Valuing travel time variability: Characteristics of the travel time distribution on an urban road

Fosgerau, Mogens and Fukuda, Daisuke

Technical University of Denmark, Tokyo Institute of Technology

2012

Valuing travel time variability: Characteristics of the travel time distribution on an urban road

Mogens Fosgerau[a,b], Daisuke Fukuda[*,c]

[a] *Technical University of Denmark, 2800 Kgs. Lyngby, Denmark*
[b] *Centre for Transport Studies, Royal Institute of Technology, Sweden*
[c] *Tokyo Institute of Technology, O-okayama, Meguro-ku, Tokyo, Japan*

Abstract

Fosgerau and Karlström [The value of reliability. Transportation Research Part B, Vol. 43 (8–9), pp. 813–820, 2010] presented a derivation of the value of travel time variability (VTTV) with a number of desirable properties. This definition of the VTTV depends on certain properties of the distribution of random travel times that require empirical verification. This paper therefore provides a detailed empirical investigation of the distribution of travel times on an urban road. Applying a range of nonparametric statistical techniques to data giving minute-by-minute travel times for a congested urban road over a period of five months, we show that the standardized travel time is roughly independent of the time of day as required by the theory. Except for the extreme right tail, a stable distribution seems to fit the data well. The travel time distributions on consecutive links seem to share a common stability parameter such that the travel time distribution for a sequence of links is also a stable distribution. The parameters of the travel time distribution for a sequence of links can then be derived analytically from the link level distributions.

Key words: value of travel time variability, travel time distribution, nonparametrics, stable distributions

1. Introduction

Travel time variability (TTV) is increasingly recognized as an important issue in the economic appraisal of transport infrastructure investment as well as transport policies such as road pricing. The importance of reducing TTV on urban and interurban roads is considered a major objective

[*]Corresponding author. Address: Department of Civil and Environmental Engineering, Tokyo Institute of Technology, 2-12-1, O-okayama, Meguro-ku, Tokyo, 152-8552 Japan. Tel.: +81 3 5734 2577; fax: +81 3 5734 3578.

Email addresses: `mf@transport.dtu.dk` (Mogens Fosgerau), `fukuda@plan.cv.titech.ac.jp` (Daisuke Fukuda)

Preprint submitted to Transportation Research Part C *August 7, 2010*

of transport policy. The traveler's marginal value of TTV, often called the value of travel time variability (VTTV), should therefore play a significant role in project evaluation. This paper contributes to this aim by investigating the empirical validity of assumptions underlying a recent theoretical derivation of the VTTV based on scheduling costs.

There are two broad modeling approaches to the travelers' valuation of TTV. The first is commonly referred to as the mean–variance approach. This approach incorporates the effects of TTV into utility or cost functions of travelers simply by taking the standard deviation or some other measure of the scale of travel time variability as an argument, jointly with mean travel time. Because of its simplicity, the mean–variance approach has been widely used (Small et al. 2005; Brownstone and Small 2005; Lam and Small 2001, among others). The mean–variance approach has however been criticized on various grounds. A main criticism is that it does not take that shape of the travel time distribution into account. Another important criticism is that the standard deviation of travel time is not an outcome of a trip. Economic theory generally defines utility directly over outcomes.

The main alternative is the scheduling approach, originally proposed by Small (1982) and extended to random travel times by Noland and Small (1995), Noland (1997) and Noland et al. (1998). The scheduling approach defines travel cost directly over outcomes, which is an advantage relative to the mean–variance approach. The scheduling approach assumes that the travelers' cost function depends in a certain way on travel time and on the arrival time relative to a preferred arrival time. Given knowledge of departure time, the distribution of travel times and the preferred arrival time, it is possible to evaluate a measure of expected travel cost that includes scheduling considerations. However, direct application of the scheduling cost function requires knowledge of the departure time and the preferred arrival time, which may be unavailable.

The assumption that travelers choose departure time optimally may replace the information on departure time and preferred arrival time. The resulting measure of expected travel cost was derived for a few special travel time distributions by Bates et al. (2001) and Noland and Polak (2002) when the travel time distribution does not depend on the departure time. It turns out that the scheduling model becomes equivalent to the mean–variance approach in these cases. These results depend, however, on specific and unrealistic assumptions concerning the distribution of random travel time.

Recently, Fosgerau and Karlström (2010) generalized these earlier results to the case where the distribution of travel times is arbitrary. Fosgerau and Karlström (2010) proved that the minimized expected cost of commuters is linear in the mean travel time and a scale measure of the travel time distribution, irrespective of the shape of the travel time distribution, provided that the travel time distribution does not depend on the departure time. Under the assumptions of their model (henceforth the FK model), the VTTV is given in terms of travelers' marginal cost of schedule delay and the average time late under the optimal departure time. The average time late is determined by the travelers' preferences and the distribution of travel times. The FK measure of VTTV may remain a good approximation when the mean and the scale of the travel time distribution depends on the time of day. Starting with observations of travel time, subtracting the mean and dividing by the scale of the travel time distribution at each time of day leaves the standardized travel time distribution. FK extended their result as an approximation when the standardized travel time distribution does not depend on the departure time.

This background motivates the present paper, which aims to carry out a check of the empirical validity of the FK assumptions regarding the distribution of travel times. It should be noted that Fosgerau and Engelson (forthcoming) have developed an alternative approach to modeling the VTTV. This approach is based on another specification of scheduling preferences, derived from Vickrey (1973). The Fosgerau–Engelson measure of VTTV is not sensitive to the shape of the travel time distribution, but like FK it does require that the travel time distribution is independent of the time of day. Furthermore, the choice between the FK model and the Fosgerau–Engelson model should be based on which formulation of scheduling preferences is thought to be the best description of the scheduling preferences of travelers. Hence the investigation of this paper remains relevant in the light of the Fosgerau–Engelson result.

The first empirical question investigated in this paper is the validity of the FK assumption that the standardized travel time can be considered to be independent of the travelers' departure time. Independence of the standardized travel time of the time of day is also a great simplification since it becomes unnecessary to account for different travel time distributions at different times of day. In this case, all the variation in the travel time distribution over the day is captured by the mean and the scale of the travel time distribution. If independence does not hold then neither FK nor the Fosgerau – Engelson result is applicable.

The next empirical question regards the distribution of standardized travel times. It is useful to be able to assume that the travel time distribution belongs to a known parametric family. Fosgerau and Karlström (2010) found in their empirical work on a single road link that the empirical distribution of the standardized travel times is asymmetric and fat right-tailed, and far from normal. Furthermore, knowledge of the travel time distribution may facilitate the aggregation of the VTTV from the link level to a sequence of links. A detailed investigation of the distributional properties of standardized travel times has not been carried out. Such an investigation is a further contribution of this paper.

We investigate these empirical questions using a large data set comprising observations of travel times on an urban road. We use minute-by-minute observations of average travel times on four consecutive links of a major radial road in Copenhagen, collected over a period of five months.

The distribution of travel times on the urban road is analyzed using a range of nonparametric techniques, including mean regression, quantile regression and kernel based estimation of conditional distributions. Nonparametric mean regression and quantile regression are employed for computing standardized travel times. The conditional distribution of standardized travel time is estimated to check whether it is independent of time of day.

We anticipate that stable distributions (see Zolotarev (1986) and Nolan (in press) for example) describe the distribution of travel times well. The family of stable distributions includes the normal as a special case. In general, this family allows distributions with skewness and heavy tails, as observed in empirical travel time distributions. Stable distributions have two important features. First, they arise as limits in the generalized central limit theorem. Second, the sum of independent stable random variables with a common stability parameter is again stable with the same stability parameter. As explained below, these two features are very attractive in relation to the FK model. In the paper we fit a stable distribution to standardized travel times and estimate the parameters that characterize the stable distribution. The goodness-of-fit for the estimated stable distribution is assessed in various ways and we examine whether the estimated stable distributions for different road links share a common stability parameter.

The paper proceeds as follows. Section 2 provides a brief description of the FK model. Section 3 explains the methodology used to investigate the statistical properties of travel time distributions. Section 4 presents our data. The empirical analysis is presented in Section 5, while Section 6 discusses the empirical results. Finally, Section 7 concludes.

2. Overview of the scheduling model

This section describes the Fosgerau and Karlström (2010) result concerning travelers' departure time choice under travel time uncertainty and the corresponding measure of VTTV. Consider a traveler about to undertake a certain trip. Without loss of generality, his preferred arrival time at the destination is taken to be zero. The traveler's scheduling cost is defined in terms of random travel time T and head start D. The head start is the duration from the departure time to the preferred arrival time and so the traveler departs at time $-D$.

The traveler is assumed consider a cost function, which depends on travel time, the head start and the lateness of arrival. A monetary travel cost is omitted for simplicity. The cost function is

$$C(D,T) = \eta D + \lambda(T-D)^{+} + \omega T,$$

where η, λ and ω are parameters, all expected to be positive, and $(T-D)^{+} = \max(T-D,\ 0)$ is the amount of time the traveler arrives late. The first term is the cost associated with departing earlier. The second term is the cost of being late and the third time is the cost of travel time per se. [1] The traveler is assumed to choose head start D to minimize the expected cost.

Express the travel time T in the convenient form $T = \mu(t) + \sigma(t)X$, where $\mu(t)$ and $\sigma(t)$ are smooth functions of the departure time t, describing the location and scale of the travel time distribution at this time. We take the location variable μ as the mean travel time. We use the interquartile range as the scale variable σ since this does not require the variance of travel time to exist. We will be considering stable distributions, which generally do not have variance. Define X as standardized travel time with probability density function ϕ and corresponding cumulative distribution function Φ. The standardized travel time distribution ϕ is assumed to be independent of D.

Fosgerau and Karlström (2010) first analyzed the case of constant μ and σ, and then extended to the case where they are variable. In the simple case, the expected cost becomes linear in μ and

[1] The present formulation is equivalent to the often used α, β, γ formulation, see Fosgerau and Karlström (2010), but is arguably more intuitive in that it has a cost of departing early rather than a cost of arriving early at the destination.

σ when travelers choose departure time to minimize expected cost. Thus, the scheduling model is equivalent to the mean–variance model. In the more general case where both μ and σ depend linearly on D, the expected cost is more complicated. Even so, the result of the first simple case can still be used as an approximation of the second case. This is briefly described in the next two subsections.

