP, the chance constraints model the likelihood that each vehicle will reach capacity before its tour is complete. In other words, they model the chance the vendor is willing to take in planning the tour, where here the chance is based on the individual customers and the probability they will be served late. This significantly changes the number and form of the chance constraints.

In presenting the two recourse models, we discussed how to compute the expected cost for a given sequence. In a chance constrained approach, not every sequence is feasible, so our modeling approach is different. Here the sequence is determined by binary decision variables in a math program where the math program imposes the chance constraints. Thus, we will need to modify the objective function calculation appropriately to incorporate these binary variables as well as incorporate these variables in modeling the chance constraints. Note that in modeling the PTSPD, the choice of variables needs to be different than those used in modeling the PTSP (for the modeling details of the PTSP, see (Laporte et al., 1994)). It is also important to note that for some problems, there may be no feasible solution.

In this paper, we present both recourse and chance constraint models because they each have their own advantages. One advantage of a chance constrained approach is that it may be hard to estimate appropriate penalty values for recourse models, where confidence levels may be easier to determine. Secondly, chance constrained models are math programs which, depending on their formulation and the size of α, may be easier to optimize than comparable recourse models. As discussed in the earlier sections, solving recourse models typically involves local search which can be slow to converge with complex objective functions.

We will present an initial non-linear version of a chance constrained model for the PTSPD, and then introduce additional variables that allow for the transformation to a linear mixed binary integer program (IP). After presenting the objective and constraints for the general form of the PTSPD, we will discuss variations of the problem that simplify the structure of the math program.

5.1. Initial Model

5.1.1. Initial Objective Formulation Like the PTSP, the objective function is based on the expected cost of a tour. Since the ordering of the tour is determined by binary decision variables, the expected cost expression presented earlier in equation 1 must be adapted for this.

The decision variables will be

$$x_{i,q} = \begin{cases} 1 \text{ if customer } i \text{ is in position } q \\ 0 \text{ otherwise.} \end{cases}$$

Given this, the expected cost associated with the arc from position q to position r in the tour, where $q < r$, can be computed by the following:

$$\sum_{i=1}^{n} \sum_{j=1, i \neq j}^{n} x_{i,q} x_{j,r} d_{i,j} p_i p_j \prod_{s=q+1}^{r-1} \left(\sum_{k=1, k \neq i,j}^{n} x_{k,s} (1 - p_k) \right). \tag{21}$$

The decision variables ensure that the expected cost is zero unless we have the appropriate i, j pair with customer i in position q and j in position r. The expected cost associated with using an arc is again based on the distance between the endpoints ($d_{i,j}$), the probability of both endpoints occurring ($p_i p_j$), and the probability of the customers in between positions q and r not occurring (captured by the product). The summation within the product will consist of only one nonzero term which is the probability that the customer assigned to position s ($q+1 \leq s \leq r-1$) will not require a delivery.

Next, we consider the costs to travel from the depot to a customer and from a customer back to the depot. The cost from the depot to the customer in the r^{th} position is essentially a simplification of equation 21 due to the fact that the depot's existence is not stochastic:

$$\sum_{j=1}^{n} x_{j,r} d_{0,j} p_j \prod_{s=1}^{r-1} \left(\sum_{k=1, k \neq j}^{n} x_{k,s} (1 - p_k) \right). \tag{22}$$

The expected cost to travel from position q to the depot is:

$$\sum_{i=1}^{n} x_{i,q} d_{i,0} p_i \prod_{s=q+1}^{n} \left(\sum_{k=1, k \neq i}^{n} x_{k,s} (1 - p_k) \right). \tag{23}$$

The differences between equations 22 and 23 are primarily in the product. To travel directly from the depot to the r^{th} customer, the customers in the first to $(r-1)^{st}$ positions must not be realized. To travel directly from the q^{th} customer to the depot, the customers in the $(q+1)^{st}$ to the n^{th} positions must not be realized.

Given these definitions, the objective function for the chance constrained model can be expressed by:

$$\text{Minimize} \sum_{r=1}^{n}\sum_{j=1}^{n} x_{j,r} d_{0,j} p_j \prod_{s=1}^{r-1}\left(\sum_{k=1,k\neq j}^{n} x_{k,s}(1-p_k)\right) + \tag{24}$$

$$\sum_{q=1}^{n-1}\sum_{r=q+1}^{n}\sum_{i=1}^{n}\sum_{j=1,i\neq j}^{n} x_{i,q}x_{j,r}d_{i,j}p_ip_j \prod_{s=q+1}^{r-1}\left(\sum_{k=1,k\neq i,j}^{n} x_{k,s}(1-p_k)\right) +$$

$$\sum_{q=1}^{n}\sum_{i=1}^{n} x_{i,q} d_{i,0} p_i \prod_{s=q+1}^{n}\left(\sum_{k=1,k\neq i}^{n} x_{k,s}(1-p_k)\right). \tag{25}$$

5.1.2. Initial Constraint Formulation The constraint set for the PTSPD models the deadline constraints via chance constraints. The remaining constraints are deterministic and are the constraints of the well-studied traveling salesman problem (see (Gutin and Punnen, 2002) for a comprehensive review).

As discussed earlier, the chance constraint restricts the probability that each customer i will be served after l_i to be less than or equal to α_i. If we let F_{rl_i} represent the probability the customer in position r will not be served later than l_i, we can express the chance constraint with:

$$F_{r,l_i} \geq (1-\alpha_i)x_{i,r} \quad \forall i,r. \tag{26}$$

The right hand side of equation 26 enforces that the F_{r,l_i} values are pushed above $(1-\alpha_i)$ only when customer i is in the r^{th} position.

Next, we will derive an expression for F_{r,l_i}. The probability that the customer in the r^{th} position will not be served later than time l_i is the sum of three components: the probability that the r^{th} customer is not realized (since a customer cannot be served late if it is not served), the probability that the tour will travel directly from the depot to the r^{th} customer and arrive by time l_i, and the probability that the tour can travel from the previous realized customer to the r^{th} and arrive by time l_i. We will refer to these pieces by F^1_{r,l_i}, F^2_{r,l_i}, and F^3_{r,l_i}, respectively. Note that we do not have to consider the return to the depot in formulating the chance constraints. The return to the depot occurs after the tour has been completed and thus does not impact deadlines.

The derivation of F^1_{r,l_i} is straightforward and follows from its definition:

$$F^1_{r,l_i} = \sum_{k=1}^{n} x_{k,r}(1-p_k). \tag{27}$$

Next, we define F^2_{r,l_i}:

$$F^2_{r,l_i} = \sum_{k=1,d_{0,k}\leq l_i}^{n} x_{k,r}p_k \prod_{s=1}^{r-1}\left(\sum_{j=1,j\neq k}^{n} x_{j,s}(1-p_j)\right). \tag{28}$$

The product portion of equation 28 captures the probability that no customers in positions prior to the r^{th} have been realized. Note that it is not necessary to consider customers k where $d_{0,k} > l_i$ because it would be impossible to leave the depot and arrive at k by l_i.

F^3_{r,l_i} is a little more complicated to express. For compactness, we require the use of a new variable:

- t_{qrl} = probability that the last realized customer through the r^{th} stop in the tour is in position q and l units of travel time or less have been accumulated at that point.

