TI 2011-180/3
Tinbergen Institute Discussion Paper

Door-to-Door Travel Times in RP Departure Time Choice Models – An Approximation Method based on GPS Data

*Stefanie Peer**
Jasper Knockaert
*Paul Koster**
*Yin-Yen Tseng**
*Erik Verhoef**

Faculty of Economics and Business Administration, VU University Amsterdam.

** Tinbergen Institute.*

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Door-to-door travel times in RP departure time choice models: An approximation method using GPS data

October 2013

Stefanie Peer[a,b,*], Jasper Knockaert[a], Paul Koster[a,c], Yin-Yen Tseng[a,d], Erik T. Verhoef[a,c]

[a] *Department of Spatial Economics, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands*

[b] *Institute for the Environment and Regional Development, Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria*

[c] *Tinbergen Institute, Gustav Mahlerplein 117, 1082 MS Amsterdam, The Netherlands*

[d] *School of Public Finance and Taxation, Southwestern University of Finance and Economics, Liutai Avenue 555, Wenjiang District, Chengdu, Sichuan 611130, China*

Abstract

A common way to determine values of travel time and schedule delay is to estimate departure time choice models, using stated preference (SP) or revealed preference (RP) data. The latter are used less frequently, mainly because of the difficulties to collect the data required for the model estimation. One main requirement is knowledge of the (expected) travel times for both chosen and unchosen departure time alternatives. As the availability of such data is limited, most RP-based scheduling models only take into account travel times on trip segments rather than door-to-door travel times, or use very rough measures of door-to-door travel times. We show that ignoring the temporal and spatial variation of travel times, and, in particular, the correlation of travel times across links may lead to biased estimates of the value of time (VOT). To approximate door-to-door travel times for which no complete measurement is possible, we develop a method that relates travel times on links with continuous speed measurements to travel times on links where relatively infrequent GPS-based speed measurements are available. We use geographically weighted regression to estimate the location-specific relation between the speeds on these two types of links, which is then used for travel time prediction at different locations, days, and times of the day. This method is not only useful for the approximation of door-to-door travel times in departure time choice models, but is generally relevant for predicting travel times in situations where continuous speed measurements can be enriched with GPS data.

Keywords: valuation of travel time and schedule delays, door-to-door travel times, departure time choice, revealed preference (RP) data, GPS data, geographically weighted regression (GWR)

*Corresponding author: speer@wu.ac.at (Stefanie Peer)

1. Introduction

Congestion tends to be strongly related to the time of the day, mainly because commuters have similar work start and ending times. Therefore, in order to evaluate policies that mitigate congestion, models are needed that take into account that (expected) travel times vary over the time of the day. The most prominent modeling approach that takes into account these variations is the so-called bottleneck model, which goes back to the early work of Vickrey (1969), and has been adapted for empirical estimation by Small (1982). It assumes that travelers optimize their departure time choices, trading off the costs of travel time against the costs of arriving earlier or later than their preferred arrival time. Empirical estimates of the drivers' willingness to pay for reductions in travel times and schedule delays can be derived from data on hypothetical and actual scheduling choices, which are usually referred to as stated preference (SP) and revealed preference (RP) data, respectively.[1] The departure time choice models (or more generally, scheduling models) are mostly estimated using discrete choice models, re-formulating the continuous departure time choice problem as a discrete problem with a finite number of alternative departure times.

SP data have the advantage that the researcher has full control over the choice set, the attribute levels of the choice alternatives, and the extent of correlation between them. On the downside, SP-based estimation results may depend on the design of the experiment and, specifically, on whether the respondents consider the attributes of the presented choice alternatives realistic. Otherwise, biases due to the hypothetical nature of the choice task may arise (e.g. Hensher, 2010). RP data, in contrast, are based on observed behavior in a real-life setting, and therefore by definition do not suffer from hypothetical biases. Nevertheless, they are used in relatively few studies, mainly because they require detailed, high-quality data, which are usually difficult to collect. One aspect is that there are only few real-life situations that incorporate time-of-day varying financial components, which in turn are required for identifying the monetary valuations of time and schedule delays. And if they exist, they often imply a strong correlation between relevant attribute values (e.g. dynamic tolls that increase when travel times increase), which renders an accurate estimation of the underlying trade-offs impossible. A second aspect is that even if an adequate RP experiment exists, it is still difficult to observe the relevant attribute values of the choice alternatives, and often, also the choices themselves. This ambiguity is thus in sharp contrast to SP data, where the choice set and the attribute levels are determined by the researcher.

This paper focuses on the second aspect, and in particular on a data requirement evident in RP-based departure time choice models: It concerns door-to-door travel times. Clearly, door-to-door travel times are the relevant travel time measure, based on which travelers can be expected to take their scheduling decisions. However, so far, many studies that estimate departure time choice models from RP data do not take into account door-to-door travel times, or only very rough measures thereof. This is mainly due to the poor availability of travel time data, especially outside the main road network. The problem of lacking travel time data is augmented by the fact that travel times need to be known not only for the chosen but also for the unchosen departure time alternatives. To our knowledge, no research has yet been undertaken that explicitly and systematically surveys the effects of using imprecise measures of (expected) travel times on the valuations of time and schedule delays.

[1] Overviews of relevant empirical research can for instance be found in Brownstone and Small (2005), Tseng et al. (2005) and Li et al. (2010).

We develop a model that allows for the (ex-post) approximation of driver-, day- and time-of-day-specific door-to-door travel times, combining loop detector and GPS data. The travel times resulting from this model are then used to determine the attributes of the departure time alternatives. We show that less precise travel time definitions, which ignore the spatial and temporal variation of speeds on some parts of the door-to-door trip, may lead to biases in the value of time. A straightforward reason is that ignoring the spatial variation of speeds typically leads to an imprecise measurement of travel times, which in turn affects the travel time coefficient to be biased downwards. However, there is a second, more subtle reason as well: Not accounting for the temporal variation of speeds means that the tendency of speeds to be correlated across links is ignored. In most situations a positive correlation is present, for instance because speeds are low on large parts of the road network during morning and evening peaks (e.g. Hackney et al., 2007; De Fabritiis et al., 2008). If these correlations are ignored, the differences between peak and off-peak travel times are underestimated, and as a result, the (absolute) time coefficient, and thus the value of time will be overestimated.