2.1. Constant mean and scale of travel times

First, we consider the case where μ and σ are constant. The traveler selects D to minimize expected cost.

$$EC^* = \min_D EC(D,T) = \min_D \left[\eta D + \lambda \int_{\frac{D-\mu}{\sigma}}^{\infty} (\mu + \sigma x - D)\phi(x)\mathrm{d}x + \omega\mu \right]. \tag{1}$$

Because the expected cost function is globally concave, the optimization problem (1) has a unique minimum and the optimal head start is given by

$$D = \mu + \sigma\Phi^{-1}\left(1 - \frac{\eta}{\lambda}\right). \tag{2}$$

Thus the optimal head start is linear in the location μ and the scale σ of the travel time distribution. The minimal expected cost is found by substituting (2) into (1) as

$$EC^* = (\eta + \omega)m + \lambda\sigma \int_{1-\frac{\eta}{\lambda}}^{1} \Phi^{-1}(\nu)\mathrm{d}\nu.$$

Now, define the functional H as:

$$H\left(\Phi, \frac{\eta}{\lambda}\right) = \int_{1-\frac{\eta}{\lambda}}^{1} \Phi^{-1}(\nu)\mathrm{d}\nu. \tag{3}$$

Note that σH is the mean lateness, such that H is the mean lateness in standardized travel time. We can rewrite the minimal expected cost as

$$EC^* = (\eta + \omega)\mu + \lambda H\left(\Phi, \frac{\eta}{\lambda}\right)\sigma. \tag{4}$$

The minimal expected cost is also linear in μ and σ for a given $H(\cdot)$. The H can be computed for a given standardized travel time distribution Φ and a traveler's scheduling preference η/λ.

The first term in (4) represents the cost of the mean travel time and the coefficient $(\eta + \omega)$ is the value of travel time. The second term represents the cost caused by the TTV and the VTTV is $\lambda H\left(\Phi, \frac{\eta}{\lambda}\right)$. The VTTV depends on the scheduling preference parameters (η and λ) and on the standardized distribution of travel time Φ. The expected cost is linear in the mean and scale of travel time for *any* fixed standardized travel time distribution Φ. This is a highly desirable property for empirical application of the FK model as it makes it very easy to compute the expected cost of trips subject to travel time risk.

2.2. Time-varying mean and scale of travel times

The assumption that the mean and the scale of the travel time distribution are constant over the time of day is not true in general. There is often pronounced systematic variation in travel times over the day caused by systematic variation in traffic demand. This means that both μ and σ will depend on the time of day. This does not exclude the possibility that the standardized travel time distribution is independent of the time of day. Fosgerau and Karlström (2010) extended the constant mean and scale model to the case where the mean travel time μ and the scale σ vary linearly with the time of day D. The distribution of the standardized travel time is still required to be independent of the time of day. In this case they found that the value of travel time is exactly the same as in the simple case but the expression for the VTTV is more complicated. They also showed that the VTTV for the case of a linearly varying mean and scale of travel time distribution can be approximated well using the VTTV for the case of constant mean and scale. They demonstrated in their empirical example, using the same data set as in the present paper, that the approximation error of the VTTV is relatively small. This result implies that it is still possible to use the result based on the constant mean and the scale of travel time to measure approximately the VTTV for time-varying mean and scale of travel times.

2.3. Remarks on the use of the theoretical model in empirical applications

The FK model is useful to define and compute the VTTV because it applies for any standardized travel time distribution. It is, however, important to note that the FK model requires that the standardized travel time distribution is constant over the time of day. [2] With this assumption, the VTTV for the time-varying mean and scale of travel times can be approximated. Hence, it is important to check empirically whether this independence assumption holds for actual travel time distributions.

[2] It is not ruled out that it is possible to establish a similar result that relaxes this condition but it has not been done.

Trips generally cover a sequence of links whereas travel time data are often recorded at the link level. So another issue in the application of the FK model in practice arises from the need for aggregating the VTTV from link to route level. This can be achieved in a simple way, if additional distributional assumptions on standardized travel times are satisfied. First, independence of travel times across links is very convenient. Second, we conjecture that standardized travel times can be described by stable distributions as explained in the next section. If this distributional assumption is plausible and if one parameter for a stable distribution is common across different road links, then addition of TTV across links becomes simple. We examine these issues empirically in the following sections.

3. Analytical framework

In this section, we explain the use of some nonparametric techniques to check whether the standardized travel time is independent of the time of day. We also examine the goodness of fit of the computed standardized travel time to the stable distributions. Express random travel time as a function of the time of day by

$$T_t = \mu(t) + \sigma(t) X_t, \tag{5}$$

where $E(X_t) = 0$ and $\sigma(t)$ is the interquartile range of travel time at time t. This is always possible. More precisely, the two functions are defined as follows:

$$\begin{aligned} \mu(t) &= \mathrm{E}[T|t] \quad \text{and} \\ \sigma(t) &= F_T^{-1}(0.75|t) - F_T^{-1}(0.25|t), \end{aligned}$$

where F_T^{-1} denotes the inverse of the distribution function of travel times conditional on the time of day.

In the following subsections, we outline nonparametric techniques to estimate $\mu(t)$ and $\sigma(t)$, which are associated with standardized travel time X_t. We estimate nonparametrically the location function $\mu(t)$ using conditional mean regression, and the scale function $\sigma(t)$ using conditional quantile regression. Nonparametric regression models, including mean, variance and quantile regressions, employ minimal constraints on the functional form of the relationship between relevant variables. Introductions to nonparametric econometrics and statistics are provided by, e.g., Härdle (1990), Pagan and Ullah (1999) and Li and Racine (2007).

3.1. Nonparametric conditional mean regression

To compute the standardized travel time conditional on the time of day, we first have to estimate the conditional mean travel time as a measure of the location of the travel time distribution. Let (T_i, t_i) be a bivariate random sample of n observations $(i = 1, \ldots, n)$. Suppose that observations are distributed over time of day with density $p(t)$. Assume that the sample realizations are i.i.d. The i.i.d. assumption for the sample realizations means that we disregard serial dependence among travel times for consecutive times of day. This is justified by noting that travelers are assumed to consider the travel time distributions over all time periods. Our analysis aims not at travel time prediction, but at estimating the travel time distribution conditional on a given time of day.

We begin by considering the regression model:

$$T_i = \mu(t_i) + \epsilon_i \qquad i = 1, \ldots, n$$

where $\mu(\cdot)$ is a smooth function of unknown form, and ϵ_i is an i.i.d. error term. We estimate $\mu(\cdot)$ nonparametrically using local constant kernel estimation. The function $\mu(t)$ is estimated by forming a weighted average of T_i around t as

$$\hat{\mu}(t) = \frac{\sum_{i=1}^{n} T_i K\left(\frac{t_i - t}{h_t}\right)}{\hat{p}(t)}, \tag{6}$$

where h_t is the bandwidth corresponding to the time of day, $K(\cdot)$ is a kernel, and $\hat{p}(t) = n^{-1}\sum_{i=1}^{n} K\left(\frac{t_i - t}{h_t}\right)$ is the kernel density estimator of $p(t)$. The bandwidth h_t determines the size of the neighborhood over which an average is taken. The selection of h_t is explained later. We use a standard normal kernel throughout the paper.

$$K(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right), \quad -\infty < u < +\infty.$$

The asymptotic normality of the estimated $\hat{\mu}(t)$ is generally guaranteed (Li and Racine 2007, p. 63) and we can compute the confidence intervals of the mean regression using the following relationship :

$$(nh_t)^{1/2}\left[\hat{\mu}(t) - \mu(t)\right] \sim \mathcal{N}\left(0,\ \sigma^2(t)\hat{p}^{-1}(t)\int_{-\infty}^{\infty} K^2(u)\mathrm{d}u\right),$$

where $\sigma^2(t)$ is the variance of travel times conditional on a given time of day t. [3] This is estimated

[3]See Pagan and Ullah (1999) for the derivation. The empirical travel time distribution has variance since travel times are bounded. Later, we shall use approximate the travel time distribution by a stable distribution for which the variance does not exist.

by performing a nonparametric mean regression of squared residuals $(T_i - \hat{\mu}(t_i))^2$ against time of day using the bandwidth from the mean regression. Note that $\int_{-\infty}^{\infty} K^2(u)\mathrm{d}u = \frac{1}{2\pi^2}$ for the standard normal kernel.

3.2. Estimating the scale of the travel time distribution

It is common to use the standard deviation (the square root of the variance) as a measure of the scale when standardizing stochastic variables. However, stable distributions, which we will consider, do not have a second moment in general. Thus, the standard deviation (or variance) may not exist. Therefore we use the interquartile range (denoted as IQR) as measure of the scale. This leads us to compute quantiles of the travel time distribution conditional on the time of day.

We first present the estimation of a conditional cumulative distribution function ("conditional distribution" hereafter) because the quantile function is obtained by inverting the conditional distribution. [4]

3.2.1. Nonparametric conditional distribution

The nonparametric kernel estimator of a conditional distribution is analogous to the local constant estimator of the conditional mean regression outlined in Section 3.1. We denote a conditional distribution function of T given t as $F(T|t)$. It is estimated without imposing any restrictive functional forms. The estimated conditional distribution is given by

$$\hat{F}(T|t) = \frac{n^{-1} \sum_{i=1}^{n} L\left(\frac{T-T_i}{h_T}\right) K\left(\frac{t_i - t}{h_t}\right)}{\hat{p}(t)}, \tag{7}$$

where $L(\cdot)$ is a kernel distribution function defined as $L(v) = \int_{-\infty}^{v} K(u)\mathrm{d}u$ and h_T denotes the smoothing bandwidth associated with travel times. The estimated conditional distribution is increasing by construction. We use the standard normal distribution for the kernel function $L(\cdot)$.

3.2.2. Nonparametric quantile regression

Once a conditional distribution function is estimated, it is straightforward to derive a conditional quantile function. The conditional ρ-quantile, $q_\rho(\cdot)$ with $\rho \in (0, 1)$is defined using the inverse

[4] We estimate the conditional distribution of travel time against the time of day. Another use for the nonparametric conditional distribution is to check the independence of the standardized travel time over the time of day. This is described later.

of the conditional distribution

$$q_\rho(t) = \inf\{T : F(T|t) \geq \rho\} = F^{-1}(\rho|t). \quad (8)$$

The estimate $\hat{q}_\rho(t)$ of $q_\rho(t)$ is computed using

$$\hat{q}_\rho(t) = \arg\min_q |\rho - \hat{F}(q|t)|,$$

where $\hat{F}(q|t)$ is taken from (7).

Finally, the interquartile range of the travel time T conditional on the time of day t is estimated, using the estimated quantile functions, by $I\hat{Q}R(t) = \hat{q}_{0.75}(t) - \hat{q}_{0.25}(t)$. We use this expression to estimate the scale function $\sigma(t)$.