The t variable contains the q index because it is important to know where the route is coming from on its way to r so that distances are included correctly. Assuming we can compute t_{qrl}, F^3_{r,l_i} can be expressed by:

$$F^3_{r,l_i} = \sum_{k=1}^{n} x_{k,r}p_k \sum_{q=1}^{r-1}\sum_{j=1,j\neq k,d_{j,k}\leq l_i}^{n} x_{j,q}t_{q,r-1,l_i-d_{j,k}}. \tag{29}$$

If customer k is in the r^{th} position and customer j is the prior realized customer in the q^{th} position, the tour can arrive at k at l_i or earlier if the tour leaves j at time $l_i - d_{j,k}$ or earlier. The value of $t_{q,r-1,l_i-d_{j,k}}$ in equation 29 is the likelihood that the tour can depart j at such a time. Note that it is not necessary to consider customers j where $d_{j,k} > l_i$ because it would be impossible to leave such a customer and arrive at i on time. Combining these three terms, our chance constraints become:

$$(27) + (28) + (29) \geq (1 - \alpha_i)x_{i,r} \quad \forall i, r > 1 \tag{30}$$

Note that when $r = 1$, the F^3_{r,l_i} term is not needed, and F^2_{r,l_i} simplifies to:

$$F^2_{1,l_i} = \sum_{k=1, d_{0,k} \leq l_i}^{n} x_{k,1} p_k. \tag{31}$$

Thus, the chance constraint for $r = 1$ is:

$$(27) + (31) \geq (1 - \alpha_i)x_{i,1} \quad \forall i \tag{32}$$

We can compute the t_{qrl} values recursively. These recursive equations are included as constraints to the math program. The recursion for t_{qrl} is defined by equations 33 and 34, with an initial value defined by equation 35.

$$t_{qrl} = \sum_{k=1}^{n} x_{k,r}(1 - p_k) t_{q,r-1,l} \quad \forall l, 1 < q < n, q < r \tag{33}$$

$$t_{r,r,l} = \sum_{k=1, d_{0,k} \leq l}^{n} x_{k,r} p_k \prod_{s=1}^{r-1} \sum_{j=1, j \neq k}^{n} x_{j,s}(1 - p_j) + \sum_{k=1}^{n} x_{k,r} p_k \sum_{q=1}^{r-1} \sum_{j=1, j \neq k, d_{j,k} \leq l}^{n} x_{j,q} t_{q,r-1,l-d_{j,k}} \quad \forall l, r > 1 \tag{34}$$

$$t_{1,1,l} = \sum_{k=1, d_{0,k} \leq l}^{n} x_{k,1} p_k \quad \forall l \tag{35}$$

Equation 33 reflects that, if the customer in the q^{th} position was the last realized customer through $r - 1$ positions, it will still be the last at the r^{th} position if the r^{th} customer is not realized. Equation 34 considers where the most recent realized customer is the one in the r^{th} position. The first part of this sum considers when no prior customers have been realized. The second part of equation 34 considers each customer that may have been realized just prior to the r^{th} and includes the probability that there is sufficient time to arrive at the r^{th} by time l. Equation 35 sets up the initial value of the recursion and is based on probabilities for the first customer on the tour. Note that the maximum l value that would need to be considered corresponds to the largest l_i value. The range of l values between 0 and $\max_i \{l_i\}$ are sufficient for defining the t terms. We note that the use of the l indices prevents this from being a compact formulation, meaning we are not guaranteed there will be a polynomial number (in n) of variables and constraints.

At this point, we have described the full set of added constraints necessary for a chance constrained formulation of the PTSPD. Thus, our initial nonlinear formulation can be described by:

$$\begin{array}{l} (25) \\ \text{subject to:} \\ (30) - (35) \\ x_{i,q} \in (0,1) \quad \forall i, q \\ x \in S_{TSP} \end{array}$$

where

- S_{TSP} = the set of all feasible solutions to the TSP.

5.2. Linear Model

5.2.1. Linear Transformation of Objective The objective function in equation 25 requires only n^2 binary variables but is nonlinear. With the addition of new continuous variables and constraints, we can linearize this expression. These new variables include an explicit representation of the expected cost to travel between two positions and a new term to represent the product portion of the objective:

- $e_{q,r}$ = expected cost of arc from position q to position r or arc to or from the depot (position 0)
- $g_{q,r}$ = probability of customers from positions q to r not occurring ($q \le r$).

Since the g variables are used in defining the expected cost terms, we will start by defining these. To reduce the number of terms that are multiplied at one time and preserve linearity, we will express the g variables in the form of a recursion. The basic recursion is found in equation 36, with a special case in equation 37. The non-negativity of g variables is enforced in equation 38. The g variables are defined with greater than or equal to constraints because of their adverse impact on the objective function.

$$g_{q,r} \ge g_{q,r-1}(1-p_i) + (1-p_i)(x_{i,r}-1) \quad \forall i, q < n, r > 1, q < r \tag{36}$$

$$g_{q,q} = \sum_{i=1}^{n} x_{i,q}(1-p_i) \quad \forall q \tag{37}$$

$$g_{q,r} \ge 0 \quad \forall q \le r \tag{38}$$

The value of $g_{q,r}$ is equal to the probability that none of the customers in positions q to $r-1$ will occur (g_{qr-1}) nor the customer in the r^{th} position. If the g_{qr-1} terms were multiplied by the $x_{i,r}$ variables, the recursion in equation 36 would be nonlinear. By including the $x_{i,r}$ terms in an additive way, the $g_{q,r}$ values are pushed to be nonzero only when the correct i, r combination is considered, and the resulting expression is linear. Equation 37 handles the special case when only one position is considered, so no recursion is needed.

To define the $e_{q,r}$ variables, we can use the g terms to replace the product portions of equations 21, 22, and 23. To remove the multiplication of the two binary variables at the beginning of the summation, we will again need to include the x variables in an additive manner. To appropriately define the these variables, we need nine expressions (39-47). Equations 45-47 enforce non-negativity.

$$e_{q,r} \ge d_{i,j}p_ip_jg_{q+1,r-1} + (x_{i,q}-1)d_{i,j}p_ip_j + (x_{j,r}-1)d_{i,j}p_ip_j \quad \forall i, j, i \ne j, q < n-1, q+1 < r \tag{39}$$

$$e_{q,q+1} \ge d_{i,j}p_ip_j + (x_{i,q}-1)d_{i,j}p_ip_j + (x_{j,q+1}-1)d_{i,j}p_ip_j \quad \forall i, j, i \ne j, q < n \tag{40}$$

$$e_{0,r} \ge d_{0,j}p_jg_{1,r-1} + (x_{j,r}-1)d_{0,j}p_j \quad \forall j, r > 1 \tag{41}$$

$$e_{0,1} \ge d_{0,j}p_j + (x_{j,1}-1)d_{0,j}p_j \quad \forall j \tag{42}$$

$$e_{q,0} \ge d_{i,0}p_ig_{q+1,n} + (x_{i,q}-1)d_{i,0}p_i \quad \forall i, q < n \tag{43}$$

$$e_{n,0} \ge d_{i,0}p_i + (x_{i,n}-1)d_{i,0}p_i \quad \forall i \tag{44}$$

$$e_{q,r} \ge 0 \quad \forall q < r \tag{45}$$

$$e_{0,r} \ge 0 \quad \forall r \tag{46}$$

$$e_{q,0} \ge 0 \quad \forall q \tag{47}$$

In equation 39, the $e_{q,r}$ term is forced above zero and to its appropriate value only when the appropriate i and j values are considered. Equation 40 is included to account for when consecutive positions are considered, and equations 41 - 44 account for traveling from or to the depot.

With the addition of $e_{q,r}$ and $g_{q,r}$ variables and their associated linear constraints, the objective now becomes simply

$$\text{Minimize} \sum_{r=1}^{n} e_{0,r} + \sum_{q=1}^{n-1} \sum_{r=q+1}^{n} e_{q,r} + \sum_{q=1}^{n} e_{q,0} \tag{48}$$

which is clearly linear.