We gathered the RP data used in this paper during a real-life peak avoidance experiment (*Spitsmijden* in Dutch) conducted in the Netherlands. Participants were eligible for a monetary reward if they avoided traveling on a specific highway link (which we refer to as 'C1–C2' as it is confined by two cameras, C1 and C2) during the morning peak. In a first phase of the experiment, 230 participants got equipped with GPS devices yielding door-to-door travel time measurements. The second phase of the experiment included approximately 2000 participants and was conducted over a period of multiple months. During that second phase only travel times along the C1–C2 link were measured. For the scheduling models estimated in this paper, we consider the observed departure time behavior of (a subgroup of 428) participants of the second phase. We will then use the GPS data gathered during the first phase to approximate door-to-door travel times faced by participants of the second phase.

In order to approximate driver-, day and time-of-day-specific door-to-door travel times, we employ geographically weighted regression (GWR). We use this method to explicitly take into account travel time correlations across trip parts, in particular between the links for which only infrequent GPS observations are available (home-C1, C2–work) and the link along which travel times are observed continuously by means of loop detectors (C1–C2). Figure 1 shows a schematic map on the relative location and (data-)properties of these links. GWR allows for spatial heterogeneity in the correlation pattern. So, for a given speed on the C1-C2 link, predicted speeds along the home–C1 (C2–work) links are probably more similar between home (work) locations that are located closer to each other, and can differ more strongly between locations that are further apart. Figure 2 underlines the relevance of this approach. It plots the starting (home) and end (work) locations across space. The color scale indicates the average speed measured from the plotted home locations to C1, and from C2 to the work locations.[2] Even without distinguishing between different times of the day, and therefore recurrent congestion levels, the spatial variation of speeds remains

[2]Note that many of these 'home' and 'work' locations are located along the highway (the East-West connection passing through C1 and C2, along which many of the data points can be found). This finding may be attributable to drivers who only switched on their GPS device after departing from home, or switched it off before arriving at work. Another explanation is that in order to warrant the quality of the GPS data, we define various criteria that need to be fulfilled for recognizing specific GPS observations as being part of a trip through C1 and C2. As a consequence, we may conservatively assume a trip to start at a location closer to C1, or to end at a location closer to C2 than it actually did in reality, resulting in the possibility that the start or end location of a trip are located along the highway.

Figure 1: Schematic map of the experimental setup

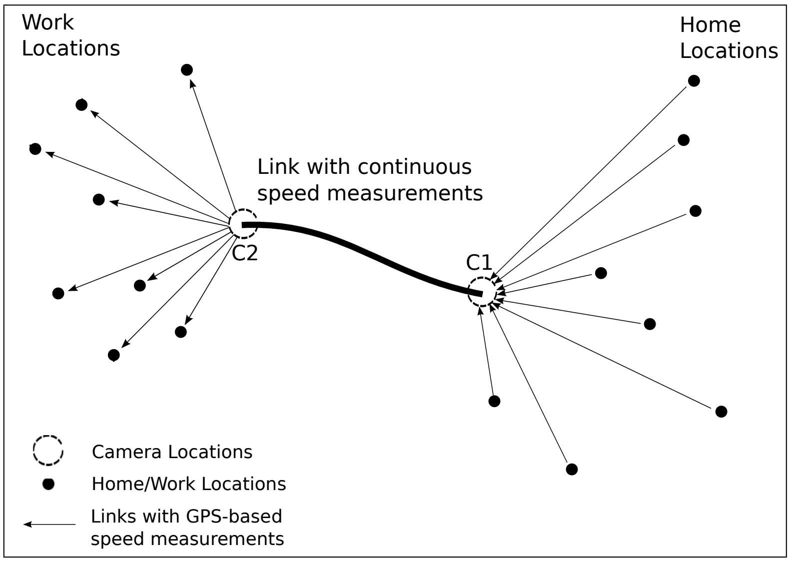

substantial. A clearly visible pattern is that home locations closer to the main highway (the East-West connection of data points, which crosses C1 and C2) face relatively high speeds, probably because they cover a higher share of their commute on the highway. Although the highway may be congested during peak hours, it allows for high travel speeds during off-peak hours.

Figure 2: Speed observations

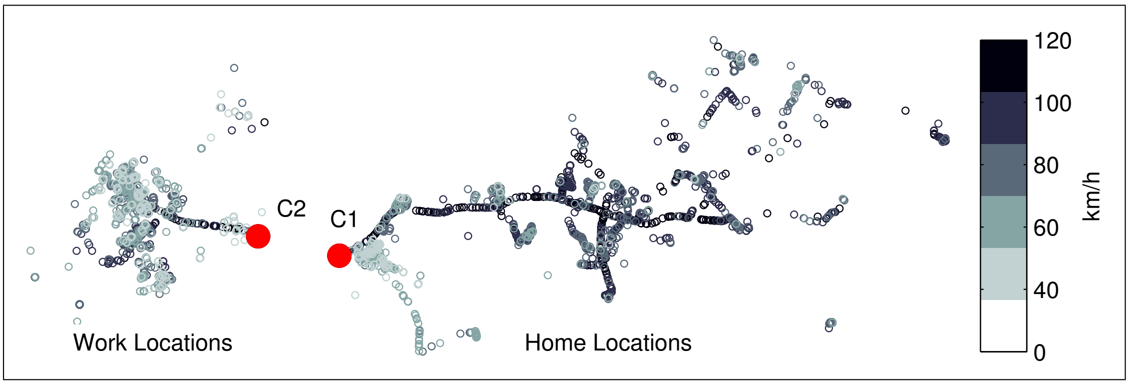

We use a combination of loop-detector and GPS data to approximate door-to-door travel times, although in principle also reported travel times or travel times generated from network models could have been used. However, both alternative data sources have in common that they are usually not available at a sufficiently fine temporal scale that is required to represent concisely the trade-offs that drivers face between departure time choice alternatives.[3] Furthermore, reported and network travel times are usually not day-specific. This means that they cannot be used for modelling travel

[3]In general, caution should prevail with the use of reported travel times as these might be systematically biased. For example, we found that travel times reported by participants of the Spitsmijden experiment were on average overreported with a factor 1.5 (Peer et al., 2013).

time variability or day-specific travel time expectations, whereas both of these features are present and show to be relevant in the models estimated in this paper.

The methodology developed here for approximating travel times on links where only sparse GPS measurements are available is not only useful in the context of departure time choice modeling. It can in principle be applied in any situation where, in addition to GPS data, continuous speed measurements are available on one or more links in the proximity. Another potential application of the method presented in this study is to derive measures of travel time and travel time variability that can be used as an input for defining the attribute values in SP choice models such that respondents perceive them as more realistic. As argued by Hensher (2010), this way one can still take advantage of the favorable properties of SP experiments, while minimizing the risk of hypothetical biases.