3.3. Conditional distribution of the standardized travel time

Once the location and the scale functions in (5) are estimated, standardized travel times are computed simply by $X_i = (T_i - \hat{\mu}(t_i))/\hat{\sigma}(t_i)$ for each observation. For the purpose of checking the independence of the standardized travel times over the time of day, we have to examine the overall shape of the standardized travel time distribution conditional on time of day. As in Section 3.2.1 for the case of the conditional travel time distribution, the conditional standardized travel time distribution $G(x|t)$ is estimated by

$$\hat{G}(x|t) = \frac{n^{-1}\sum_{i=1}^{n} L\left(\frac{x-x_i}{h_X}\right) K\left(\frac{t_i-t}{h_t}\right)}{\hat{p}(t)}, \quad (9)$$

where h_X is the bandwidth associated with standardized travel times.

Given values of h_X and h_t, it is easy to compute the conditional distribution with (9). Furthermore, it is possible to inspect the overall shape of the conditional probability density or the conditional distribution by drawing graphs such as contours or iso-quantiles of the probabilities. Recall that the FK model requires that the standardized travel time distribution is independent of the time of day. In this case, the contours of the distribution would be completely horizontal. We use this fact as an informal check of the independence. [5]

[5] It is also possible to use cross-validation for the conditional distribution/density to detect whether the time of day is relevant to the standardized travel times, though the computation of cross-validation is generally very time consuming for large data sets. See Hall et al. (2004) and Li and Racine (2007) for details. Ichimura and Fukuda (2010) have developed a faster method for computing least-squares cross-validations for nonparametric conditional kernel density functions.

3.4. Bandwidth selection

While nonparametric kernel estimation is relatively insensitive to the choice of kernel, the choice of bandwidths does have significant effect on results. The time of day is binned by minute in our data, which means that observations do not become dense on the time axis as the number of observations increases. This violates the assumption of cross-validation methods. We therefore determine the bandwidths for the mean and interquartile range regressions using the plug-in method (Pagan and Ullah, 1999; Li and Racine, 2007). This method seeks relatively larger bandwidths than cross-validation methods for our large data set and this smoothes out some less credible fluctuations of the estimated travel time curves. [6]

The plug-in bandwidths with respect to the time of day in nonparametric mean regressions are given by

$$h_t^{plug,m} = 1.06\sigma_t n^{-1/5}, \tag{10}$$

where σ_t, the standard deviation of travel times in the population, is replaced by the sample standard deviation.

The plug-in bandwidths in the nonparametric conditional distribution, which are used for estimating the interquartile curves, are computed as

$$\begin{aligned} h_t^{plug,cd} &= 1.06\sigma_t n^{-1/6} \quad \text{and} \\ h_T^{plug,cd} &= 1.06\sigma_T n^{-1/6}, \end{aligned} \tag{11}$$

where σ_T is the standard deviation of travel times, and also estimated by the sample standard deviation.

[6] We did attempt to use the bandwidths selected by cross-validation. (See Li and Racine (2007) for cross-validation of bandwidths: their chapter 2 for the mean regression and chapter 6 for the conditional distribution and quantile regression.) However, the bandwidths for the time of day (h_t) in the mean regressions from the least squares cross-validation turned out to be less than three minutes for our all data sets. Furthermore, we found that the bandwidths of travel times in the quantile regressions (h_T), which were computed using log-likelihood cross validation for the conditional distribution, were around 0.1 minutes, which is less than the bin-width of 1 minute. These very small bandwidths lead to unlikely patterns of the estimated mean and interquartile range of travel times. For example, we observed many large bumps in the mean or interquartile range of travel time, which might be caused by a small number of incidents that occurred during the observation period. Hence, we find that the cross-validation method tends to select unreasonably small bandwidths for our large data sets. For this reason we do not use cross-validation to compute an exact test of independence.

3.5. *Fitting stable distributions to standardized travel times*

Consider now the case where we accept the independence of the standardized travel times of the time of day. We next investigate whether a stable distribution fits the data. This section presents some basic properties of stable distributions.

Stable distributions allow asymmetry (skewness) of the probability density and heavy fat tails that would be caused by rare events with extreme values. The class of stable distributions encompasses the Gaussian normal, Lévy and Cauchy distributions as special cases (Zolotarev, 1986; Nolan, in press). A univariate random variable X with a stable distribution is described by four parameters as $X \sim S(\alpha, \beta, \gamma, \delta)$. The parameters are a stability parameter $\alpha \in (0, 2]$, a skewness parameter $\beta \in [-1, 1]$, a scale parameter $\gamma > 0$ and a location parameter $\delta \in \mathbb{R}$. The stability parameter α governs the tail behavior of the distribution; the tail becomes heavier as α decreases. The parameter β describes the degree of skewness. In the case of $\beta = -1$, the distribution is maximally skewed to the left and vice versa for the case of $\beta = 1$. The distribution is symmetric when $\beta = 0$. The parameter γ determines the scale of the distribution, but it is not equivalent to the standard deviation. The location parameter δ is not generally the mean.

Stability property. A favorable characteristic of stable distributions for our analysis is the stability property. This property implies that the sum of independent stable random variables also follows a stable distribution if (and only if) they share a common stability parameter α. The convoluted distribution shares the same stability parameter and expressions exist to compute the other parameters.

Let $T^j \sim S(\alpha, \beta^j, \gamma^j, \delta^j),\ j = 1, ..., J$ be J mutually independent random variables that follow stable distributions with common stability parameter α. In our analysis, these random variables would correspond to the travel times for a set of consecutive road links. The average of the independent stable random variables $\bar{T} = (1/J) \sum_{j=1}^{J} T^j$ also follows a stable distribution (Nolan, in press). The distribution of the average of these random variables is

$$\bar{T} \sim S(\alpha, \bar{\beta}, \bar{\gamma}, \bar{\delta}),$$

where

$$\bar{\beta} = \frac{\sum_{j=1}^{J} \beta^j \left|\gamma^j/J\right|^{\alpha}}{\sum_{j=J}^{m} \left|\gamma^j/J\right|^{\alpha}},$$
$$\bar{\gamma} = \left(\sum_{j=1}^{J} \left|\gamma^j/J\right|^{\alpha}\right)^{1/\alpha}, \tag{12}$$
$$\bar{\delta} = \begin{cases} \sum_{j=1}^{J} \delta^j/J + \left(\tan\frac{\pi\alpha}{2}\right)\left[\bar{\beta}\bar{\gamma} - \sum_{j=1}^{J} \beta^j\gamma^j/J\right] & (\alpha \neq 1) \\ \sum_{j=1}^{J} \delta^j/J + \frac{2}{\pi}\left[\bar{\beta}\bar{\gamma}\log\bar{\gamma} - \sum_{j=1}^{J} \beta^j\gamma^j/J \log\left|\gamma^j/J\right|\right] & (\alpha = 1) \end{cases}.$$

It is useful for our purposes to note that linear combinations of stable random variables with the same stability parameter α is also stable with the same α. In particular, if $\sigma \neq 0$ and $X \sim S(\alpha, \beta, \gamma, \delta)$, then $\sigma X \sim S(\alpha, \text{sign}(\sigma)\beta, |\sigma|\gamma, \sigma\delta)$. We check the equivalence of the stability parameters among different road links in the empirical analysis. If their estimates are not significantly different, we could convolute standardized travel time distributions for a set of road links. For example, if two travel times are distributed as $\mu_1 + \sigma_1 X_1$ and $\mu_2 + \sigma_2 X_2$, where X_1 and X_2 are stable with the same α, then the distribution of the sum is readily computed.

Generalized central limit theorem. Another important property of stable distributions is the role they play in the generalized central limit theorem (GCLT). The classical central limit theorem states that the normalized sum of independent random variables with finite variances weakly converges to a standard normal distribution as the number of variables increases. Gnedenko and Kolmogorov (1954) generalized this idea to the case where random variables have infinite variances. Roughly speaking, the GCLT implies that the only possible limiting distribution of the normalized sum of any independent random variables is stable (Zolotarev, 1986; Nolan, in press).

Now, it is not difficult to imagine an urban road network with a large number of links where the associated standardized travel times might have heavy right tails because of a very few, but serious incidents. The distributions of standardized travel times might be obviously different from normal because they seem to be skewed to the right and fat tailed. The GCLT assures that as the sums of standardized travel times for these links accumulate over a long-range period, they might converge to a stable distribution. This would enable estimation of the standardized travel time distributions corresponding to some routes and further improve the measurement of the VTTV at the route level.

There exist closed-form expressions of stable distributions only for some special cases with some specific parameterizations (e.g., Gaussian normal [$\alpha = 2$], Cauchy [$\alpha = 1$ and $\beta = 0$] and Lévy [$\alpha = 0.5$ and $\beta = 1$]). In general, there are no explicit forms for stable densities or distributions. On the other hand, it is possible to express explicitly the characteristic function $\phi(\tau) = \mathrm{E}(\exp(i\tau X))$ for any stable distribution.

Zolotarev's (M) parameterization (Zolotarev, 1986) is preferable for numerical purposes because the characteristic functions, densities and distribution function are jointly continuous in all four parameters (Nolan, in press). With this parameterization, the characteristic function is expressed as:

$$\phi(\tau) = \begin{cases} \exp\left\{-\gamma^{\alpha}\left|\tau\right|^{\alpha}\left[1 + i\beta\left(\mathrm{sign}\tau\right)\left(\tan\frac{\pi\alpha}{2}\right)\left(\left(\gamma\left|\tau\right|\right)^{1-\alpha} - 1\right)\right] + i\delta\tau\right\} & (\alpha \neq 1) \\ \exp\left\{-\gamma\left|\tau\right|\left[1 + i\beta\left(\mathrm{sign}\tau\right)\tan\frac{2}{\pi}\left(\ln\left|\tau\right| + \ln\gamma\right)\right] + i\delta\tau\right\} & (\alpha = 1) \end{cases} . \tag{13}$$

The function $\phi(\tau)$ characterizes the stable distribution of X. Based on (13), Nolan (1997) gave a computational formula for spline approximation to stable densities and also developed program code to compute numerically the density function of a general one-dimensional stable distribution. Nolan (2001) outlined a procedure of maximum likelihood for estimating stable parameters by approximation with a numerical quadrature. [7]

4. Data

This section describes the traffic data used for the analysis. All data are provided by the TRIM system of the Danish Road Directorate. [8] They measure the speed and traffic flows on some consecutive congested links of the Danish road network using cameras and automatic vehicle identification (number plate matching).

The *Frederikssundsvej* data are recorded on four consecutive links with a total length of 11.263 km. It is a main radial road in Greater Copenhagen connecting the city center and the north-west region. Figure 1 shows the location of the targeted road.

The data comprise minute-by-minute observations of average travel time on each link over about five months. We use data from weekdays between 6 a.m. and 10 p.m. during the period 16th January to 8th May, 2007, in the direction toward Copenhagen.