5.2.2. Linear Transformation of Constraints The nonlinear constraints can also be linearized with the use of additional variables. Equations 28 and 34 contain products of $r-1$ terms. We will replace this product with:

- $k_{i,r}$ = the probability that none of the first $r-1$ customers have been realized if $x_{i,r} = 1$, 0 otherwise.

To prevent multiplication between continuous and binary variables, we also use:

- $\hat{t}_{q,r,l,i,j} = t_{q,r,l}$ if $x_{i,q} = 1$ and $x_{j,r} = 1$, 0 otherwise .

The k terms can easily be defined by equations 49 - 51.

$$k_{i,r} \leq x_{i,r} \quad \forall i, r > 1 \tag{49}$$

$$\sum_{i=1}^{n} k_{i,r} \leq \sum_{j=1}^{n} (1-p_j) k_{j,r-1} \quad \forall r > 2 \tag{50}$$

$$\sum_{i=1}^{n} k_{i,2} \leq \sum_{j=1}^{n} (1-p_j) x_{j,1} \tag{51}$$

Since the k terms appear in the left-hand side of the chance constraints, we determine their value through the use of less than or equal to constraints. Only when $x_{i,q} = 1$ will $k_{i,q}$ be able to be nonzero (equation 49). The sum in equation 50 defines the recursion when $x_{i,q} = 1$. Equation 51 establishes an initial value.

Similarly, we define the $\hat{t}_{q,r,l,i,j}$ variables by equations 52-56. The new $\hat{t}$ variables replace the t variables in the original formulation. The $\hat{t}$ variables are used in satisfying the chance constraints. Thus, we set their values using less than or equal to constraints.

$$\sum_{j=1}^{n} \hat{t}_{qrlij} \leq x_{i,q} \quad \forall l, i, q \leq r \tag{52}$$

$$\sum_{i=1}^{n} \hat{t}_{q,r,l,i,j} \leq x_{j,r} \quad \forall l, j, q \leq r \tag{53}$$

$$\sum_{i=1, i\neq j}^{n} \hat{t}_{q,r,l,i,j} \leq (1-p_j) \sum_{i=1}^{n} \sum_{k=1}^{n} \hat{t}_{q,r-1,l,i,k} \quad \forall l, j, q < r, r > 1 \tag{54}$$

$$\sum_{j=1, d_{0,j} \leq l}^{n} \hat{t}_{r,r,l,j,j} \leq \sum_{j=1}^{n} p_j k_{j,r} + \sum_{j=1}^{n} p_j \sum_{q=1}^{r-1} \sum_{i=1, i\neq j, d_{i,j} \leq l}^{n} \sum_{k=1, k\neq i,j}^{n} \hat{t}_{q,r-1,l-d_{i,j},i,k} \quad \forall l, r > 1 \tag{55}$$

$$\sum_{j=1, d_{0,j} \leq l}^{n} \hat{t}_{1,1,l,j,j} \leq \sum_{j=1}^{n} x_{j,1} p_j \quad \forall l \tag{56}$$

Equations 52 and 53 force the new variables to zero unless both $x_{i,q} = 1$ and $x_{j,r} = 1$. Equation 54 is analogous to equation 33 in that the probability associated with $\hat{t}$ values when $q < r$ is based on the probability of the r^{th} also not occurring. Equations 55-56 are the same as equations 34-35 modified for the new variables.

Last, given our new k and $\hat{t}$ variables, we can replace our nonlinear chance constraint (equation 30) with equation 57.

$$\sum_{j=1}^{n} [x_{jq}(1-p_j) + \sum_{k=1, d_{0,k} \leq l_i}^{n} p_k k_{k,q} + p_k \sum_{r=1}^{q-1} \sum_{j=1, j\neq k, d_{j,k} \leq l_i}^{n} \sum_{h=1, h\neq j,k}^{n} \hat{t}_{r,q-1,l_i-d_{j,k},j,h}] \geq (1-\alpha_i) x_{i,q} \quad \forall i, q > 1 \tag{57}$$

Likewise, equation 32 becomes:

$$\sum_{j=1}^{n} [x_{j,q}(1-p_j) + \sum_{j=1, d_{0,j} \leq l_i} x_{j,1} p_j \geq (1-\alpha_i) x_{i,q} \quad \forall i \tag{58}$$

With this, our new linear objective and constraint formulation becomes:

$$\begin{array}{l} (48) \\ \text{subject to:} \\ (36)-(47) \\ (49)-(58) \\ x_{i,q} \in (0,1) \quad \forall i, q \\ x \in S_{TSP} \end{array}$$

Observe that even with the additional variables, there are still only n^2 binary variables.

5.3. Special Cases for Chance Constrained Model

Next, we present a selection of special cases resulting from various assumptions about the instance data.

5.3.1. Unit Distance In this special case, we assume $d_{i,j} = 1$ for every i and j. The ordering of customers, in this case, will not impact the expected cost of the tour, so the problem becomes strictly one of feasibility.

Unit distances further simplify the problem since we do not have to keep track of where the tour is coming from to know if it will be late at the next customer. Lateness can be determined simply by the number of customers that have been realized, not which customers specifically have been realized. Because we do not need to track where the tour is coming from at all times, the basic t variables can lose an index and can be replaced by:

- $t_{k,q}$ = probability that exactly k of the first q customers are realized

Now we easily omit chance constraints for certain i, q combinations. If $l_i \geq q$ and i is in the q^{th} position, it is impossible for the tour to be late at i. Thus, for customers with l_i values of n or greater, no chance constraints will be needed.

The formulation with unit distances can be linearized as with heterogeneous distances. The key difference is that the $\hat{t}$ variables can function with fewer indices. Now we can use:

- $\hat{t}_{k,q,i}$ = probability that exactly k of the first $q-1$ are realized if i is in the q^{th} position, 0 otherwise.

Since the t and $\hat{t}$ values both lose their l index, both formulations are now compact and polynomial in n. The k variables are still needed for the linearization, and they function as before.

5.3.2. Unit Distance, Homogeneous p If the p values are homogeneous, the values for the t (and $\hat{t}$) variables discussed above are based on a binomial distribution and can be computed independently of the sequence of customers. These constants can be computed ahead of time, as well as the probability of the q^{th} customer being the first realized customer. Thus, the basic chance constraint becomes :

$$x_{i,q}[(1-p) + p(1-p)^{q-1} + p\sum_{k=1}^{l_i-1} t_{k,q-1}] \geq (1-\alpha_i)x_{i,q} \quad \forall i, q \geq 2, l_i \leq q-1. \tag{59}$$

Equation 59 is linear if we precompute values for p^q and $(1-p)^q$ where $q = 1$ to n. With the use of pre-computed constants, the only variables required are the binary x variables.

5.3.3. Unit Distance, Homogeneous p, Homogeneous α If l_i values are heterogeneous but both p and α values are homogeneous, we should order the customers by increasing l_i value to maximize the likelihood of feasibility. With this ordering, we can quickly verify if the chance constraint is satisfied for each customer without solving an IP. If this ordering is not feasible, no ordering is feasible.

5.3.4. Unit Distance, Homogeneous p, Homogeneous l Similarly, if α_i values are heterogeneous but both p and l values are homogeneous, the customers should be ordered by increasing α_i value and then feasibility validated for each customer.

5.3.5. Unit Distance, Homogeneous p, Homogeneous α, Homogeneous l If we further know that all α and l values are identical, then solving the problem becomes trivial. With these restrictions on the problem, we can easily verify that if the probability of being late at the last customer on a sequence is less than α, then any ordering is feasible. If it is greater than α, then no ordering is feasible. Thus, we need only to examine whether

$$(1-p) + p(1-p)^{q-1} + p\sum_{k=1}^{l-1} t_{k,n-1} \geq (1-\alpha) \tag{60}$$

where t_{kn-1} is again a constant based on a binomial distribution and is not sequence dependent. If $l \geq n$, then no chance constraints are required.