The structure of the paper is as follows. Section 2 reviews the relevant literature. In Section 3, we describe the GWR methodology used for the approximation of door-to-door travel times. Section 4 provides an overview of the data used in the GWR model, and the results of the GWR models are then shown in Section 5. Section 6 discusses the set-up of the departure time choice models, the respective data, as well as alternative model specifications. Section 7 presents the results of the choice models, and Section 8 concludes.

2. Related literature

A limited number of studies exist that estimate departure time choice models using RP data, either as only data source, or in combination with SP data[4]. Most of the older RP-based studies on departure time choice behavior assume that travel times do not vary between days, ignoring travel time variability. For instance, Abkowitz (1981), Cosslett (1977) and Small (1987, 1982) use a dataset from the San Francisco Bay Area for the Urban Travel Demand Forecasting Project (UTDFP). By means of interpolation between peak and off-peak network values, driver-specific door-to-door travel times are computed for 12 alternative travel times, which do not vary across days. A similar strategy is employed by McCafferty and Hall (1982) and Bhat (1998a,b).

More recent studies on RP-based departure time choice models do account for travel time variability. While this renders the models more realistic, it also raises the requirements for travel time data. Either day-specific travel times data need to be gathered, or a good way of approximating them or the (time-of-day-specific) travel time distribution has to be found. As a consequence, only few of these studies take into account door-to-door travel times. One example is Börjesson (2008), who combines day-independent traffic simulation data and enriches them with camera data to generate travel times that vary by time of day. However, she assumes that travel times on secondary roads do not differ between days, arguing that these are hardly ever congested. She does not comment on the possibility of varying travel times due to supply-related influences such as weather conditions.

Lam and Small (2001) observe travel times directly along a 10 mile corridor for their joint route and departure time choice model, but not for the remaining links of the door-to-door trips. They estimate two models: The first model only takes into account travel times along this corridor, whereas the second model assumes that drivers face also on the rest of their travel the same speeds as observed along the corridor (except for a 5 mile access link, which is assumed to be uncongested).

[4]E.g. Ghosh (2001); Börjesson (2008); Small et al. (2005); Brownstone and Small (2005)

They find that the value of time decreases by 50% if they move from the first specification to the second specification. It can be expected that the second specification inflates the assumed travel time differences between travel moments, making the result in itself plausible. Besides the first model of Lam and Small (2001), also Knockaert et al. (2012) and Tseng et al. (2013) only account for travel times on the specific highway link along which the regarding experiment took place.

To our knowledge this study is the first to use GPS data for defining the attribute values of the choice alternatives in a scheduling model. Nevertheless, some authors have used GPS data to analyze other travel choices, in particular route choice behavior (e.g. Carrion and Levinson, 2013; Ramos et al., 2012; Bierlaire et al., 2010). Moreover, Fifer et al. (2011) use GPS data for defining realistic attribute values in their SP-based study on crash-risk reduction.

This paper fits well into the recently emerging stream of literature on the use of GPS data for travel time measurements. Independent of the context of choice modeling, many studies investigate the problems arising from a limited availability of GPS data, which can manifest itself in two different ways: First, the sampling frequency of GPS observations for a given trip may be low, such that the spacing between two subsequent GPS observations from the same device becomes rather wide. Possible causes are external conditions such as tall buildings or tunnels (Chen et al., 2005) as well as the technical specifications of the GPS device (Jenelius et al., 2013; Westgate et al., 2011). Naturally, the consequence of low-frequency sampling is that it becomes difficult to infer which route a driver has followed – an information that is often considered desirable or even necessary for the computation of GPS-based travel times. Various methods to deal with this so-called map matching problem have therefore been developed (e.g. Bierlaire et al., 2010; Quddus et al., 2007). This first kind of GPS data sparsity is not evident in this paper, as GPS observations are sampled at a fairly high frequency (almost every second).

The second type of GPS data sparsity is present when the number of vehicles equipped with GPS devices is relatively low compared to the overall network size. This type of sparsity clearly applies to the GPS data used in this study, since we only use GPS data from 230 drivers. The respective literature, however, suggests that GPS data from 4-6% of total vehicles on a given network are necessary for travel time measurements to be fairly precise (El Esawey and Sayed, 2012).[5] Different solutions have been presented in the literature to deal with the sparsity of GPS data due to limited coverage with GPS devices. One is to combine the GPS data with other travel time data sources such as loop detector data, using Kalman filtering (Chu et al., 2005) or Bayesian methods (Choi and Chung, 2002). Another one is to use available travel time information from nearby network links, taking into account that link travel times tend to be correlated. A distinction can then be made between models that emphasize temporal correlation patterns and those that focus on spatial correlation patterns. The ones belonging to the first category usually rely strongly on time-series data. Examples are models with Markov properties (e.g. Herring et al., 2010; Castro et al., 2012) or space–time autoregressive integrated moving average (STARIMA) models (e.g. Kamarianakis and Prastacos, 2005).[6] Models that emphasize the spatial correlation pattern often use geostatistical kriging models (El Esawey and Sayed, 2012; Zou et al., 2012; Miura, 2010; Aultman-Hall and Du, 2006) or local regression models (e.g Hackney et al., 2007; Idé and Kato, 2009).

In this paper, to deal with the very limited number of GPS-equipped vehicles, we employ both

[5]Wang et al. (2013) emphasize that this percentage depends on the acceptable accuracy of the measurements as well as the sampling frequency of the GPS devices.

[6]Note that these models are not *per se* specific to GPS data, but can also be used for other traffic data sources. For instance Kamarianakis and Prastacos (2005) apply their model to loop-detector data.

strategies presented above: the combination with other data sources as well as the estimation of models that account for travel correlations across links. We first combine GPS data with data from loop detectors. However, in contrast to the papers of Chu et al. (2005) and Choi and Chung (2002), we use GPS data for one part of the network and loop detector data on another part of the network, rather than combining both data sources on a given link. Second, we take into account travel time correlations across links. As the GPS observations are very sparse in our dataset, models requiring time series data are infeasible. Among those models that focus on the spatial correlation across links, we choose to use geographically weighted regression (GWR), which is a spatial version of local regression, and goes back to the seminal papers of Stone (1977) and Cleveland (1979).