[7] The program package has already been implemented as "STABLE" (Robust Analysis, Inc., 2006). We use this package for our empirical analysis.

[8] "TRIM" is the Danish acronym for "Traffic Management on the Motorways around Copenhagen".

The road consists of four links: (1) Måløv Byvej; (2) Ballerup Byvej; (3) Herlev Hovedgade; and (4) Frederikssundsvej. We also analyze data concerning traffic that passes through all four consecutive links (5). Table 1 reports summary statistics of travel time data together with the computed plug-in bandwidths that were explained in the previous section. We also present summary statistics of travel time for each link in Table 2.

5. Empirical results and discussion

This section describes our empirical analysis for travel time distribution. All computations are carried out using Ox (Doornik, 2001), R (R Development Core Team, 2007) and STABLE (Robust Analysis, Inc., 2006).

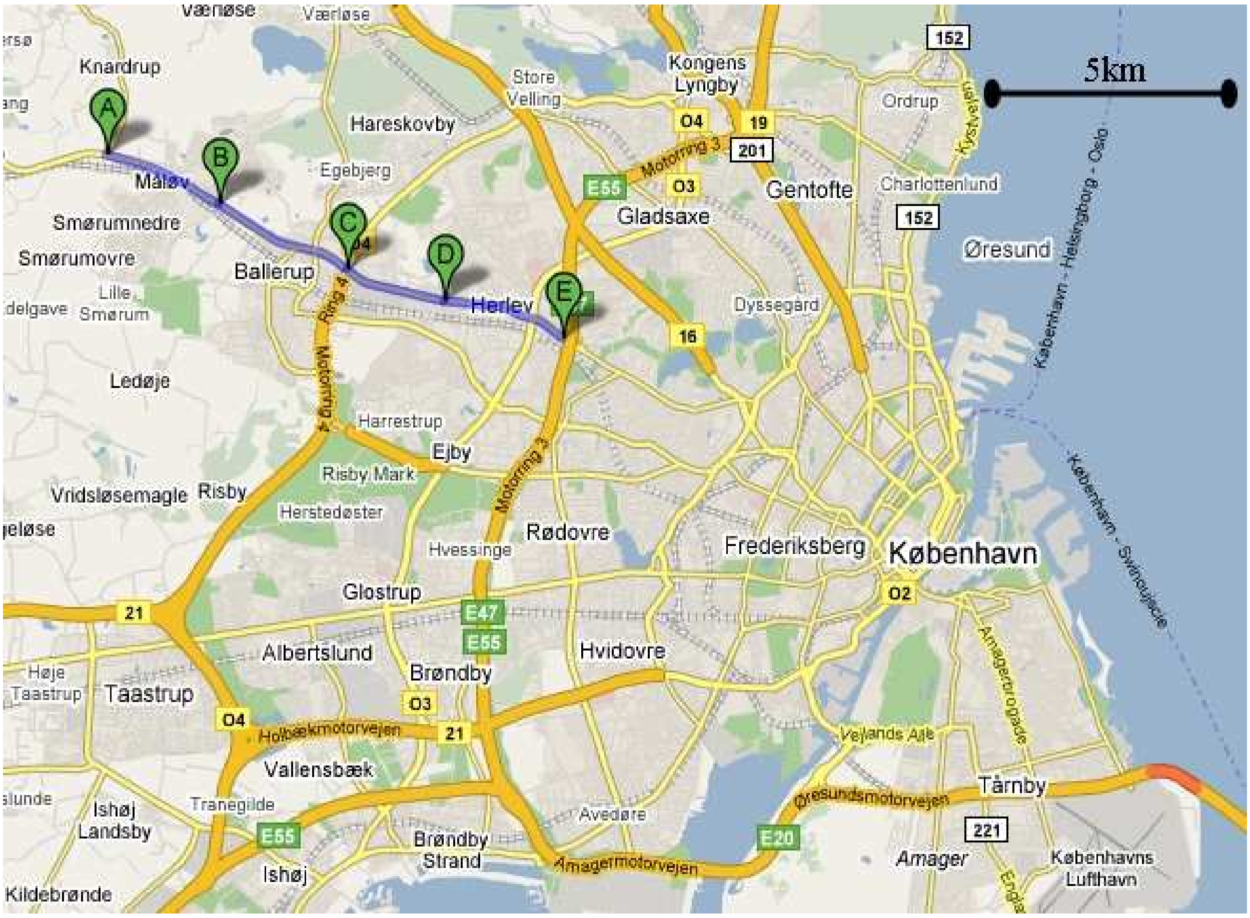

Figure 1: Targeted link of the urban road in Copenhagen (Frederikssundsvej)

Table 1: Outline of the urban road, observations and the computed plug-in bandwidths

Link ID	Direction	Length (km)	Obs.	$h_T^{plug,m}$	$h_T^{plug,cd}$	$h_t^{plug,cd}$
1	A $\rightarrow$ B	2.725	60669	32.9	47.5	0.162
2	B $\rightarrow$ C	3.279	59950	32	46.1	0.406
3	C $\rightarrow$ D	2.508	57759	32.1	46.2	0.183
4	D $\rightarrow$ E	2.751	54462	32.6	46.9	0.339
5	A $\rightarrow$ E	11.263	24271	37.9	53.1	0.895

Note: The unit for plug-in bandwidths is minute.

Table 2: Summary statistics of travel times (in minutes)

Link ID	Mean	S.D.	Min.	L.Q.	Median	U.Q.	Max.
1	2.967	0.957	0.98	2.49	2.69	3.14	24.6
2	4.854	2.395	1.55	3.45	3.94	5.22	27.4
3	3.037	1.074	0.1	2.38	2.66	3.3	19.5
4	4.442	1.967	1.4	3.16	3.84	5.05	28.59
5	15.399	4.543	8.76	12.15	13.83	17.67	47.5

5.1. Mean and scale regressions

Figure 2 shows the nonparametric kernel regression of mean travel time together with 95% confidence bands (upper panels) and the estimated interquartile range of travel times (lower panels) over the time of day. Both curves are smoothed using the plug-in bandwidths defined by (10) for mean and (11) for the interquartile range. In the two road links further from downtown (Figure 2 (a) and (b)), we see that there are distinct travel time peaks in the morning period. In contrast, the remaining links closer to the city center (Figure 2 (c) and (d)) show a peak in the mean travel time around 5 p.m. that would be caused by daily traffic congestion around the city center in the evening hours. As for the traffic data that ran the whole links (Figure 2 (e)), we only see the morning peak of the mean travel time. The narrow confidence bands for the mean travel time curves indicate that μ is quite precisely estimated because of our large data set.

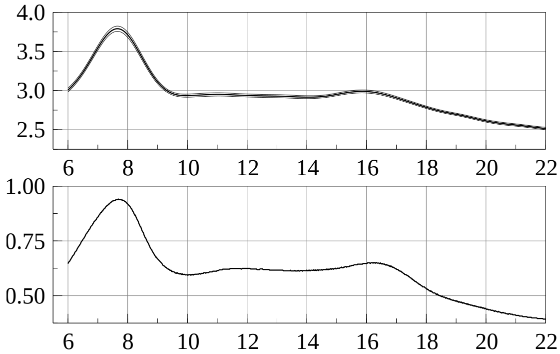

(a) Link 1

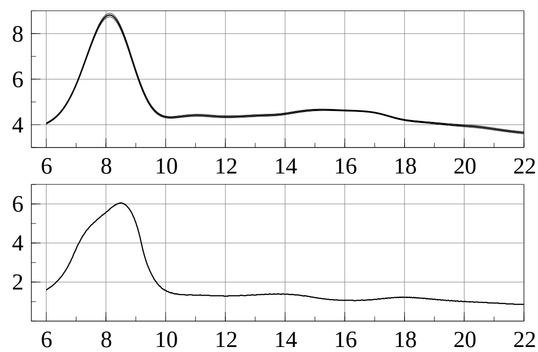

(b) Link 2

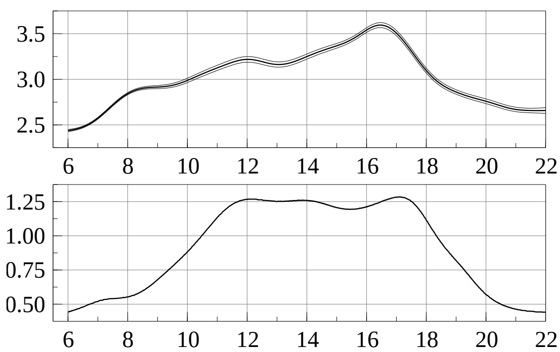

(c) Link 3

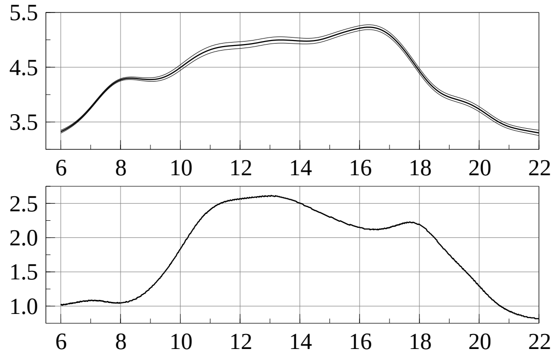

(d) Link 4

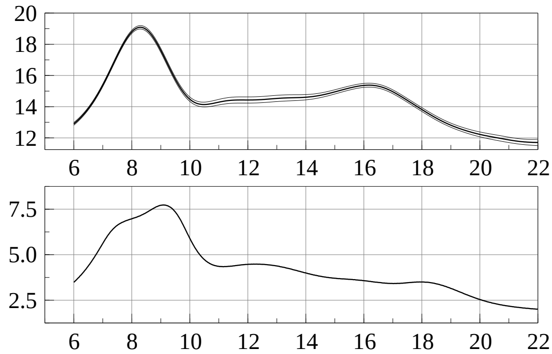

(e) Link 5

Figure 2: Mean regression (upper) and interquartile range regression (lower) of travel time over time of day