5.3.6. Homogeneous p, Homogeneous α, Homogeneous l Without the restriction on customer distances, the objective value will be sequence dependent. If p , α , and l values are homogeneous, the chance constraints can be enforced, though, via the last customer in the sequence. If the last customer satisfies the chance constraint, so will all preceding customers. This allows the number of constraints in the formulation to be greatly reduced.

5.3.7. Straight Line Distance Recall that with a straight line distance assumption, we can order the customers $1,...,n$ such that $d_{i,k} = d_{i,j} + d_{j,k}$ for $i < j < k$. From a cost perspective, it is clearly optimal for customers to be served in the identity order, $1,...,n$. Depending on the values of p,α, and l, though, this sequence may not be feasible. If distances obey the straight line assumption and the identity tour satisfies the chance constraints, the identity tour is clearly optimal in terms of the chance constrained model. It is also important to observe that the identity tour is the tour that is the most likely to be feasible, regardless of the choice of these values.

Consider any non-identity tour τ. As a non-identity tour, τ contains at least one customer who is not in topological order. Let i be the first customer not in topological order in τ. Let τ' be the tour such that i is removed from its position in τ and inserted back in the tour so that i is topological order. For any realization ω, the cost of τ' is less than the cost of τ, as shown in 3.2.3. At the same time, for any realization, we arrive to no customer later in τ' than in τ, but we likely arrive to i earlier. Thus, by moving i, we not only decrease the cost, but we also potentially increase the probability that the customers will satisfy the chance constraints. Continuing the proposed tour improvement scheme, we can eventually transform a tour τ into the identity tour, making it a tour with lower cost and more likely to be feasible for all customers.

6. Illustrative Example

In this section, we explore an example which highlights how the three PTSPD models evaluate the expected cost of two tours. This example also demonstrates that significant cost savings can be achieved by taking advantage of low probability events. First consider $N = \{1, \dots, 4\}$, where the customers are located in $\mathcal{R}^2$ with the coordinates $(4,0), (2,-2), (0,-1)$, and $(1,1)$, respectively. We assume that the depot is located at $(0,0)$ and that travel time is equivalent to distance. Assume that only customer 1 has a deadline and that deadline is $l_1 = 4$. Note that, if customer 1 is realized, it must be served directly after the depot for its deadline to be satisfied. Let Tour I be $0 \to 1 \to 2 \to 3 \to 4$, illustrated in Figure 1.A. If the customers require service with probabilities $P = \{0.1, 1, 1, 0.5\}$ and assuming that customers who are not realized are skipped, Tour I has an expected travel cost of 7.79.

Because customer 1 occurs with a relatively low probability, we also consider a second tour $0 \to 4 \to 1 \to 2 \to 3$, illustrated in Figure 1.B. We call this tour Tour II. Note that, if both customers 1 and 4 are realized, customer 1's deadline will be violated. However, the realization of both customers 1 and 4 has a probability of only .05. Thus, Tour II satisfies a 95% chance constraint. It also has an expected travel cost of 7.28 time units, a savings over the Tour I expected cost of 6.5%. When evaluated using Recourse I, a fixed-charge penalty of greater than 10.19 or a per-unit-charge of 17.68 is necessary before Tour I becomes preferable to Tour II. Further, when we evaluate Tour II using our Recourse II evaluation function, for any penalty less than or equal to 13.02, Tour II has a lower expected cost than Tour I. In the case of Recourse II, the penalty necessary for equivalency with Tour I is larger than for Recourse I because Recourse II skips any customer whose deadline is violated. Thus, Recourse II does not incur a cost, in addition to the penalty, for traveling to the customer whose deadline is violated.

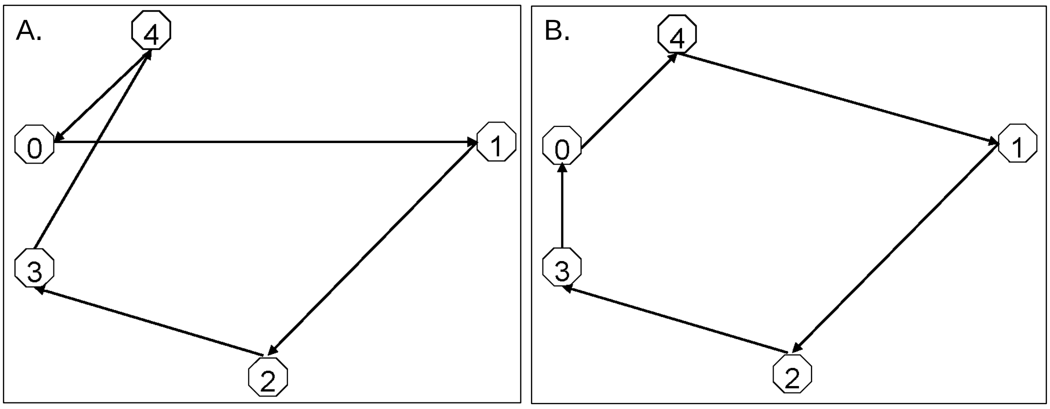

Figure 1 **Two Tours for Customers located in $\mathcal{R}^2$**

7. Computational Experiments

The above illustrative example raises the important question of when is it particularly important to model the presence of customers stochastically when deadlines are involved. Since a stochastic model is much more complex to model and solve than a TSPTW, for example, it is worth exploring what instance characteristics yield different tours when customers are modelled stochastically instead of deterministically. To help address this question, we will focus on the Recourse I model with per-unit-time penalties and look at instance characteristics such as selection of deadlines and customer realization probabilities. We will compare the solutions found from solving the PTSPD with the solutions found by solving the traveling salesman problem with soft deadlines (TSPD). In contrast to the PTSPD, the TSPD model does not account for the probability of a customer needing service and rather assumes that each customer will be visited. We use soft deadlines so that feasible solutions can be found for all of the instances considered. At the end of this section, we will also address the tradeoffs in run times between solving stochastic and deterministic models.

7.1. Data Sets and Solution Procedure

Because the PTSPD is a new problem, our computational experiments require the development of data sets. We begin with the 20-, 40-, and 60-customer, time-window width 20 units TSPTW instances first proposed by Dumas et al. (1995). From each of the 15 sets, we generate two instances which differ only in the deadlines. For the first instance, we set the deadline equal to the opening time of the time window unless that time is zero in which case the deadline is set to the closing time of the time window. For the second instance, we set all deadlines equal to the closing time of each customers' time window. These instances will hereafter be referred to as "early" and "late" deadlines, respectively. The early deadline instances are used to represent situations where feasible solutions with respect to deadlines are very unlikely to exist if all customers are realized. With the later deadlines, it is much more likely that feasible solutions will exist even when all customers are realized.

For each of the now 30 instances, we then consider four different probability settings. Two of these settings are homogeneous settings in which all probabilities are set to 0.1 and 0.9, respectively. These two settings represent when each customer is unlikely to be realized or, alternately, very likely to be realized. Probabilities of 0.1 may be appropriate when goods are delivered to the home or to small businesses, since daily deliveries are unlikely. Likewise, probabilities of 0.9 may be appropriate for medium-sized businesses who receives packages almost every day. The other two settings are heterogeneous. In the first case, the probability of each customer is generated from a uniform random number between 0 and 1. This helps us understand how the results change when there are more options in terms of customer probabilities. In the second heterogeneous case, we randomly assign probabilities of either 0.1 or 1. This case addresses the situation in which large and small businesses are served by the same vehicle. These two data sets will be referred to in the tables by the labels "range" and "mixed," respectively.

