Transportation Research Part B 102 (2017) 22–37

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

An optimal stopping approach to managing travel-time uncertainty for time-sensitive customer pickup

Neža Vodopivec*, Elise Miller-Hooks

Sid and Reva Dewberry Department of Civil, Environmental, and Infrastructure Engineering, George Mason University, 4400 University Drive, MS 6C1, Fairfax, VA 22030, USA

ARTICLE INFO

Article history:
Received 11 September 2016
Revised 28 April 2017
Accepted 29 April 2017
Available online 16 May 2017

Keywords:
On-line routing
Optimal stopping
Stochastic travel times
Dial-a-ride

ABSTRACT

In dynamic vehicle routing, it is common to respond to real-time information with immediate updates to routes and fleet management. However, even if routes are updated continuously, in practice, some decisions once made are difficult to reverse. At times, it may thus be valuable to wait for additional information before acting on a decision. We use the theory of optimal stopping to determine the optimal timing of a recourse action when vehicles are likely to miss customer deadlines due to travel-time stochasticities and backup services are available. The factors involved in making this decision – that is, the likelihood that the primary vehicle will arrive late, the location of the backup vehicle, and value of waiting for additional travel-time information – each change dynamically over time. We develop a recourse model that accounts for this complexity. We formulate the optimal recourse policy as a stochastic dynamic program. Properties of the optimal policy are derived analytically, and its solution is approximated with a binomial lattice method used in the pricing of American options. Finally, we develop a two-stage stochastic optimization approach to show how the opportunity to take recourse dynamically might be integrated into a priori scheduling and routing. The framework is demonstrated for a stochastic dial-a-ride application in which taxis serve as backup to ridesharing vehicles.

1. Introduction

Advances in information technology have made up-to-the-minute travel-time forecasts and real-time vehicle location updates widely available through mobile devices. Such technology has allowed vehicle routing decisions to be made dynamically as new information is learned. However, even if routes are updated continuously, in practice, some decisions once made are difficult to reverse. After a vehicle's destination has been chosen, a customer has been assigned, or the order of the customers has been determined, it is frequently impractical or even infeasible to overturn the decision. In light of this, the best way to exploit dynamically revealed information might be at times simply to wait. Delaying a decision could mean waiting for multiple dynamic requests to be received before rerouting vehicles or waiting to see if the next customer will be a no-show before making a new assignment.

In this work, we use the theory of optimal stopping to determine the optimal timing of an irreversible recourse action when vehicles are likely to miss customer deadlines due to travel-time stochasticities and backup services are available. The factors involved in making such a recourse decision – that is, the likelihood that the primary vehicle will arrive late, the

* Corresponding author.
E-mail addresses: nvodopiv@gmu.edu (N. Vodopivec), miller@gmu.edu (E. Miller-Hooks).

http://dx.doi.org/10.1016/j.trb.2017.04.017

location of the backup vehicle, and value of waiting for additional travel-time information – each change dynamically over time. We develop a recourse model that accounts for this complexity. The model studied can help companies that transport customers, collect goods at customer locations, or those providing customers with at-home services improve adherence to time-window requirements. Example applications include: dial-a-ride services; scheduling of nurses, physical therapists, and other at-home service providers who travel to clients; and carriers that provide time-critical pickup services, such as FedEx and UPS. To provide a clear exposition, we give a tangible example of our proposed decision-making approach. Without loss of generality, concepts are developed in relation to a dial-a-ride application.

Dial-a-ride services are demand-responsive door-to-door ridesharing services. Their typical users are mobility-impaired passengers seeking rides to doctors' appointments and travelers accessing airports. The Dial-A-Ride Problem (DARP) considers dial-a-ride routes and schedules that are efficient and meet quality-of-service requirements, such as a maximum time aboard and adherence to pickup time windows. Regardless of how robust DARP solutions are, in practice, travel-time stochasticities often cause schedule perturbations. As a result, customers are frequently picked up late (Vodopivec et al., 2015).