GWR is very flexible in its use and specification, and often easier to calibrate compared to models that employ kriging (Harris et al., 2010). In defining the correlation neighborhoods, we choose to use Euclidean distances rather than network distances. Some previous studies (e.g. Hackney et al., 2007; Zou et al., 2012) have defined neighborhoods based on (undirected) network distances, and found that these perform better in predicting speeds compared to the Euclidean distances. However, using the more simple Euclidean neighborhood definition seems to be justifiable in our case, since we are not interested in trips throughout the entire network but only in those leading from the home locations to camera C1, and from camera C2 to the work locations. Due to the limited number of routes leading to these destinations (e.g. a limited number of highway entries and exits), two persons with home (work) locations close to each other in terms of Euclidean distance are likely to choose a similar route towards C1 (work) anyways, reducing the need to model network distance explicitly. For a similar reason, our method does not involve any map-matching on the home–C1 and C2–work links, meaning that we only consider speeds and distances rather than the actually chosen route along the GPS-observed links in our analysis. Note that therefore in this paper a 'link' refers to an abstract connection between a home location and C1, or between C2 and a work location, without pinpointing this link to a specific route. Our method is therefore generally well applicable in situations where GPS data are sparse due to low sampling frequency, and map-matching would be difficult.

The problems related to GPS data sparsity are expected to decrease over the next decade, as more and more people adopt GPS-enabled smartphones and navigation devices, and technical advancements allow for higher sampling frequencies. Also various issues related to privacy, data ownership, standardization and data processing methodology are likely to be solved in the coming years (e.g. Leduc, 2008; Krause et al., 2008). However, the methodology developed in this paper is expected to remain relevant, since it provides for a generic way for using GPS data to measure and predict travel times.

3. Geographically Weighted Regression

3.1. Methodology

The basic idea is to use geographically weighted regression (GWR) to explain speeds on those parts of the network for which infrequent GPS observations are available (from 230 drivers; along the home–C1 and C2–work links) by the speeds along the C1–C2 link where continuous speed measurements are available from loop detectors. The estimated parameters that describe the relationship between the GPS-measured speeds and the continuously measured speeds are then used to approximate time-of-day- and day-specific speeds between the home locations of the participants of the peak avoidance experiment and C1, and C2 and their work locations. If their home or work location was not in the original sample, based on which the GWR model was estimated, spatial interpolation is used to compute the relevant parameters necessary to approximate their door-to-door travel times.

GWR is a form of local regression that aims at the analysis of spatial data. Unlike the ordinary least squares (OLS) model, local regression models do not yield a global set of coefficients but local coefficients that result from fitting models to localized subsets of the data. In the GWR model, coefficients can therefore differ over space (in our application across home and work locations). The basic idea is that location specific coefficients tend to be more similar for locations situated close to each other, and to diverge more strongly for locations further apart. GWR has the advantageous property that all spatial information with respect to the start and ending points of the trips is fully used. No aggregation into areas (e.g. ZIP code areas) is required, and full advantage is taken of the spatial variation of the data.

The GWR model implies that for each observation a weighted least squares regression model is estimated, using a spatial weight matrix (e.g. Fan and Gijbels, 1996; Charlton and Fotheringham, 2009; Brunsdon et al., 1998). The weight matrix is determined by the relative geographic locations, implying that higher weights are attached to observations that are closer to the reference observation. The coefficients estimated from this model can then be used to predict speeds on the GPS observed links for different days and times of the day. By means of spatial interpolation, we can also utilize them to approximate speeds for (home and work) locations for which no GPS observations are available. The estimator of the locally weighted least squares model at the reference location u is thus given by

$$\hat{\lambda}(u) = [X^T W(u) X]^{-1} X^T W(u) y, \tag{1}$$

where $\hat{\lambda}(u)$ is the local parameter estimate, y is the N x 1 column vector including the values of the dependent variable (the speeds on the GPS observed links) for all observations $i = 1, \ldots, N$, and X is the N x p matrix of covariates (for p covariates: most importantly, the speeds on the continuously observed link). $W(u)$ denotes the spatial weight matrix. Its diagonal elements correspond to the weights between the reference location u and all other observations.

$$W(u) = \begin{pmatrix} w_1(u) & 0 & \cdots & 0 \\ 0 & w_2(u) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & w_n(u) \end{pmatrix} \tag{2}$$

The weights at location u with respect to observation i, $w_i(u)$, are determined by the kernel function f, which takes the Euclidean distance $d_i(u)$ between the locations associated with u and i and the bandwidth parameter h, as inputs. The kernel function yields a weight that is decreasing in the distance between u and i:

$$w_i(u) = f(d_i(u), h) \tag{3}$$

The functional form of $f(d_i(u), h)$ is assumed to be Gaussian[7].

$$w_i(u) = \exp\left[-\frac{1}{2}\left(\frac{d_i(u)}{h}\right)^2\right]. \tag{4}$$

[7]The scale of the weights does not matter as $[X^T(\alpha W(u))X]^{-1}X^T(\alpha W(u))y = 1/\alpha[X^T W(u)X]^{-1}X^T \alpha W(u)y = [X^T W(u)X]^{-1}X^T W(u)y$.

The bandwidth h is a fixed, unknown parameter. It determines the distance decay of the weights and is expressed in the same units as the distances $d_i(u)$. As the bandwidth increases, the weights become uniform and the local GWR model approaches the global OLS model. The optimal bandwidth trades off model fit (resulting from large bandwidths) and variance (resulting from small bandwidths). We use cross-validation (Bowman et al., 1998; Wheeler and Páez, 2010) to determine the optimal bandwidth. Cross validation defines that bandwidth as optimal that minimizes the root mean squared prediction errors (RMSPE), using a subset of the data for prediction. In the standard application of this method, the reference observation i is left out from the estimation in order to prevent a perfect fit of the model. However, in our application we leave out all observations that are attributed to the same driver z as observation i can be attributed to:

$$\hat{h} = \arg\min_{h} \sqrt{\sum_{i=1}^{n} [y_i - \hat{y}_{-z(i)}(h)]^2} \tag{5}$$

This 'leave-one-driver-out' cross-validation criterion is adopted to account for the panel structure of the dataset. Observations attributable to the same driver tend to be clustered spatially as well as temporally (i.e. a driver often has similar departure times across working days). If only the reference observation i is left out, the optimal bandwidth is likely to be too small, yielding very precise predictions for all trips undertaken by the corresponding driver $z(i)$, however, not for other drivers with similar start or end locations, who might depart at different times of the day. Since we intend to use the coefficients for out-of-sample predictions, it is crucial that the model is able to predict travel times well for all times of the day. We do not account for the panel setup directly in the estimation, for instance, by using fixed effects, since start and end points of the door-to-door trips for a given driver can differ between days.[8] Introducing person-specific effects is therefore not useful as they might capture location-specific rather than driver-specific effects. Furthermore, we expect speed differences between individual drivers to be negligible, as speed-limits are enforced rather strictly and many links are congested during most morning peaks, leaving little room for speed differences between drivers.