Finally, for each instance, we consider two different per-unit-time penalties. In one case, the penalty is set to 5 and in the other to 50. These choices of penalties represent small versus large costs for failure to satisfy the customer deadlines.

In the end, we tested a total of 240 instances. The data sets are available from `http://myweb.uiowa.edu/bthoa/research.html`.

The focus of our computational tests is not to develop efficient solution procedures, but rather to demonstrate the difference between TSPD and PTSPD solutions. The development of such efficient approaches requires a significant additional effort both theoretically and computationally and is the subject of a followup paper (Campbell and Thomas, 2007). Because the PTSPD is a new problem and because there is little research focusing on TSPD or even the TSPTW with soft time windows, some experimentation is required to find an effective method for finding good solutions to both problems. For ease of implementation, we focus on local-search heuristics. While Ohlmann and Thomas (to appear), Cheh et al. (1991), and Carlton and Barnes (1996) find the 1-shift neighborhood most effective for the TSPTW, our tests find that this neighborhood does not by itself consistently find good solutions for either problem. Bertsimas and Howell (1993) find that the 2-Opt neighborhood was effective over a range of probabilities for the PTSP. Yet, like the 1-shift neighborhood, the 2-Opt does not consistently find good solutions to either the PTSPD or the TSPD.

To overcome these shortcomings, we implement a greedy randomized adaptive search procedure (GRASP) with a 1-shift neighborhood and a restricted candidate list of size three (see Feo and Resende (1995) for a detailed discussion of GRASP). To avoid premature convergence, we run a best-improving 2-Opt neighborhood each time the 1-shift neighborhood fails to find an improving solution. If the 2-Opt procedure finds an improving solution, we run the GRASP procedure again starting from this new solution.

To solve the TSPD, we run the heuristic 20 different times, each time seeding the heuristic with a new random solution. In general, each of the 20 GRASP runs converged to the same solution. This fact coupled with the random nature of GRASP search suggests that our search procedure is returning good solutions for the TSPD.

Given our desire to demonstrate differences between the TSPD and PTSPD solutions, in solving the PTSPD, we seed the same heuristic with the TSPD solutions. Because of greatly increased run times resulting from the evaluation of the PTSPD objective, we run the heuristic only 10 times when solving the PTSPD. In all cases, we report the best found solution over all of the runs. For both the PTSPD and TSPD solutions, the reported solution value is the with respect to the PTSPD objective to allow comparison between TSPD and PTSPD solutions.

7.2. Results

The results of our computational experiments can be found in Tables 1-4. Table 1 describes the results when the instances with the early deadlines are used, and the penalty is set to 5. As discussed above, it is much harder to find feasible solutions to these problems, so the objectives contain a larger penalty portion than when later deadlines are used (Table 3). When probabilities are set to 0.1, we find much bigger differences, on average, between the PTSPD and TSPD solutions than when probabilities are set to 0.9. When probabilities are set to 0.9, most customers are realized, and not surprisingly, the PTSPD solutions are close to the deterministic solutions. When probabilities are 0.1, the solutions begin to take advantage of low probability events in the tour construction. The improvements relative to the TSPD solutions tend to be higher for the larger data sets because the advantages of low probability events are compounded over a larger number of customers. The

Table 1 Experiments with Earlier Deadlines and Penalty=5

Prob		0.1			0.9			Range			Mixed		
Data set		PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff
n = 20	1	71.1	72.9	2%	581.8	581.8	0%	302.1	320.2	6%	291.7	300.1	3%
	2	66.9	68.0	2%	411.8	413.1	0%	164.2	171.4	4%	165.8	183.3	10%
	3	82.3	83.3	1%	513.3	513.3	0%	354.7	367.8	4%	452.0	452.0	0%
	4	65.7	67.3	2%	585.2	585.2	0%	256.7	306.3	16%	270.1	332.4	19%
	5	69.2	69.9	1%	406.1	406.1	0%	234.1	253.2	8%	229.7	289.7	21%
n = 40	1	114.9	125.9	9%	336.3	337.8	0%	281.5	292.4	4%	233.3	258.7	10%
	2	93.5	96.7	3%	361.6	361.6	0%	265.5	270.3	2%	205.1	212.6	4%
	3	100.2	100.9	1%	528.1	554.5	5%	274.6	285.9	4%	368.0	391.5	6%
	4	99.0	108.8	9%	627.1	657.7	5%	240.5	294.5	18%	221.2	271.3	18%
	5	109.5	120.3	9%	393.4	393.4	0%	250.1	257.6	3%	258.4	290.0	11%
n = 60	1	124.5	129.8	4%	588.6	594.9	1%	324.2	354.1	8%	251.4	275.3	9%
	2	163.4	174.8	7%	902.8	911.8	1%	397.2	424.9	7%	552.6	594.6	7%
	3	166.2	178.8	7%	665.5	668.0	0%	457.2	501.5	9%	400.6	464.8	14%
	4	135.5	156.3	13%	625.4	625.4	0%	323.6	382.9	15%	432.9	436.9	1%
	5	136.8	145.9	6%	567.1	580.6	2%	337.9	347.3	3%	327.4	358.9	9%

two heterogeneous settings lead to even larger differences between PTSPD and TSPD solutions, with the mixed setting yielding the most dramatic results. For example, when there are only 20 customers, considering customers probabilistically can lead to a savings in expected cost of up to 21% (see data set 5 when $n = 20$). This result follows from the fact that the tour can be ordered such that customers with higher presence probabilities have low probabilities of having their deadlines violated in the resulting solutions.

Table 2 illustrates the effect of larger per-unit-time penalty values, and it is easy to see that the objective functions increase for all of the solutions. For instances where probabilities equal 0.9, for example, there is an increase by an order of magnitude. At same time, the difference between PTSPD and TSPD solutions also increase. For instance, with probabilities of 0.1 and $n = 40$, the fourth data set goes from having a 9% difference to a 34% difference. Differences such as these result because the TSPD objective is now dominated by the penalty cost which leads to tours that are designed to minimize penalty cost at the expense of travel cost. However, the PTSPD takes advantage of the fact that many events occur with low probability, and hence the PTSPD is able to reduce penalty costs in conjunction with routing costs.

When later deadlines are considered in combination with smaller penalties in Table 3, we find less difference in general between PTSPD and TSPD solutions than when earlier deadlines are used. This result follows from the fact that feasible solutions exist even when all customers are realized and accordingly there is the possibility of incurring no penalty cost. Note, though, that some solutions still contain a small penalty cost because the tradeoff in potential savings in distance is worth the cost of a deadline violation.

When larger penalties are combined with the later deadlines in Table 4, we do not see the significant change due to penalties as when earlier deadlines are used. The PTSPD and TSPD objective values are consistently equal to or just slightly higher than when lower penalties are used. This outcome is because few, if any, penalties are incurred in the realization of the solutions. When probabilities of 0.9 are considered, the differences between PTSPD and TSPD solutions are consistently zero with both sizes of penalties. This outcome suggests that modeling the stochastic presence of customers is less important when feasible solutions are possible and probabilities of customers being realized are high.