As a recourse measure, we consider dynamically reassigning customers to taxis when ridesharing vehicles fall behind schedule. In urban settings, wheelchair-enabled taxis are becoming increasingly available through government subsidies (Weiser and Flegenheimer, 2013). These taxis could readily back up shared mobility service providers. We focus not on route development or revision, but on an effective strategy for calling taxis as backup to a vehicle deployed along a given route. We assume that ridesharing operators must negotiate between two costs: the added fare for calling a taxi and a fine for picking up a customer late. Moreover, if a taxi is called but does not arrive on time, the operator must pay both the added fare and the fine. As time passes, the vehicle approaches its destination and the operator knows with greater certainty whether calling a taxi is warranted. However, the more the decision is postponed, the greater the chance that a taxi, if called, will itself arrive late. We refer to the problem of determining if and when to call a taxi in this context as the Taxi Recourse Problem (TRP).

The decision to call a taxi is fundamentally different from the decision to continue with the vehicle—the first is irreversible and the second always contains a hidden opportunity for later recourse. This suggests our analysis will benefit from a decision-making model that accounts for the value of the information that might be gained by postponing a decision. With this in mind, we model the TRP as an optimal stopping problem. An optimal stopping problem can be understood as a game in which a gambler observes a succession of random variables and is offered a reward based on each outcome. Whenever a reward is offered, the gambler must decide either to take the reward and quit the game or to turn down the reward and continue to play. Optimal stopping theory seeks a stopping rule that will maximize the expected reward.

In the TRP, the decision to stop and call a taxi can be made at any point on a continuous, finite interval of time. Inside the interval, the stopping cost is known *a priori* and increases linearly with time as the probability of taxi lateness increases. If stopping has not occurred by the end of the interval, the decision process terminates automatically and the terminal stopping cost is a binary value determined by whether the vehicle has reached its destination.

Optimal stopping methods have been widely applied to problems involving irreversible decisions made in stochastic environments. In typical examples, a decision maker considers when to take an action that may be postponed, but once performed is difficult to reverse. Optimal stopping became well-known when its application to the pricing of financial options was discovered (e.g. Karatzas, 1988). It has also been used to determine when to make an investment (Dixit and Pindyck, 1994), when to harvest a forest (Alvarez, 2004), when to perform surgery on hemodialysis patients (Chakraborty et al., 2010), and when to stop testing software before its release (Morali and Soyer, 2003). It has been applied to the field of transportation, for example, to determine expirations for promotional airfares (Feng and Gallego, 1995), develop strategies for investing in transportation projects (Chow and Regan, 2011), avoid aircraft collisions (Zapotezny-Anderson and Ford, 2011), and decide when to sell a car (Cirillo et al., 2015). However, to the best of our knowledge, it has not been used in the context of vehicle fleet management.

In this paper, we explore the optimal stopping policy for taking a recourse action in the context of the TRP. In Section 2, we review literature on on-line routing decisions given travel-time stochasticity; we focus on models in which travel-time uncertainties are managed by way of recourse. A formal description of the TRP is given in Section 3. Section 4 considers a series of stopping policies, each characterized by a unique set of stopping states at which a recourse action is to be taken. In particular, we discuss the optimal stopping policy and formulate its stopping set as the solution to a stochastic dynamic program. In Section 5, properties of the optimal stopping policy are derived analytically. A numerical binomial lattice technique is used to obtain the optimal stopping set in Section 6. In Section 7, we give an illustrative example to demonstrate the benefits of an optimal stopping solution. We compare this solution to results obtained through myopic decision making. Finally, the recourse model is incorporated within a two-stage stochastic programming framework that integrates strategic use of taxi as recourse with optimal use of slack in planning routes in the first stage (Section 8).

2. Literature review

We extend our literature review from Section 1 to place our problem within the broader context of on-line routing with stochastic travel times. We begin with works that consider the specific case of the DARP.

Although numerous works have addressed the DARP, most models assume that exact travel times are known in advance. Two DARP works (Xiang et al., 2008 and Schilde et al., 2014) consider real-time algorithms that dynamically recover from