3.2. Functional form of the local specification

In order to render the estimation results easily applicable for out-of-sample prediction of speeds for home and work locations that are not in the estimation sample, two models are estimated: One that measures the correlation between speeds on the home–C1 links and the C1–C2 link, and a second that measures the correlation between speeds on the C1–C2 link and the C2–work links.

We first need to establish the relationship between the time when the GPS-observed trip had started (at home for home–C1 trips; at C2 for C2–work trips), $\tilde{t}$, and the passage time at C1 that defines the C1–C2 speed used as explanatory variable in the GWR model, $\bar{t}$. The corresponding speeds on the home–C1 (C2–work) link for a specific home (work) location u are then denoted by $v_{\tilde{t}}^{gps}(u)$, and the speeds on the C1–C2 link by $v_{\bar{t}}^{cont}$.

In the home–C1 case, for an observed departure time from home, $\tilde{t}$, we use the C1–C2 speed also at time $\tilde{t}$ as explanatory variable ($\bar{t} := \tilde{t}$). If we defined the C1–C2 speed upon arrival at C1, recursiveness would result, since then the passage time at C1 would be determined by the

[8]Such differences can result for multiple reasons: Variations in the parking location of the car, switching on (or off) the GPS device during the trip, or trip and GPS signal interruptions. See also Footnote 2.

home–C1 speeds but at the same time the home–C1 speeds would determine the passage time at C1. However, using the definition of $\bar{t} := \tilde{t}$, speeds on the home–C1 link determine the passage time at C1 but not the other way round. In the C2–work case, recursiveness does not occur since the C1–C2 link is passed before the GPS-observed link. We therefore define $\bar{t}$ as the passage time at C1 that results in the observed passage time at C2, $\tilde{t}$, given C1–C2 speeds[9], v_t^{cont}:

$$\begin{aligned} &\bar{t} := \tilde{t} && \text{for home–C1 links} \\ &\bar{t} := \arg\min_t |\tilde{t} - \left(t + \frac{\text{distance}^{C1-C2}}{v_t^{cont}} \right)| && \text{for C2–work links} \end{aligned} \tag{6}$$

The (location-specific) model that relates the GPS observed speeds to the C1–C2 speeds assumes a simple, linear form:

$$v_{\tilde{t}}^{gps}(u) = \lambda_0(u) + \lambda_1(u) v_{\tilde{t}}^{cont} \tag{7}$$

Eq. 7 establishes a simple direct relation between the speeds on the two links. It proved to perform better in out-of-sample prediction than more complicated model structures that we tested, including models that use the departure time as explanatory variable both in linear [10] as well as nonlinear[11] specifications. While the overall fit of these models is comparable to the fit of the model based on Eq. 7, the more complicated functions tend to yield large outliers at the begin and at the end of the peak.

Furthermore, we also investigated a model where the weight matrix was not only location- but also time-specific. While this method leads to small improvements in the predictive power of the models[12], it has the disadvantage that a different set of coefficients is valid for different times of the day, making the model less general. Eq. 7, on the other hand, is applicable to all times of the day. Moreover, it does not require any assumption on the shape of the peak. Note that weekday- and weather-dependent influences on door-to-door travel times, which tend to affect the entire network, are captured through by the observed speeds on the continuously observed link and the correlations of the home–C1 and C2–work speeds therewith.

4. Data

4.1. Experiment

The data used in this study were gathered from a peak avoidance experiment taking place in the Netherlands. In Dutch it is referred to as *Spitsmijden* experiment. Participants could receive a daily reward of 4 Euro for avoiding a specific highway link of 9.21 km length (C1–C2) during morning peak hours (6:30–9:30).[13] The experiment was set up with the goal to mitigate the negative effects of roadworks undertaken along this link during the time of the experiment.

[9] The distance between C1 and C2 is equal to 9.21 km.

[10] E.g. $v_{\tilde{t}}^{gps}(u) = \lambda_0(u) + \lambda_1(u) v_{\tilde{t}}^{cont} + \beta_3(u)\bar{t} + \beta_4(u)\bar{t}^2$

[11] E.g. $v_{\tilde{t}}^{gps}(u) = \lambda_1(u) - \beta_2(u)(1 + \beta_3(u) v_{\tilde{t}}^{cont}) e^{-(\beta_4(u)\bar{t} - \beta_5)^2} + \beta_6 v_{\tilde{t}}^{cont}$

[12] The root mean squared error (RMSE) decreases on average by 0.45 km/h for the home–C1 model, and by 1.18 km/h for the C2–work model.

[13] Drivers participating in the experiment could not earn a reward on weekends or school vacation days, but also not if they had already exceeded the maximum number of days per week for which a reward could be obtained. This maximum is driver-specific, and equal to the average number of weekly trips undertaken during the pre-experimental period (reference behavior).

Car drivers who were observed to pass the C1–C2 link multiple times a week were invited to participate in the experiment. In the first phase (11/2008–4/2009), 230 drivers obtained customized smartphones, also named Rabomobiel. These phones were equipped with a GPS receiver and transmitted information about the location of the phone to a central database. The GPS data gathered during this first phase are then used to approximate door-to-door travel times for the second phase of the experiment (09/2009-12/2009), where no GPS data were collected any longer and therefore travel time measurements were only available for the C1–C2 link. The number of participants in the second phase was close to 2000, of which 428 are considered in this paper. These are the ones for whom the preferred arrival time at work, which is required for the estimation of the scheduling model, is known from a questionnaire. An in-depth description of the Spitsmijden experiment can be found in Knockaert et al. (2012).

4.2. GPS-based speed measurements

In their raw form, GPS measurements are data points with a location, speed and time stamp attached to them. From sequential data points from a given device, speeds and distances can be computed.[14] We take into account GPS-observed trips that pass either through camera location C1, C2 or both. We construct separate datasets for trips observed between home and C1, and C2 and work. A trip is defined to start when the speed at which the GPS receiver moves exceeds 5 km/h, and to end when the speed drops to less than 5 km/h for at least 10 minutes. To determine the average speed, the distance covered between the start and the end time needs to be known. We define the distance as the sum of the distances between the GPS observations belonging to the same trip.[15] We only consider trips along the home-C1 and C2–work links that were observed during the morning period (between 5:00 and 11:00 a.m.) and on working days, corresponding to our analysis of departure time choices in the morning and on weekdays. We exclude trips that result in average speeds lower than 30 km/h or above 120 km/h, and those for which no close neighboring observations – with similar start or end locations[16] – are available. Also, for a given trip the largest distance between 2 GPS measurements must not exceed 2 km.