While PTSPD solutions can offer significant advantages over TSPD solutions, they do so with a significant computational cost. Table 5 presents run times in CPU seconds for the PTSPD and TSPD heuristics on datasets with later deadlines and a penalty of 5, labeled with the respective headings PTSPD and TSPD. The run time represents the total run time for 10 PTSPD iterations

Table 2 Experiments with Earlier Deadlines and Penalty=50

Prob		0.1			0.9			Range			Mixed		
Data set		PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff
n = 20	1	199.3	206.0	3%	3385.9	3385.9	0%	1402.3	1598.6	12%	1397.8	1559.8	10%
	2	92.5	93.1	1%	2155.0	2155.0	0%	412.4	432.8	5%	538.9	550.7	2%
	3	128.0	128.4	0%	2589.4	2589.4	0%	1379.9	1379.9	0%	1932.5	1932.5	0%
	4	105.5	115.2	8%	3405.3	3405.3	0%	903.0	1124.3	20%	980.3	1363.0	28%
	5	137.3	141.6	3%	2265.8	2287.2	1%	863.8	931.3	7%	727.3	759.5	4%
n = 40	1	117.1	129.5	10%	600.5	600.5	0%	489.9	500.8	2%	249.3	255.2	2%
	2	95.2	102.4	7%	1271.4	1271.7	0%	468.2	780.4	40%	234.4	249.0	6%
	3	119.9	125.6	5%	2632.9	2712.7	3%	639.4	676.0	5%	1172.4	1840.5	36%
	4	110.6	166.4	34%	3922.2	3922.2	0%	468.8	691.2	32%	260.9	450.4	42%
	5	113.9	125.1	9%	1324.2	1359.6	3%	288.9	311.9	7%	641.0	789.8	19%
n = 60	1	171.2	192.9	11%	3004.6	3008.8	0%	946.7	1098.2	14%	360.8	603.3	40%
	2	285.2	296.2	4%	5499.1	5636.2	2%	1485.7	1516.7	2%	2849.7	2953.4	4%
	3	308.4	336.8	8%	3535.6	3582.8	1%	2065.9	2368.1	13%	1151.1	2105.9	45%
	4	189.9	216.6	12%	3011.6	3009.4	0%	1007.8	1166.4	14%	1295.4	1299.5	0%
	5	217.8	229.4	5%	2671.4	2774.4	4%	828.1	869.5	5%	926.2	1075.8	14%

Table 3 Experiments with Later Deadlines and Penalty=5

Prob		0.1			0.9			Range			Mixed		
Data set		PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff
n = 20	1	56.2	56.5	1%	210.7	210.7	0%	132.3	133.0	1%	123.9	124.1	0%
	2	63.6	64.4	1%	191.2	191.2	0%	122.7	133.2	8%	119.0	139.6	15%
	3	76.4	77.9	2%	243.4	243.4	0%	217.3	225.8	4%	240.0	245.9	2%
	4	61.2	61.6	1%	242.3	242.3	0%	181.4	182.6	1%	183.6	187.1	2%
	5	61.7	61.8	0%	192.2	192.2	0%	160.3	161.1	1%	173.3	173.5	0%
n = 40	1	114.4	118.1	3%	303.3	303.3	0%	230.4	243.7	5%	200.1	228.7	12%
	2	92.4	95.6	3%	238.5	238.7	0%	202.1	203.6	1%	191.6	191.7	0%
	3	97.8	98.0	0%	277.7	278.0	0%	213.9	231.3	8%	208.9	210.3	1%
	4	96.8	99.6	3%	258.6	258.6	0%	213.2	213.9	0%	216.4	216.5	0%
	5	107.0	117.6	9%	278.2	278.2	0%	225.7	242.2	7%	181.9	195.0	7%
n = 60	1	118.8	120.2	1%	313.3	313.4	0%	247.3	257.3	4%	230.7	233.4	1%
	2	139.0	153.6	10%	345.4	345.7	0%	274.5	283.2	3%	292.4	305.8	4%
	3	145.2	156.0	7%	328.6	328.6	0%	276.4	276.4	0%	263.9	268.0	2%
	4	128.8	147.9	13%	349.9	349.9	0%	235.0	290.0	19%	334.9	335.0	0%
	5	125.3	134.9	7%	323.5	323.8	0%	254.1	280.3	9%	258.6	274.6	6%

Table 4 Experiments with Later Deadlines and Penalty=50

Prob		0.1			0.9			Range			Mixed		
Data set		PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff	PTSPD	TSPD	% Diff
n = 20	1	56.3	56.5	0%	234.6	234.6	0%	132.3	133.0	1%	124.0	124.2	0%
	2	63.7	65.4	3%	207.5	207.5	0%	134.4	142.2	5%	121.6	139.6	13%
	3	76.8	77.9	1%	243.4	243.5	0%	225.8	225.8	0%	245.9	245.9	0%
	4	61.2	61.6	1%	242.3	242.3	0%	181.4	182.6	1%	186.0	187.1	1%
	5	61.7	62.4	1%	211.5	211.5	0%	162.2	167.0	3%	174.1	177.4	2%
n = 40	1	114.4	118.1	3%	303.5	303.6	0%	230.5	243.7	5%	204.8	245.2	16%
	2	92.4	95.6	3%	238.5	238.7	0%	202.1	203.6	1%	191.6	191.7	0%
	3	97.9	98.0	0%	283.3	283.8	0%	231.3	235.1	2%	208.9	213.5	2%
	4	97.7	99.6	2%	263.0	263.0	0%	213.8	215.3	1%	216.4	216.9	0%
	5	108.4	117.6	8%	278.2	278.2	0%	226.0	242.4	7%	181.9	195.1	7%
n = 60	1	118.8	120.2	1%	310.6	310.8	0%	250.5	257.3	3%	230.7	235.8	2%
	2	149.0	153.6	3%	345.4	345.7	0%	274.9	283.2	3%	295.9	305.8	3%
	3	150.3	156.0	4%	330.2	330.5	0%	276.4	281.4	2%	264.0	274.0	4%
	4	129.1	147.9	13%	349.9	349.9	0%	235.2	290.0	19%	334.9	335.0	0%
	5	127.4	134.9	6%	323.6	323.6	0%	258.2	280.3	8%	258.7	274.8	6%

Table 5 Run Times for Experiments with Later Deadlines and Penalty=5

Prob		0.1			0.9			Range			Mixed		
Data set		PTSPD	TSPD	PTSP	PTSPD	TSPD	PTSP	PTSPD	TSPD	PTSP	PTSPD	TSPD	PTSP
n = 20	1	12	0	1	6	0	1	9	0	0	9	0	0
	2	26	0	0	8	0	0	23	0	0	29	0	0
	3	36	0	0	9	0	0	11	0	0	14	0	1
	4	70	0	0	10	0	0	14	0	0	27	0	0
	5	119	1	0	6	1	0	9	1	0	9	1	0
n = 40	1	991	1	3	335	1	2	610	1	3	1027	1	3
	2	1065	1	4	292	1	1	477	1	3	270	1	2
	3	809	1	2	306	1	2	528	1	3	497	1	3
	4	970	1	4	247	1	2	282	1	2	288	1	3
	5	712	2	2	238	2	2	508	2	3	509	2	2
n = 60	1	6929	7	29	2252	7	6	4769	7	12	4867	7	11
	2	20403	6	31	2699	6	16	7460	6	18	8448	6	17
	3	15064	6	29	2115	6	9	2090	6	11	4098	6	17
	4	11923	5	21	2065	5	18	5812	5	11	2306	5	25
	5	11340	6	20	2278	6	12	6086	6	10	5416	6	14

of the heuristic and 20 iterations of the heuristic for the TSPD. As the table shows, the run time differences are large and increase nonlinearly with problem size. Such a disparity follows even though the PTSPD heuristic is run only for half as many iterations as that for the TSPD.