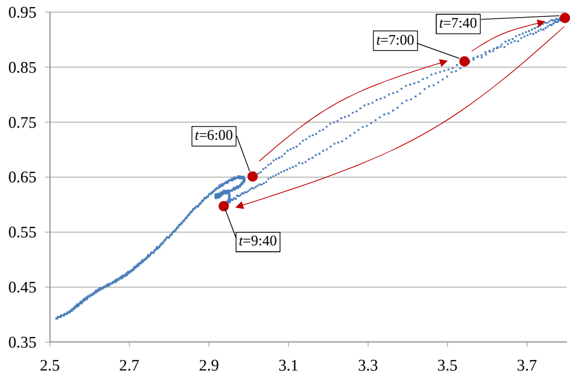

(a) Link 1

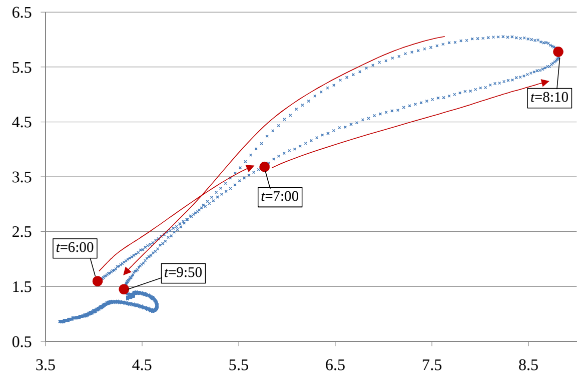

(b) Link 2

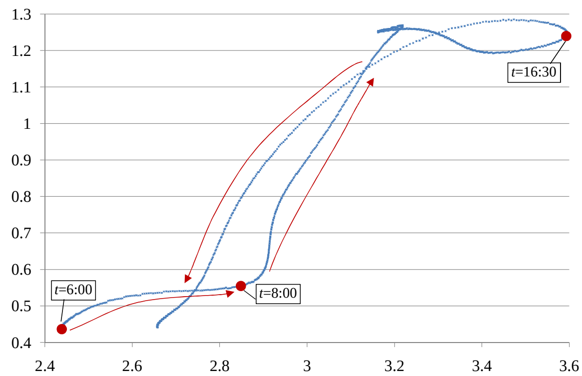

(c) Link 3

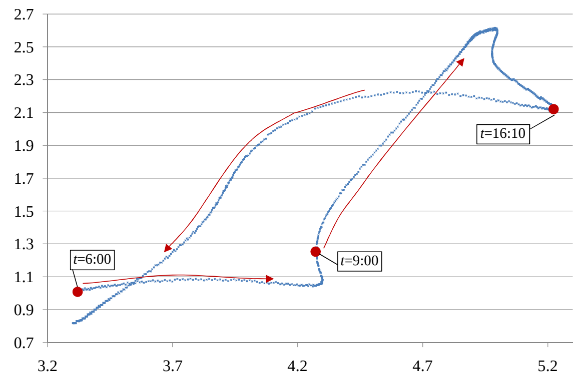

(d) Link 4

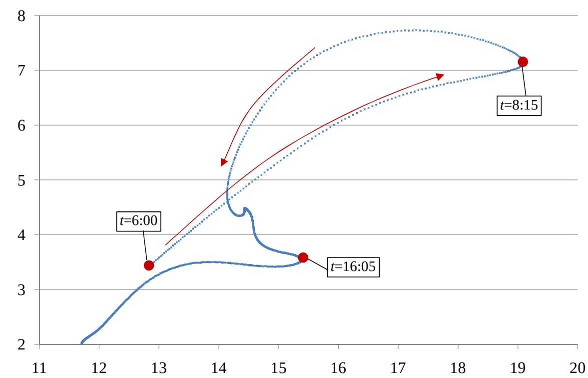

(e) Link 5

Figure 3: Scatter plot of mean and interquartile of travel times

In Figure 2, we see clear variation in $IQR(t)$ over the time of day. We also confirm the clear correlation between μ and IQR as it is evident from the scatter plot of IQR against μ in Figure 3. There are significant positive correlations between μ and IQR meaning that the larger the mean travel time, the larger the variation in travel time. In many cases, we also find that: (1) the variation in travel time measured as the interquartile range increases more slowly than the mean travel time; (2) they almost simultaneously reach their maximum in the peak period; and (3) the mean travel time decreases faster than the scale of it after the peak period. [9]

5.2. *Checking the standardized travel times conditional on time of day*

Next, we standardize the travel times following the procedure described in Section 3.3. Figure 4 presents the contours of the conditional CDFs of standardized travel time over the time of day. Each horizontal curve corresponds to a computed quantile (10% to 90%) of the standardized travel time on a given time of day. If standardized travel time is strictly independent of the time of day, all contour lines would be completely horizontal. We find that most of the estimated contour lines in every road link seem to be roughly horizontal across the day. In some road links, there exist infrequent but very big incidents such as serious traffic accidents, which result in extremely large travel times. The corresponding standardized travel time is large and this creates bumps of the contour lines of the larger quantiles in Figure 4, particularly during the morning or evening periods of traffic congestion. Although we see some unevenness in the contour lines for the larger (e.g., 90%) quantile, most of the contour lines seem to be about parallel. Hence, the essential assumption in the FK model that the standardized travel time is independent of the time of day would not be inappropriate to make as a rough approximation.

5.3. *Density estimation of standardized travel times*

Now we are able to estimate the unconditional standardized travel time distribution. We estimate the four parameters characterizing stable distributions using the numerical maximum likelihood estimation method (Nolan, 1997). The estimation procedure is carried out separately for each road link.

Table 3 outlines the estimation results. We also show the maximum likelihood estimates of the stable parameters for the data for the whole link (link 5) in the table. In every link, the estimates

[9] Fosgerau (2010) shows how this pattern arises due to the dynamics of congestion.

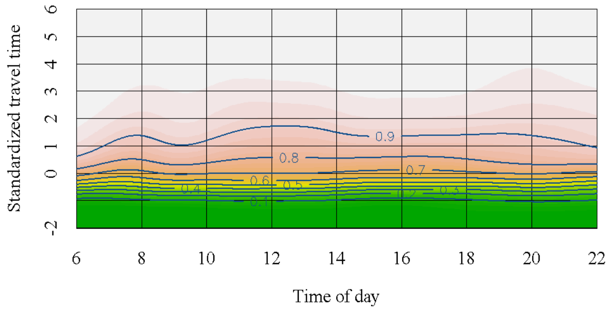

(a) Link 1

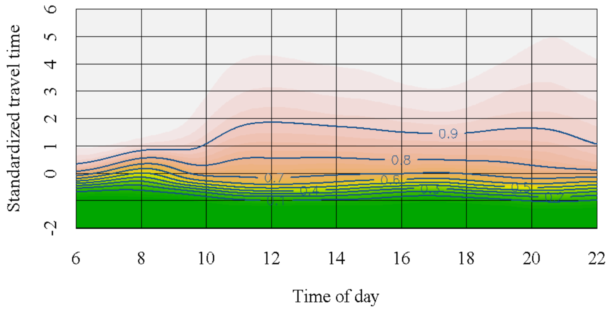

(b) Link 2

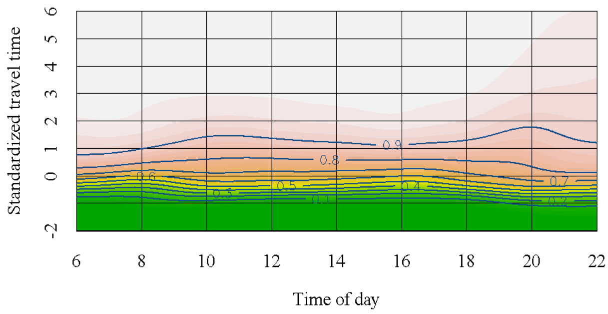

(c) Link 3

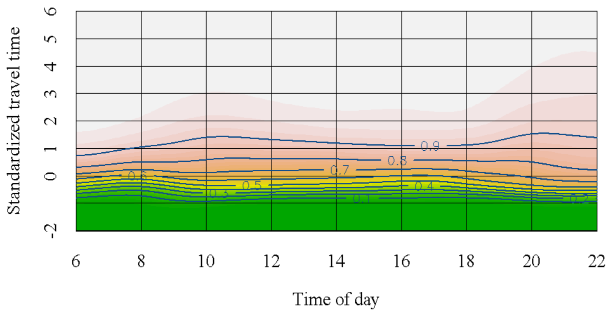

(d) Link 4

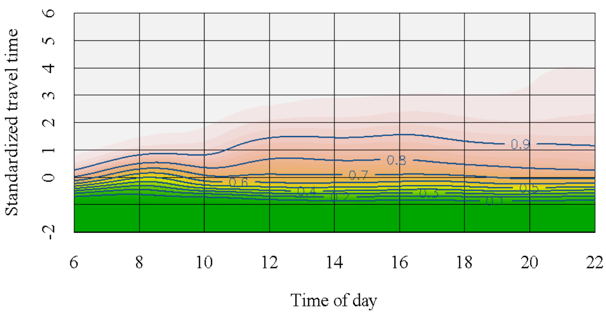

(e) Link 5

Figure 4: Conditional distribution of standardized travel times

Table 3: Estimated stable parameters

Link	α	β	γ	δ	$LL_{\max}$	$LL^{\alpha}_{\max}$	$-2(LL^{\alpha}_{\max} - LL_{\max})$	p-value	Obs.
1	1.1585	0.8824	0.3265	-0.528	-67600.5	-67605.2	9.43	0.002	60669
2	1.113	0.9089	0.2825	-0.5181	-59883.1	-59890.2	14.2	0.0002	59950
3	1.1385	0.9172	0.3153	-0.484	-61490.8	-61491.2	0.823	0.360	57759
4	1.118	0.99	0.3043	-0.4762	-55424.1	-55428.3	8.4	0.004	54462
5	1.3	1	0.3049	-0.3785	-21940.4	–	–	–	24271

of the four stable parameters are statistically significant. All estimated stability parameters ($\hat{\alpha}$) are significantly less than two (normal distribution), showing leptokurtosis in standardized travel times. If $0 < \alpha < 1$, the first moment of the stable distribution diverges to infinity. On the contrary, all of our estimates of the αs are significantly greater than one.

All estimates of the skewness parameter (β) are close to one: the upper bounds of the skewness parameter in stable distributions. This means that the estimated stable distributions are very skewed to the right. The estimates of the location parameter (δ) take similar negative values and the estimated scale parameter (γ) are also close to each other. The fitted stable distributions for these four consecutive links are shown in Figure 5 together with the data histogram. The bin width of each histogram is given by $3.5\sigma_X/n^{1/3}$ which is known as "Scott's choice rule" (Scott, 1979). The representation of data sets as histograms shows heavy tails on the right.

In Table 3, we observe that the estimated $\hat{\alpha}$ for the four consecutive links (1–4) take similar values with an average average of $\bar{\alpha} = 1/4\sum_{j=1}^{4}\hat{\alpha}_j = 1.1320$. We conduct a likelihood ratio test to check the equality of the stable parameter α across the four road links. To do this, we compute the maximal log likelihood of stable distributions ($LL^{\alpha}_{\max}$) under the restriction that $\alpha = \bar{\alpha}$ and compute the test statistic $-2(LL^{\alpha}_{\max} - LL_{\max})$ as shown in Table 3. Because of the very large sample size, the statistical power in our empirical analysis is quite strong. Hence, the null hypothesis that the stable parameter is equal to $\bar{\alpha}$ is rejected even at the 0.1% significance level ($< \chi^2_{\text{d.f.}=1} = 10.83$), except for link 2. We conclude that difference is statistically significant but not large.