Table 1 gives an overview of the descriptive statistics of the data for the home–C1 and C2–work links. Trips along the home–C1 links tend to be longer and faster (in terms of speed) compared to trips along the C2–work links. This finding can be attributed to the fact that, on average, a larger share of the home–C1 links consists of highways rather than local roads. Due to the higher average speeds on the home–C1 link, the mean and maximum distance between two subsequent GPS measurements are on average higher on the home–C1 link. Also, average speeds on the C1–C2 link (after passing the home–C1 or before passing the C2–work link) are shown. They are similar to the speeds on the GPS-observed links but significantly lower than the free-flow speed of 100 km/h (i.e. the speed-limit along the C1–C2 link). In contrast to the standard deviation of speeds

[14] Also the chosen routes can be inferred from these data. However, as previously stated, we do not make use of any map-matching procedures in this paper.

[15] Since location measurements using GPS can be somewhat imprecise, it is possible that the GPS device transmits different locations although the vehicle is standing still. If these distances are added, the overall trip distance might be overestimated, resulting in a structural downward bias of travel times (e.g. Zito et al., 1995). For this reason, if speed drops below 5 km/h, we only take into account the distance between the location at which speed dropped below 5 km/h and the location at which the vehicle resumed a speed of above 5 km/h. It is reassuring that this approach yields very similar results compared to defining distance based on a 'fastest network distance' (retrieved from `openrouteservice.org`).

[16] Less than 10 observations within a perimeter of 5 km of the reference location.

for the GPS-observed links, which is also a consequence of heterogeneity between driver-specific links and the corresponding free-flow speeds, the standard deviation of speeds for the C1–C2 link can almost exclusively be attributed to congestion (between-driver-variability is expected to play only a minor role).

Moreover, Table 1 shows that departure (arrival) times as well as home (work) locations differ much more strongly across drivers than for a given driver (across trips), providing a good argument for the use of the 'leave-one-driver-out' cross validation criterion for bandwidth selection, as discussed in Section 3.1. Moreover, Table 1 demonstrates the importance of using door-to-door travel times. While the length of the C1–C2 link is 9.21 km, the average length of the home–C1 and C2–work links are equal to 30.91 km and 16.34 km, respectively. Hence, the omission of these parts of the trip that cannot be observed by means of continuous speed measurements would result in the omission of substantial parts of the door-to-door trips.

Table 1: Descriptives GPS data

Characteristics	Home–C1	C2–Work
Nr. of trips	2896	5896
Nr. of drivers	91	163
Nr. of days	144	144
Mean travel time (per trip)	23:08 min	15:08 min
Std. dev.	14:11 min	5:50 min
Mean distance (per trip)	30.91 km	16.34 km
Std. dev.	19.36 km	3.80 km
Mean speed (per trip)	78.69 km/h	69.29 km/h
Std. dev.	21.88 km/h	16.65 km/h
Mean speed on C1–C2 (per trip)	80.02 km/h	74.46 km/h
Std. dev.	20.30 km/h	19.21 km/h
Mean number of GPS measurements (per trip)	1288	784
Std. dev.	775.26	252.64
Mean time btw. 2 GPS measurements	1.09 s	1.15 s
Std. dev.	0.14 s	0.15 s
Mean distance btw. 2 GPS measurements	23.36 m	21.61 m
Std. dev.	5.94 m	4.00 m
Maximum distance btw. 2 GPS measurements (on a given trip)	85.14 m	115.43 m
Std. dev.	119.69 m	198.37 m
Overall std. dev. dep. (arr.) time at home/work	90:36 min	87:36 min
Median driver-specific std. dev. dep. (arr.) time at home/work	32:12 min	34:02 min
Overall median distance btw. home (work) locations	16.37 km	4.24 km
Median distance btw. driver-specific home (work) locations	1.43 km	0.35 km

4.3. Loop-detector based speed measurements

The GPS observed speeds discussed in the previous section are matched to C1–C2 speeds according to the temporal relations between the starting time of the home–C1 (C2–work) trip and the passage times on the C1–C2 link established in Eq. 6. The speeds on the C1–C2 link are available for any time of the day from a dense network of loop detectors. These speeds are aggregated in time (15-minute intervals) and space (from the single detectors towards the entire link), employing the trajectory method. A general description of this method can for instance be found in Van Lint (2010), whereas the specific method applied on the data used in this study is described in Modelit (2009). It can be shown that the speeds on the C1–C2 link do not differ significantly between the two project phases (11/2008–4/2009 and 09/2009-12/2009), providing confidence that we can safely use the relation derived between the GPS-based speeds and the camera observations during the first phase to make travel time predictions during the second phase. Moreover, we could verify that the speeds derived from the loop detectors are not significantly different from the GPS-measured speeds along the C1–C2 link.

5. Results: GWR model

In this section we present the results obtained from the GWR model. We first determine the optimal bandwidths using the 'leave-one driver-out' cross-validation approach. Figure 3 shows that a minimum exists with respect to the RMSPE for both the home–C1 and the C2–work link. The optimal bandwidths are determined accordingly: 2.64 km for the home–C1 model and 1.05 km for C2–work model. Figure 4 then displays the corresponding kernel weights attached to the