To understand how much computational effort is required to evaluate the penalty portion of the PTSPD, Table 5 also presents run times for the PTSP (labeled accordingly PTSP). The PTSP is the PTSPD in which the deadlines have been removed. The PTSP solutions were found with the same algorithm as for the PTSPD and seeded with the TSPD solutions. This seeding accords no advantage to the PTSP whose solution, like those of the traveling salesman problem, differs greatly from their time-constrained counterparts.

As the table shows, the run times for the PTSP are certainly larger than those for the TSPD, but always significantly smaller than those for the PTSPD. This difference implies that the evaluation of the penalty for the PTSPD requires the bulk of the computation time. This conclusion motivates the need for further PTSPD research.

8. Conclusions and Future Work

The PTSPD is an interesting and challenging problem. In this paper, we have presented three different models to represent three different ways in which deadlines can be enforced. The Recourse I model requires that we visit each realized customer, but pay a penalty for any violation of a customer's deadline. The Recourse II model differs from the Recourse I model in that we skip any realized customer whose deadline would be violated, but a penalty is incurred for the skipped customer. The chance constrained model requires all realized customers to be visited, but limits the probability that a customer's deadline is violated. We have also identified special polynomially solvable cases for all three models as well as special cases which reduce constraints and variables for the chance constrained model.

Our illustrative example and computational experiments on the Recourse I model offer the following insights:

- When there are no feasible solutions with respect to deadlines if all customers are realized, modeling customers stochastically greatly impacts the solutions. The impact is even greater when customers have a low probability of being present. This result indicates that such an approach is critical in planning deliveries to homes and/or small businesses.
- Likewise, if feasible solutions exist and all customers have a high likelihood of being realized, solutions are minimally impacted when customers are modeled stochastically. This outcome is not

surprising, but indicates a stochastic approach is not necessary when only large businesses are being considered and deadlines are achievable.

• When customers exist with a combination of high and low probabilities, solutions are significantly impacted by modeling customers stochastically. The impact follows from the stochastic model's ability to prioritize higher probability customers in the solution. This result indicates an approach such as ours may be particularly helpful if individual vehicles are serving combinations of small and large businesses.

• The solutions to the PTSPD come with the cost of increased computation time. This increased run time is dominated by the time required to compute the cost of the penalties.

Because time constraints have not been studied in conjunction with a priori routing, there are many remaining questions for further research. The next steps in our future research will be the exploration of efficient solution procedures. Given that exact solutions to the PTSP are limited to small problem sizes, solution approaches to the recourse models should focus on heuristic methods such as local search or metaheuristics. In doing local search, we can take advantage of established time-saving techniques such as storage and recursion (Bianchi et al., 2005, Bianchi and Campbell, 2007) or approximation (Campbell, 2006, Tang and Miller-Hooks, 2004) to evaluate search moves faster. How to apply such time-saving techniques to the penalty terms of Recourse I is the subject of our follow-up computational study on the PTSPD (Campbell and Thomas, 2007).

By transforming the initial chance constrained model into a linear integer program, we can potentially take advantage of existing software to solve integer programs such as CPLEX. Due to the number and form of constraints, however, simple branch and bound approaches are unlikely to be efficient for large problem instances. Future research will explore the benefits of alternative formulations, as well as how to use relaxations of this IP in solving large problem instances. We can round the solution to the relaxed version of the problem to create an initial solution. This could be a used as a primal heuristic to create a good upper bound in jumpstarting a branch and bound approach. Alternately, this initial solution could be improved through local search techniques. Unlike recourse models, local search techniques here will have to consider whether or not neighboring moves are feasible and either not allow certain moves or include penalties in evaluating their costs.

There are also numerous variants of the PTSPD which do not presently appear in the literature. For example, we have considered deadlines rather than a time window with both early and late arrival restrictions. This small difference in the problem appears to have significant impact on our models. The inclusion of both pickups and deliveries, as well as tour length or vehicle capacity restrictions, are also important real world constraints and are important to consider for future research.

Acknowledgment

The authors would like to thank three anonymous referees for their useful suggestions. This work was partially supported by the National Science Foundation through grant number 0237726(Campbell).

References

E. Baker, "An Exact Algorithm for the Time Constrained Traveling Salesman Problem," *Operations Research* 31, 938–945 (1983).

J. J. Bartholdi, L. K. Platzman, R. L. Collins, and W. H. Warden, "A Minimal Technology Routing System for Meals on Wheels," *Interfaces* 13, 1–8 (1983).

C. Bastian and A. H. G. Rinnooy Kan, "The Stochastic Vehicle Routing Problem Revisited," *European Journal of Operational Research* 56, 407–412 (1992).

D. J. Bertsimas, *Probabilistic Combinatorial Optimizations Problems*, Ph.D. thesis, Massachusetts Institute of Technology (1988).

D. J. Bertsimas, “A Vehicle Routing Problem with Stochastic Demand,” *Operations Research* 40, 574–585 (1992).

D. J. Bertsimas, P. Chervi, and M. Peterson, “Computational Approaches to Stochastic Vehicle Routing Problems,” *Transportation Science* 29, 342–352 (1995).

D. J. Bertsimas and L. H. Howell, “Further Results on the Probabilistic Traveling Salesman Problem,” *European Journal of Operational Research* 65, 68–95 (1993).

D. J. Bertsimas, P. Jaillet, and A. R. Odoni, “A Priori Optimization,” *Operations Research* 38, 1019–1033 (1990).

D. J. Bertsimas and D. Simchi-Levi, “A New Generation of Vehicle Routing Research: Robust Algorithms, Addressing Uncertainty,” *Operations Research* 44, 286–303 (1996).

L. Bianchi and A. M. Campbell, “Extension of the 2-p-opt and 1-shift algorithms to the Heterogeneous Probabilistic Traveling Salesman Problem,” *European Journal of Operational Research* 176, 131–144 (2007).

L. Bianchi, J. Knowles, and N. Bowler, “Local Search for the Probabilistic Traveling Salesman Problem: Correction to the 2-p-opt and 1-shift algorithms,” *European Journal of Operational Research* 162, 206–219 (2005).

J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*, Springer-Verlag, New York (1997).

J. Bramel, E. G. Coffman, P. W. Shor, and D. Simchi-Levi, “Probabilistic Analysis of the Capacitated Vehicle Routing Problem with Unsplit Demands,” *Operations Research* 340, 1095–1106 (1992).

A. Campbell, “Aggregation for the Probabilistic Traveling Salesman Problem,” *Computers & Operations Research* 33, 2703–2724 (2006).

A. M. Campbell and B. W. Thomas, “Solution Approaches for the Probabilistic Traveling Salesman Problem with Deadlines,” (2007), submitted for publication.

W. B. Carlton and J. W. Barnes, “Solving the Traveling-Salesman Problem with Time Windows Using Tabu Search,” *IIE Transactions* 28, 617–629 (1996).

A. Charnes and W. W. Cooper, “Chance-Constrained Programming,” *Management Science* 6, 73–79 (1959).

A. Charnes and W. W. Cooper, “Deterministic Equivalents for Optimizing and Satisficing Under Chance Constraints,” *Operations Research* 11, 18–39 (1963).

K. Charnsirisakskul, P. M. Griffin, and P. Keskinocak, “Order Selection and Scheduling with Leadtime Flexibility,” *IEE Transactions* 36, 697–707 (2004).

K. Cheh, J. Goldberg, and R. Askin, “A note on the effect of neighborhood structure in simulated annealing,” *Computers & Operations Research* 18, 537–547 (1991).

P. Chervi, *A Computational Approach to Probabilistic Vehicle Routing Problems*, Master’s thesis, Massachusetts Institute of Technology (1988).

N. Christofides, A. Mingozzi, and P. Toth, “State Space Relaxation Procedures for the Computation of Bounds to Routing Problems,” *Networks* 11, 145–164 (1981).