We sketch the overall shapes of the estimated density curves in Figure 5. It seems that the estimated densities plots provide us with a stable distribution. Although the plotting results are likely to indicate the stability of the standardized travel time, it is less informative on the behavior

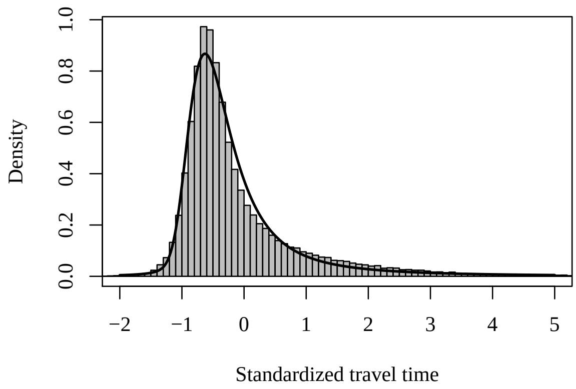

(a) Link 1

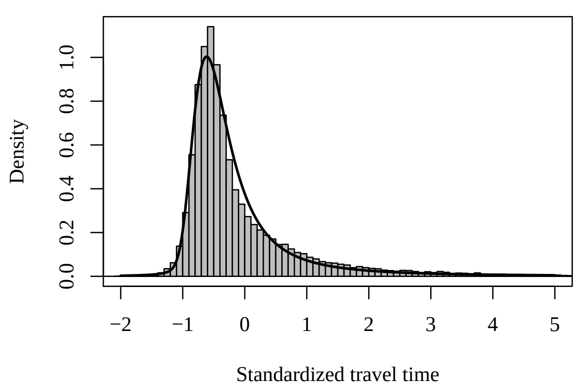

(b) Link 2

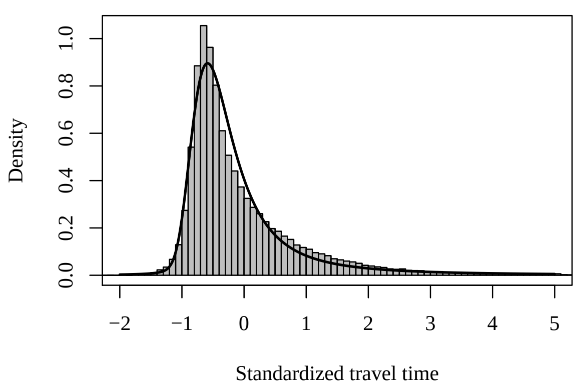

(c) Link 3

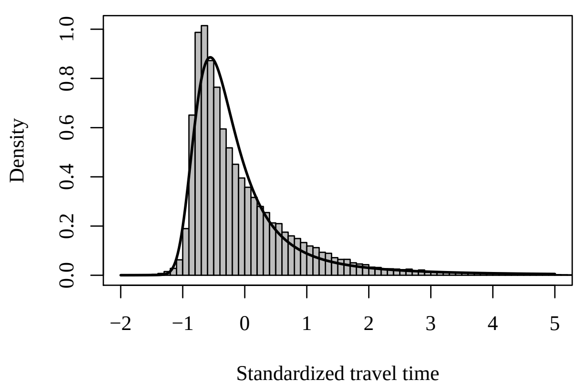

(d) Link 4

Figure 5: Fitting standardized travel times to stable distributions

of the tail probabilities. Figures 6 and 7 show variance stabilized P–P (probability–probability) plots (Michael, 1983) and Q–Q (quantile–quantile) plots of each data set respectively. Because too many data points add little to the plots, we show thinned P–P and Q–Q plots with 1,000 values. Nolan (2001) recommended using the variance stabilized P–P plots instead of standard P–P plots arguing that the use of the variance stabilized P–P plot is better than the standard P–P plot because it detects a poor fit near the extremes of the data.

The variance stabilized P–P plots show a reasonable fit around the modes for all data. However, we see in Figure 6 that there is a slight discrepancy between the data and the fitted distributions around the tail probabilities (i.e., 0 or 1) in all road links. This is more distinctive in the Q–Q plots in Figure 7. It can be seen that there is too much mass in the stable tails compared to the empirical distribution.

5.4. Computing H

We further compute the value of the functional H, defined by (3) for various values of η/λ under different distributional assumptions on standardized travel times. We consider three distributions: (1) normal; (2) empirical; and (3) stable. We compute H for the normal distributions using the sample mean and standard deviation for each road link. The H for stable distributions are computed on the basis of the maximum likelihood estimates of stable parameters shown in Figure 3. [10] The result of the computation is illustrated in Figure 8. Figure 8 also contains the results of some modified Hs corresponding to truncations of the fitted stable distributions (see Section 6.2 later).

Figure 8 summarizes the result of computing H. There are differences in H by distributional assumptions as well as across different road segments. The changes in H for different η/λ under normality are more distinctive than the other two distributional assumptions. For example, the normal H for $\eta/\lambda = 0.5$ in the road link 1 is 0.538. This is nearly 1.54 times larger than the empirical H for $\eta/\lambda = 0.5$. On the contrary, the normal H for $\eta/\lambda = 0.05$ in the road link 1 is 0.140 and smaller than the empirical H for $\eta/\lambda = 0.05$. Similar tendencies can be seen in other road links. On the other hand, H for the empirical and the stable distribution do not change so much with η/λ.

[10] The numerical integral in (3) for normality and stability is computed using the trapezoidal rule.