Figure 3: Optimal bandwidth choice

(a) Home–C1 (b) C2–Work

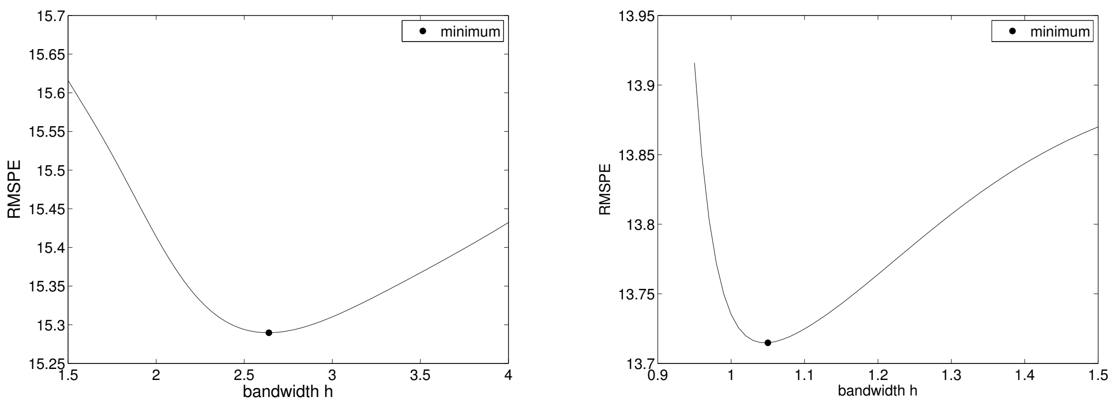

distance between reference location u and another observation i: The weights are subject to a steeper distance decay for the case of the C2–work links than for the home–C1 links. This is likely be the joint result of a low variation in the local traffic conditions along the C2–work links, the higher number of C2–work observations and the more densely located work locations (compared to the home locations; see Figure 2). A comparison of the optimal GWR model to an OLS model with generic parameter values reveals that indeed the root mean squared error (RMSE) decreases, which of course is expected since OLS is a restricted case of GWR. Also the R-squared increases substantially if one moves from OLS to GWR (Table 2). These results are more pronounced for the home–C1 links, confirming that spatial heterogeneity plays a larger role for the home–C1 link.

Figure 5a shows a map of the area within which the home and work locations are situated. Figures 5b and 5c show the corresponding spatial pattern that results if speeds are predicted using the location-specific coefficients derived from the GWR model, making different assumptions on

Figure 4: Kernel weights

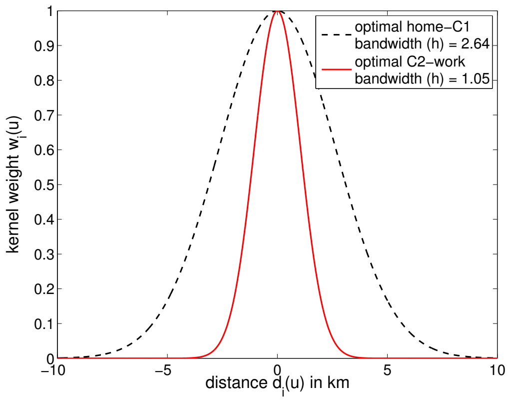

Table 2: Results OLS and GWR model

	OLS		GWR	
	Home–C1	C2–Work	Home–C1	C2–Work
λ_0	54.62	47.76	–	–
λ_1	0.29	0.29	–	–
Bandwidth	–	–	2.64	1.05
mean(λ_0)	–	–	52.91	50.11
st.dev.(λ_0)	–	–	16.22	10.05
mean(λ_1)	–	–	0.31	0.26
st.dev. (λ_1)	–	–	0.10	0.08
N	2896	5896	2896	5896
RMSPE	21.49	15.81	15.29	13.72
RMSE	21.06	15.69	13.62	12.53
R-squared	0.07	0.11	0.61	0.43

the C1–C2 speed: In Figure 5b, a C1–C2 speed of 50 km/h is assumed ('peak scenario'), whereas in Figure 5c a C1–C2 speed of 100 km/h ('off-peak scenario') is assumed. Note that the speeds are given in terms of departure location for the home–C1 links and in terms of arrival location for the C2–work links. As expected, higher speeds are predicted for observations starting/ending along the highway, while lower speeds are predicted for observations starting or ending along local roads.

6. Departure time choice models

As a next step, we estimate departure time choice models. We use the location-specific coefficients derived from the GWR model to approximate door-to-door travel times, and based on these the values of the other attributes of the utility function. Moreover, we show to which extent the estimated coefficients are different if less sophisticated estimates of travel times are used.