M. Dror, “Modeling Vehicle Routing with Uncertain Demands as Stochastic Programs: Properties of the Corresponding Solution,” *European Journal of Operational Research* 64, 432–441 (1993).

M. Dror, G. Laporte, and P. Trudeau, “Vehicle Routing with Stochastic Demands: Properties and Solution Frameworks,” *Transportation Science* 23, 166–176 (1989).

M. Dror and P. Trudeau, “Stochastic Vehicle Routing with Modified Savings Algorithm,” *European Journal of Operational Research* 23, 228–235 (1986).

Y. Dumas, J. Desrosiers, E. Gelinas, and M. M. Solomon, “An Optimal Algorithm for the Traveling Salesman Problem with Time Windows,” *Operations Research* 43, 367–371 (1995).

FedEx, “Rules/Accessorial Tariff via All Motor Routes Naming Rules, Regulations and Claims Procedures Applying on Surface Expedited Services Between Points in North America (Except

Mexico)," http://customcritical.fedex.com/us/serviceinfo/documents/pdf/tarifffdcc101g.pdf?link=4 (2003), accessed on April 11, 2005.

FedEx, "Service Info: Money Back Guarantee," http://www.fedex.com/us/services/express/ (2004), accessed on August 9, 2004.

T. A. Feo and M. G. C. Resende, "Greedy Randomized Adpative Search Procedures," *Journal of Global Optimization* 6, 109–134 (1995).

F. Focacci, A. Lodi, and M. Milano, "A Hybrid Exact Algorithm for the TSPTW," *INFORMS Journal on Computing* 14, 403–417 (2002).

T. A. Foster, "Expedited Explodes," *Logistics Management and Distribution Report* 38, 69–73 (1999).

M. Gendreau, A. Hertz, and G. Laporte, "New Insertion and Postoptimization Procedures for the Traveling Salesman Problem," *Operations Research* 40, 1086–1094 (1992).

M. Gendreau, A. Hertz, G. Laporte, and M. Stan, "A Generalized Insertion Heuristic for the Traveling Salesman Problem with Time Windows," *Operations Research* 46, 330–335 (1998).

M. Gendreau, G. Laporte, and R. Séguin, "An exact algorithm for the vehicle routing problem with stochastic demands and customers," *Transportation Science* 29, 143–155 (1995a).

M. Gendreau, G. Laporte, and R. Séguin, "Stochastic Vehicle Routing," *European Journal of Operational Research* 88, 3–12 (1996).

M. Gendreau, G. Laporte, and M. M. Solomon, "Single-Vehicle Routing and Scheduling to Minimize the Number of Delays," *Transportation Science* 29, 56–62 (1995b).

G. Gutin and A. P. Punnen, *The Traveling Salesman Problem and Its Variations*, volume 12 of *Combinatorial Optimization*, Kluwer Academic Publishers, Dordrecht, The Netherlands (2002).

W. J. Hopp and M. L. Spearman, *Factory Physics: Foundations of Manufacturing Management*, Irwin/McGraw-Hill, Boston, second edition (2000).

P. Jaillet, *Probabilistic Traveling Salesman Problems*, Ph.D. thesis, Massachusetts Institute of Technology (1985).

P. Jaillet, "A Priori Solution of the Traveling Salesman Problem in which a Random Subset of Customers are Visited," *Operations Research* 36, 929–936 (1988).

A. Langevin, M. Desrochers, J. Desrosiers, S. Gélinas, and F. Soumis, "A Two-Commodity Flow Formulation for the Traveling Salesman and Makespan Problems with Time Windows," *Networks* 23, 631–640 (1993).

G. Laporte, F. V. Louveaux, and H. Mercure, "Models and Exact Solutions for a Class of Stochastic Location-Routing Problems," *European Journal of Operational Research* 39, 71–78 (1989).

G. Laporte, F. V. Louveaux, and H. Mercure, "A Priori Optimization of the Probabilistic Traveling Salesman Problem," *Operations Research* 42, 543–549 (1994).

S. Nahmias, *Production and Operations Analysis*, Irwin/McGraw-Hill, Boston, fourth edition (2001).

J. W. Ohlmann and B. W. Thomas, "A Compressed Annealing Approach to the Traveling Salesman Problem with Time Windows," *INFORMS Journal on Computing* (to appear).

G. Pesant, M. Gendreau, J.-Y. Potvin, and J.-M. Rousseau, "An Exact Constraint Logic Programming Algorithm for the Traveling Salesman Problem with Time Windows," *Transportation Science* 32, 12–29 (1998).

G. Pesant, M. Gendreau, J.-Y. Potvin, and J.-M. Rousseau, "On the Flexibility of Constraint Programming Models: From Single to Multiple Time Windows for the Traveling Salesman Problem," *European Journal of Operational Research* 117, 253–263 (1999).

W. B. Powell, P. Jaillet, and A. Odoni, "Stochastic and Dynamic Networks and Routing," in *Network Routing*, M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser (eds), volume 8 of *Handbooks in Operations Research and Management Science*, 141–295, North-Holland, Amsterdam, 1995.

M. W. P. Savelsbergh, "Local Search in Routing Problems with Time Windows," *Annals of Operations Research* 4, 285–305 (1985).

M. W. P. Savelsbergh and M. Goetschalckx, "A Comparison of the Efficiency of Fixed Versus Variable Vehicle Routes," *Journal of Business Logistics* 46, 474–490 (1995).

T. R. Scherck, "A View of the Future for the U.S. Expedited Transportation Industry," `http://www.colography.com` (2003), accessed on December 26, 2003.

J. D. Schulz, "Next Day, Unionized," *Traffic World* 267, 26–27 (2003).

J. Shanahan, "The Need for Speed," *Logistics Management* 42, 49–52 (2003).

S. A. Slotnick and M. J. Sobel, "Manufacturing Lead-Time Rules: Customer Retention Versus Tardiness Costs," *European Journal of Operational Research* 163, 825–856 (2005).

W. R. Stewart and B. L. Golden, "Stochastic Vehicle Routing: A Comprehensive Approach," *European Journal of Operational Research* 14, 371–385 (1983).

H. Tang and E. Miller-Hooks, "Approximate Procedures for the Probabilistic Traveling Salesman Problem," *Transportation Research Record* 1882, 27–36 (2004).

S. Y. Teng, H. L. Ong, and H. C. Huang, "An Integer L-Shaped Algorithm for the Time-Constrained Traveling Salesman Problem with Stochastic Travel Times and Service Times," *Asia-Pacific Journal of Operational Research* 21, 241–257 (2004).

F. Tillman, "The Multiple Terminal Delivery Problem with Probabilistic Demands," *Transportation Science* 3, 192–204 (1969).

United Parcel Service, "Calculating Time and Cost FAQ," `http://www.ups.com/content/us/en/resources/service/` (2004), accessed on August 9, 2004.

U.S. Department of Transportation Federal Highway Administration, "Freight Transportation: Improvements and the Economy," `http://www.ops.fhwa.dot.gov/freight/documents/improve_econ.pdf` (2004), available online since July 12, 2004.

R. Wolfler Calvo, "A New Heuristic for the Traveling Salesman Problem with Time Windows," *Transportation Science* 34, 113–124 (2000).

J. C. F. Wong, J. M. Y. Leung, and C. H. Cheng, "On a Vehicle Routing Problem with Time Windows and Stochastic Travel Times: Models, Algorithms, and Heuristics," Technical Report SEEM2003-03, Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, 2003.

W.-H. Yang, K. Mather, and R. H. Ballou, "Stochastic Vehicle Routing Problem with Restocking," *Transportation Science* 34, 99–112 (2000).