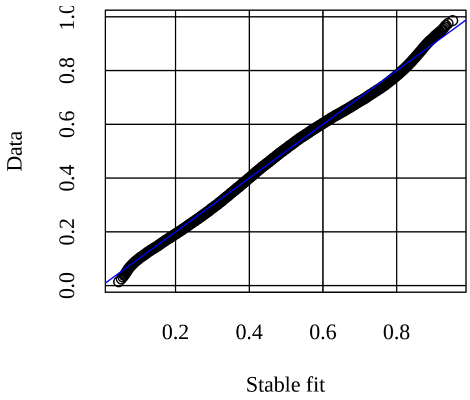

(a) Link 1

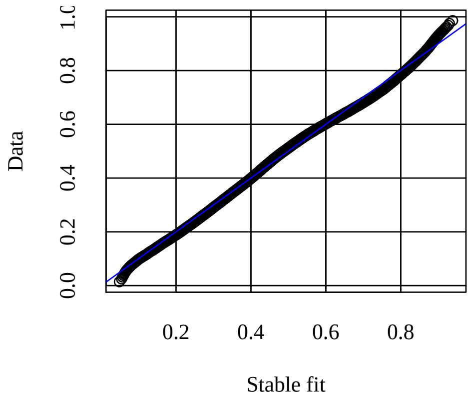

(b) Link 2

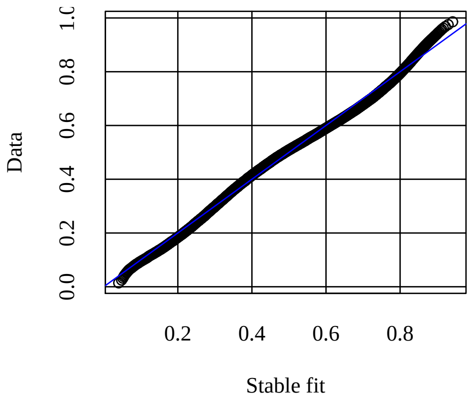

(c) Link 3

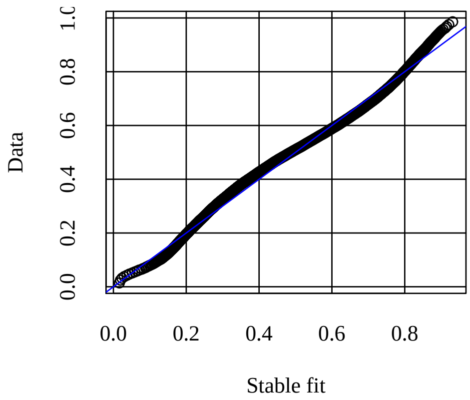

(d) Link 4

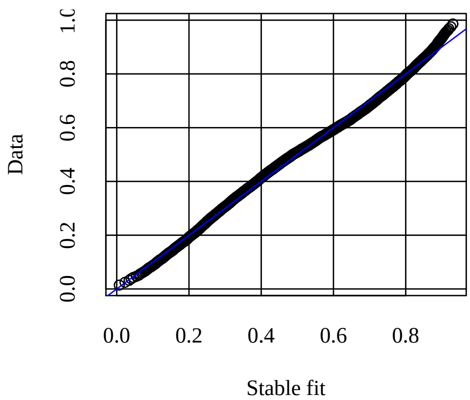

(e) Link 5

Figure 6: Variance stabilized P–P plot of stable distributions

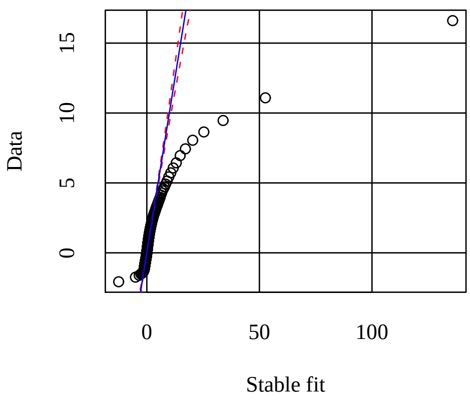

(a) Link 1

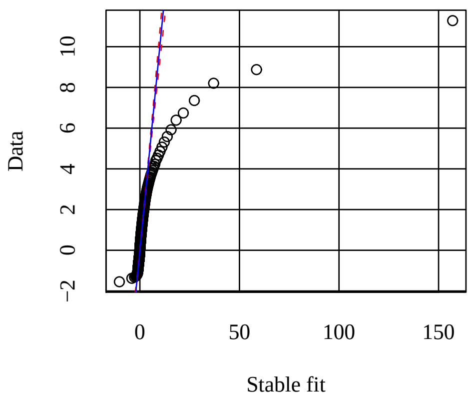

(b) Link 2

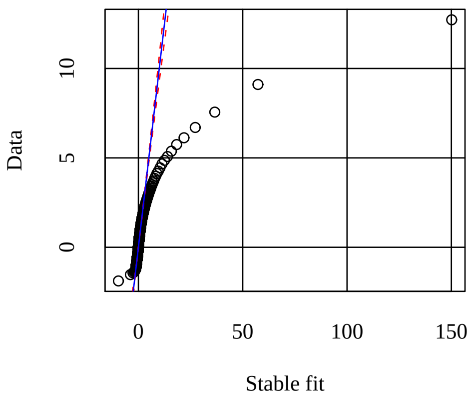

(c) Link 3

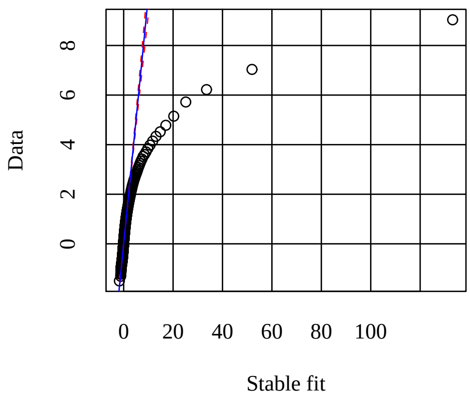

(d) Link 4

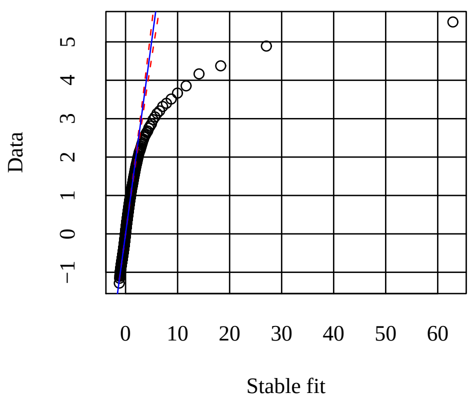

(e) Link 5

Figure 7: Q–Q plot of stable distributions

The computed H for the stable distributions in any road links are larger than for normal and empirical. We find that the computed H are notably larger than that for the empirical. There exists a significant difference in the right tail probabilities between the stable and empirical distributions as shown in Figures 6 and 7. These difference in the tails between the empirical and stable would be influential when Hs are computed.

6. Discussion

The main purpose of the present paper is to investigate to which degree the empirical characteristics of travel time distributions conform to the requirements of the FK model in applications of valuing travel time variability.

6.1. Independence of standard travel time distributions and the time of day

The first fundamental hypothesis of the FK model is that the distribution of standardized travel time, after removing changes in the mean and scale of travel time across the time of day, is independent of the time of day. To investigate this hypothesis, we analyzed traffic data that were collected on an urban road over a long period using nonparametric techniques.

The nonparametric regression results for the mean and the interquartile range of travel times given a time of day (Figures 2 and 3) indicate that the mean and the scale of travel time are not constant over the time of day in every road link. This is expected since traffic varies with the time of day and such variation leads to variation in the mean and scale of travel times.

To check the independence assumption of standardized travel times against the time of day, we studied the nonparametric distribution of standardized travel times conditional on the time of day. Strict independence would require all contours of the probability distribution being completely horizontal. Figure 4 shows that in every road link the contour lines for the probability distributions of standardized travel times are not very different from horizontal. The fluctuations are largest at the highest quantiles and may be due to a small number of incidents. So we would feel justified in accepting that standardized travel time is roughly independent of the time of day.

6.2. Fitting standard travel times to stable distributions

The second hypothesis we investigate is that the standardized travel time follows a stable distribution. If this hypothesis is supported, practical applications are facilitated by the favorable

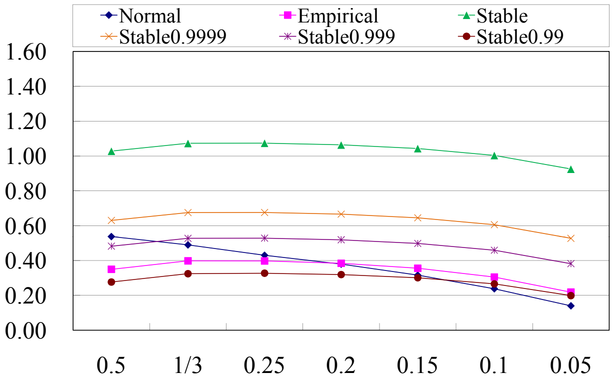

(a) Link 1

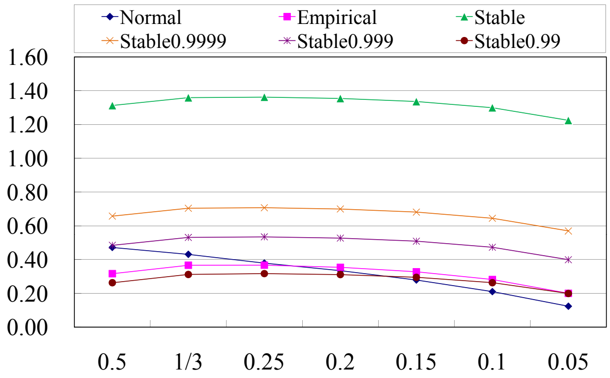

(b) Link 2

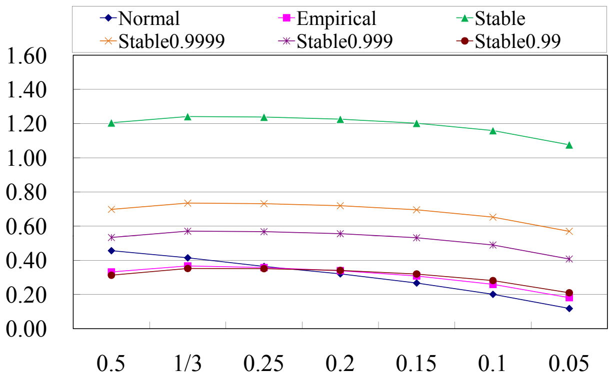

(c) Link 3

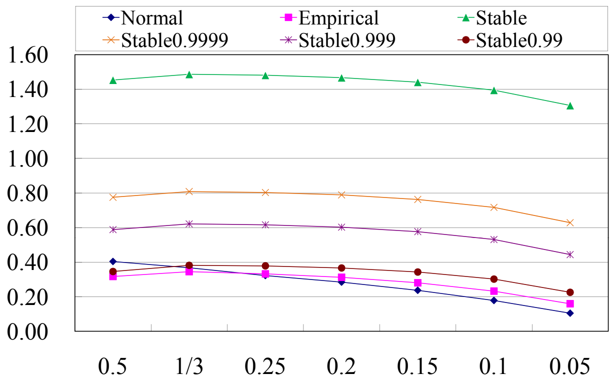

(d) Link 4

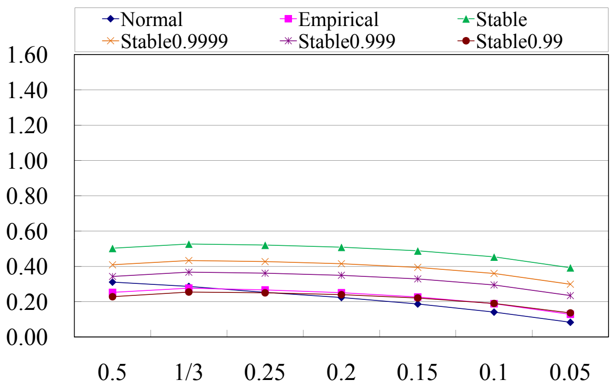

(e) Link 5

Figure 8: Computed H at various values of η/λ

properties of stable distributions. To check stability, we estimate stable parameters for each road link using maximum likelihood and evaluated some diagnostics.

The parameter estimates (Table 3) and the plotted stable densities (Figurs 5) show that the data sets of the standardized travel times for any road links are far from normal. All skewness parameters are estimated to be close to the upper bounds ($\hat{\beta} = 0.8824 \sim 1.0000$) indicating that the distributions are very skewed to the right. With this skewness, the estimated stable densities fit the data around the modes of the distributions as shown in Figure 5. Furthermore, the estimated stability parameters are scattered around 1.1320 as explained in Section 5.3. The estimates are closer to the stability parameter of a Cauchy distribution ($\alpha = 1$) than to a normal distribution ($\alpha = 2$).

These results might be caused by the typical characteristics of travel times on urban roads: (1) there would exist a lower bound of travel time because of physical and environmental constraints; and (2) the maximal standardized travel times, on the other hand, would be very large because there would be a small but significant possibility that severe incidents might occur.

As for the behavior of the tails on the other hand, there seem to be significant differences between the data and the estimated stable distributions. The Q–Q plots in Figure 7 show that the extreme tails of the standardized travel time data are thinner than the stable densities. Thus the fitted stable distributions tend to overestimate the tail probabilities. This fact significantly affects the computational results of the functional H. As shown in Figure 8, the value of H obtained for the stable distributions is larger than for the empirical distribution on each road link.

This difference is related to the fact that stable but non-normal distributions have infinite variances. In our empirical results, the estimated stability parameters are all near 1.1320, and hence the distributions are far from normal. In contrast, empirical travel times are bounded and so have finite variance. This would provide much larger tail probabilities in the fitted stable densities than in the empirical distributions. In other words, the fitted stable distribution will predict too high probability of outrageously high travel times.

6.3. Assumption of the maximum travel time in the distributions

A possibility for circumventing the above-mentioned problem in the use of stable distributions is to reconsider the scheduling model by imposing a "maximum" travel time when the traveler evaluates the expected cost. We assume that the traveler only considers travel times below this maximum. This assumption corresponds to replacing the upper integral limit in (1) by a finite positive number.

Denote the maximum of standardized travel times as $X_{\max}$. Furthermore, denote the probability that a standardized travel time is equal to or less than $X_{\max}$ as $p_{X_{\max}} = \text{Prob}(x \leq X_{\max}) = \Phi(X_{\max})$. The scheduling model (1) is rewritten as:

$$EC^* = \min_D EC(D,T) = \min_D \left[\eta D + \lambda \int_{\frac{D-\mu}{\sigma}}^{X_{\max}} (\mu + \sigma x - D)\phi(x)\mathrm{d}x + \omega\mu \right].$$

The first order condition of the scheduling model (2) is replaced by the following similar formula:

$$D' = \mu + \sigma\Phi^{-1}\left(p_{X_{\max}} - \frac{\eta}{\lambda}\right).$$

Furthermore, the new functional H' becomes:

$$H'\left(\Phi, \frac{\eta}{\lambda}, p_{X_{\max}}\right) = \int_{p_{X_{\max}}-\frac{\eta}{\lambda}}^{p_{X_{\max}}} \Phi^{-1}(\nu)\mathrm{d}\nu. \tag{14}$$

Notice that $H'\left(\Phi, \frac{\eta}{\lambda}, 1\right)$ tends to $H\left(\Phi, \frac{\eta}{\lambda}\right)$ as $\lim_{X_{\max}\to\infty} p_{X_{\max}} = 1$.

The choice of $p_{X_{\max}}$ is somewhat arbitrary. We can however check the resulting H' for stable distributions to empirical ones. We can also find appropriate values of $p_{X_{\max}}$ by checking the goodness-of-fit of stable H's with respect to the empirical Hs.

Figure 8 presents the computed H's with three different values of $p_{X_{\max}}$ for the fitted stable distributions. If $p_{X_{\max}} = 99.99\%$, for example, the probability that the standardized travel time becomes greater than $X_{\max}$ is 0.01% and travelers are assumed to disregard such a large travel time in their scheduling choice. This result shows that the restriction on the upper limit integral in (14) would significantly reduce the deviations from the empirical Hs. In our applications, we expect that the appropriate $p_{X_{\max}}$ would be between 99.0% and 99.9%.

6.4. Equality of stability parameters

As shown in Table 3, the estimated $\hat{\alpha}$s do not differ much from each other. Thus, it may not be inappropriate to assume that the standardized travel time distributions share a common stability parameter across the different road links. From the comparison of the stable parameters for the links 1–4 and the one for the link 5, we see that these difference are not so large but significant bacause of large samples. For this result, we speculate that some correlation might exist among standardized travel times across different road links. Because the traffic congestion upstream propagates to downstream, it is likely that the travel times for the consecutive roads are positively correlated.

Recall that the standardized travel times should be independent in the convolutions of stable distributions. To check this informally, we have plotted the pair of standardized travel times for the consecutive two links and have drawn a bivariate joint density in Figure 9. The pairs are identified based on the date and the time of day for each link. We find that there does not seem to be significant conditional dependence between the two standardized travel times and so the independence assumption of standardized travel times could be reasonable.[11]

7. Concluding remarks

This paper has analyzed some empirical characteristics of the travel time distribution on an urban road with the purpose of checking the degree to which the travel time distribution conforms to the assumption in the Fosgerau and Karlström (2010) model. A number of nonparametric techniques were employed to estimate the distribution of standardized travel time conditional on the time of day.

First, we found that the FK assumption that the standardized travel time is independent of the time of day seems reasonable as an approximation. This is crucial for the application of the FK model.

Second, the standardized travel time distribution is far from normal but close to a stable distribution. Like the normal distribution, the stable distribution arises in a central limit theorem, but requires weaker assumptions on the variances of the random variables of which it is a limit. The stable distribution is able to reproduce the high skewness and fat tails of empirical travel time distributions.

Third, the extreme right tails of the stable distribution are fatter than in the empirical distributions. This suggests that the stable distribution is not appropriate as a description of extreme delays. In reality, these are bounded from above; this is not true of the stable distribution. This suggests using some truncation of the stable distribution. Truncating the stable distribution yields values of the standardized mean lateness factor H that are close to the empirical values.

[11] Some statistical tests (e.g. Su and White (2007)) would be applicable to check this formally.

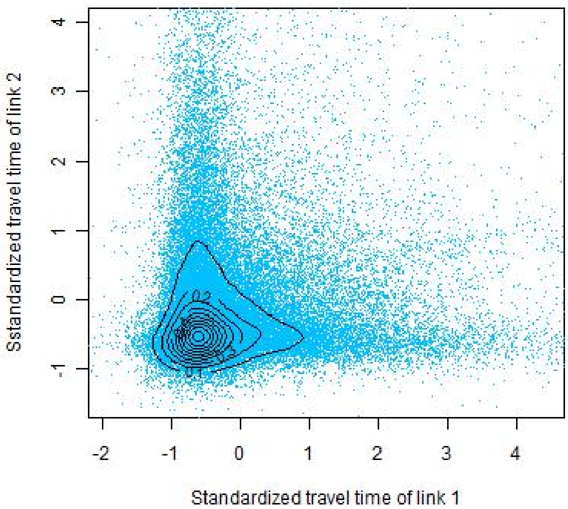

(a) Links 1 and 2 (54,310 points)

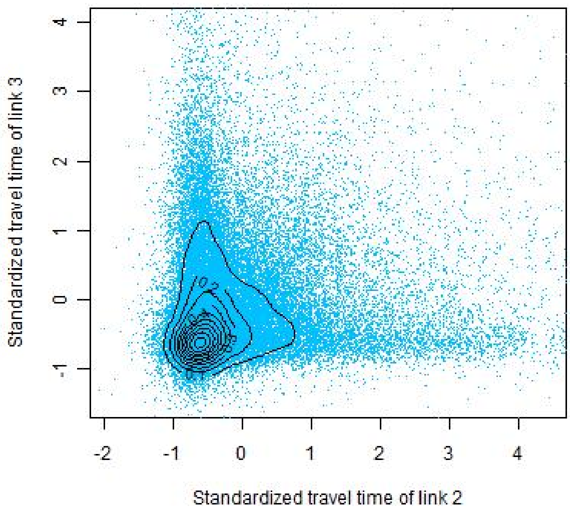

(b) Links 2 and 3 (50,499 points)

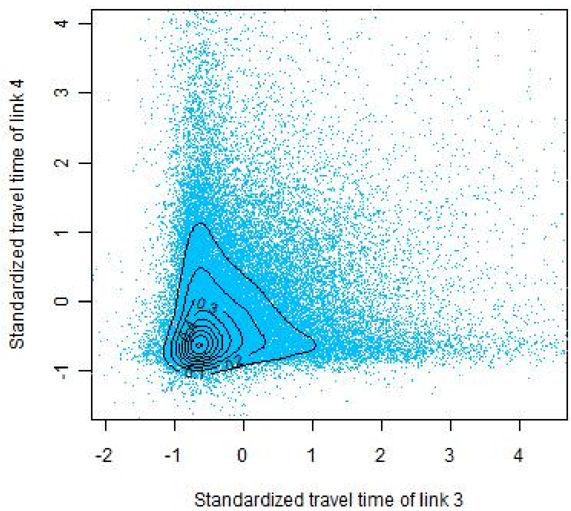

(c) Links 3 and 4 (52,649 points)

Figure 9: Scatter plot and joint density of two standardized travel times

Fourth, the stability parameter α seems to be roughly constant across road links. Furthermore, standardized travel times seem to be about independent across links. Therefore, computing the travel time distribution for a route as the convolution of travel times on individual links may be considered reasonable for practical purposes.

Acknowledgements

The authors are thankful for comments from Ken Small and Anders Karlström. Mogens Fosgerau has received support from the Danish Social Science Research Council. Part of the work was carried out during a research stay of Daisuke Fukuda at the Technical University of Denmark, funded by the Kajima Foundation. We also thank the Tokyo Institute of Technology for use of the "TSUBAME" grid supercomputer for numerical computations.

References

Bates, J., Polak, J., Jones, P., Cook, A., 2001. The valuation of reliability for personal travel. Transportation Research Part E: Logistics and Transportation Review 37 (2–3), 191–229.

Brownstone, D., Small, K., 2005. Valuing time and reliability: assessing the evidence from road pricing demonstrations. Transportation Research Part A: Policy and Practice 39 (4), 279–293.

Doornik, J. A., 2001. Ox: An Object-Oriented Matrix Language. London.

Fosgerau, M., 2010. On the relation between the mean and variance of delay in dynamic queues with random capacity and demand. Journal of Economic Dynamics and Contril 34 (4), 598–603.

Fosgerau, M., Engelson, L., forthcoming. The value of travel time variance. Transportation Research Part B: Methodological.

Fosgerau, M., Karlström, A., 2010. The value of reliability. Transportation Research Part B: Methodological 43 (8–9), 813–820.

Gnedenko, B. V., Kolmogorov, A. N., 1954. Limit Distributions for Sums of Independent Random Variables. Addison-Wesley.

Hall, P., Racine, J., Li, Q., 2004. Cross-validation and the estimation of conditional probability densities. Journal of the American Statistical Association 99 (468), 1015–1027.

Härdle, W., 1990. Applied Nonparametric Regression. Econometric Society Monograph Series. Cambridge University Press.

Ichimura, T., Fukuda, D., 2010. A fast algorithm for computing least-squares cross-validations for nonparametric conditional kernel density functions. Computational Statistics & Data Analysis 54 (12), 3404–3410.

Lam, T. C., Small, K., 2001. The value of time and reliability: measurement from a value pricing experiment. Transportation Research Part E: Logistics and Transportation Review 37 (2-3), 231–251.

Li, Q., Racine, J., 2007. Nonparametric Econometrics: Theory and Practice. Princeton University Press.

Michael, J. R., 1983. The stabilized probability plot. Biometrika 70 (1), 11–17.

Nolan, J. P., 1997. Numerical calculation of stable densities and distribution functions. Stochastic Models 13 (4), 759–774.

Nolan, J. P., 2001. Maximum likelihood estimation of stable parameters. In: Barndorff-Nielsen, O. E., Mikosch, T., Resnick, S. I. (Eds.), Lévy Process: Theory and Applications. Birkhäuser, Boston, pp. 379–400.

Nolan, J. P., in press. Stable Distributions – Models for Heavy Tailed Data. Birkhäuser, Boston.

Noland, R. B., 1997. Commuter responses to travel time uncertainty under congested conditions: expected costs and the provision of information. Journal of Urban Economics 41 (3), 377–406.

Noland, R. B., Polak, J. W., 2002. Travel time variability: a review of theoretical and empirical issues. Transport Reviews 22 (1), 39–54.

Noland, R. B., Small, K. A., 1995. Travel-time uncertainty, departure time choice, and the cost of morning commutes. Transportation Research Record 1493, 150–158.

Noland, R. B., Small, K. A., Koskenoja, P. M., Chu, X., 1998. Simulating travel reliability. Regional Science and Urban Economics 28 (5), 535–564.

Pagan, A., Ullah, A., 1999. Nonparametric Econometrics. Cambridge University Press.

R Development Core Team, 2007. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.

Robust Analysis, Inc., 2006. User Manual for STABLE 4.0 S–Plus/R Version.

Scott, D. W., 1979. On optimal and data-based histograms. Biometrika 66 (3), 605–610.

Small, K. A., 1982. The scheduling of consumer activities: work trips. American Economic Review 72 (3), 467–479.

Small, K. A., Winston, C., Yan, J., 2005. Uncovering the distribution of motorists' preferences for travel time and reliability. Econometrica 73 (4), 1367–1382.

Su, L., White, H., 2007. A consistent characteristic function-based test for conditional independence. Journal of Econometrics 141 (1), 807–834.

Vickrey, W. S., 1973. Pricing, metering, and efficiency using urban transportation facilities. Highway Research Record 476, 36–48.

Zolotarev, V., 1986. One-Dimensional Stable Distributions. American Mathematical Society.