Figure 5: Speed predictions

(a) Map

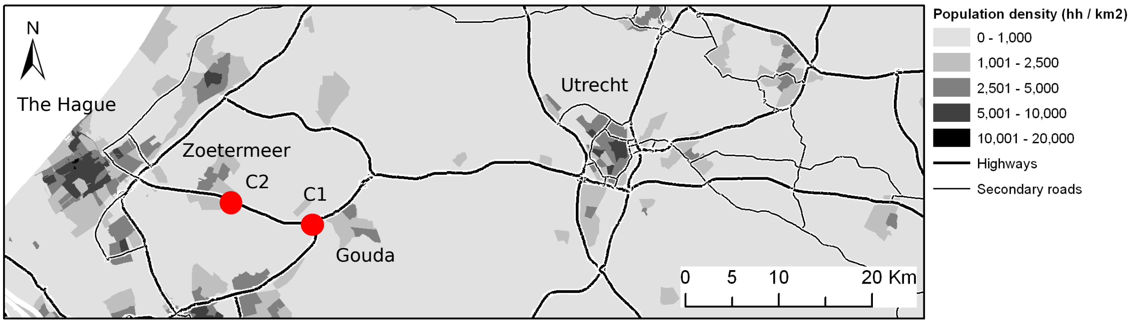

(b) Predictions: C1–C2 speed = 50 km/h

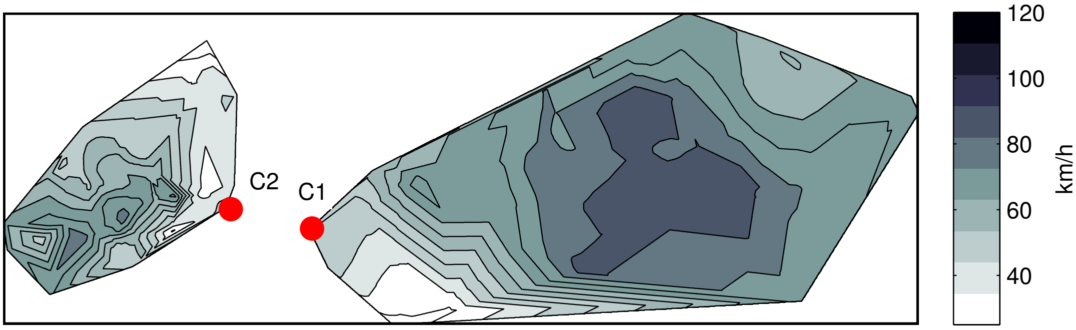

(c) Predictions: C1–C2 speed = 100 km/h

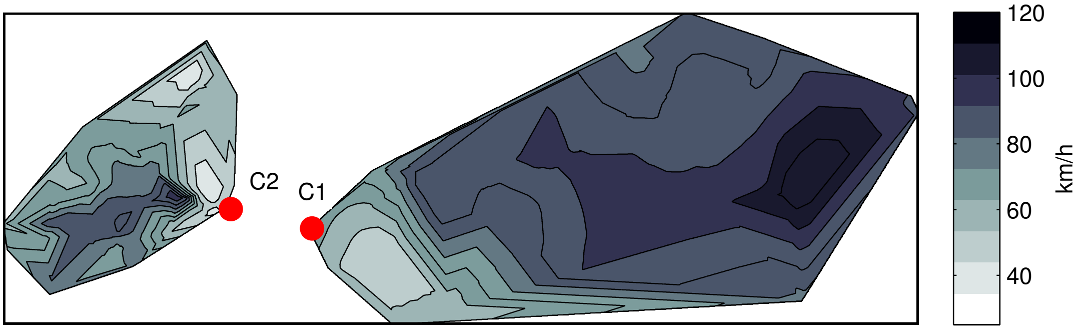

6.1. Utility function

The continuous departure time choice problem is re-formulated as a discrete problem with a finite number of departure time intervals. We use a standard multinomial logit (MNL) model for the estimation. A driver z faces $k = 1, \ldots, K$ choices among the alternatives $j = 1, \ldots, J$, where $J = 17$, since departure time intervals of 15 minutes between 5:30 and 9:45 a.m. are considered. K is equal to the duration of the experiment of 75 working days, however, no driver was observed to pass the C1–C2 link on all 75 days. Since we are interested in departure time choices rather than the choice on whether to travel at all on a given day, we only use in the estimations the driver-specific subset of days during which a driver has been observed to travel. On a day k driver z then chooses the alternative j that maximizes the following random utility function:

$$U_{zkj} = V_{zkj} + \epsilon_{zkj}, \tag{8}$$

The random utility function consists of a deterministic component V_{zkj} and a random component ϵ_{zkj} that follows a Gumbel distribution, with errors distributed identically and independently (iid) across observations. To account for a bias in the standard errors as a result of the panel nature of the data, the panel sandwich estimator is used (e.g Daly and Hess, 2011). The choice probability of alternative $\breve{j} \in J$ is then given by:

$$P_{zk\breve{j}} = \frac{\exp(V_{zk\breve{j}})}{\sum_{j=1}^{J} \exp(V_{zkj})}. \tag{9}$$

The formulation of the deterministic utility component builds on the scheduling model of Vickrey (1969) and Small (1982). It accounts for the trade-off between travel times and schedule delays. Since we take into consideration that travel times may vary across days, the attributes of the utility function are the expected measures of reward, ER_{zkj}, travel time, ET_{zkj}, schedule delay early, $ESDE_{zkj}$, and schedule delay late, $ESDL_{zkj}$. The corresponding coefficients are denoted by $\beta_R, \beta_T, \beta_E$ and β_L.

$$V_{zkj} = \beta_R * ER_{zkj} + \beta_T * ET_{zkj} + \beta_E * ESDE_{zkj} + \beta_L * ESDL_{zkj}, \tag{10}$$

We also tested more advanced model specifications that take into account travel time variability *per se* in addition to the travel time variations implied by the scheduling terms. However, when travel time variability is expressed in terms of the travel time variance or in terms of percentile differences, it turns out insignificant. If expressed in terms of the standard deviation, it assumes a very high coefficient in some models while the travel time coefficient turns insignificant. This reflects that travel times and the standard deviation are strongly correlated, and can therefore not be estimated separately. For these reasons, we leave out the variability indicator in the estimations, and keep in mind that the time coefficient β_T might also pick up some disutility attributable to travel time variability *per se*.

The values of time (VOT) and schedule delay early (VSDE) and late (VSDL) are then defined as the ratios between the coefficients of time and schedule delays and the reward coefficient, multiplied by (-1) to indicate that the values represent the willingness-to-pay for *reductions* in travel time and schedule delays.

$$\text{VOT} = -\beta_T/\beta_R \qquad \text{VSDE} = -\beta_E/\beta_R \qquad \text{VSDL} = -\beta_L/\beta_R \tag{11}$$

When determining the expected values of the attributes $A = \{\text{R, T, SDE, SDL}\}$, we take into account that travelers usually have more knowledge of travel times than their averages, even if time-of-day specific. We model this by assuming that the expected values of the attributes are based on a compound measure of average travel time across the entire experiment, thus across days $k = 1, \ldots, K$[17] and the realized travel time on the day of travel, $\breve{k}$, which we refer to as 'current' travel times. Even though current travel times are typically unknown in advance (hence, at the time the departure time decision is taken), they represent an upper threshold for the maximum possible extent of information available to drivers. The average travel times, on the contrary, represent the long-run pattern of travel times over the time of the day, which participants of the experiment are likely to be aware of. We denote the (estimated) relative weight attached to the current travel times by θ. The expectation for the attribute A of the utility function is then given by:

$$EA_{z\breve{k}j} = \theta * A_{z\breve{k}j} + (1 - \theta) * \frac{1}{K} \sum_{k=1}^{K} A_{zkj} \tag{12}$$

If a driver z is eligible for a reward on day k, the reward linked to departure time choice j, R_{zkj}, is equal to 4 Euro if it results in a passage time at C2 before 6:30 or a passage time at C1 after 9:30, and it is equal to 0 in all other cases. Schedule delays are a function of the difference between the arrival time, which depends on departure time t_{zkj} and travel time T_{zkj}, and the preferred arrival time PAT_z, which is defined as the self-reported (in a questionnaire) preferred arrival time in case no congestion would ever occur.

$$\begin{aligned} SDE_{zkj} &= \max(0, \text{PAT}_z - t_{zkj} - T_{zkj}) \\ SDL_{zkj} &= \max(0, t_{zkj} + T_{zkj} - \text{PAT}_z) \end{aligned} \tag{13}$$

It is likely that travel times are to some extent endogenous, since they are determined by the choice of the residential location (whereas the choice of the work location is arguably exogenous, in particular because we only focus on work locations in and close by The Hague). However, this does not automatically translate into an endogeneity bias in the parameter estimates, as the estimates are determined by attribute differences rather than absolute attribute values. So, for an endogeneity bias to become evident, drivers would have to choose their residential location taking into account the actual recurrent congestion pattern over the time of the day. But this seems to be rather unlikely, as residential location choices are long-term, infrequent decision processes, whereas the recurrent congestion pattern can change quite easily over time, for instance as a consequence of traffic management measures, or road capacity expansion. And even if recurrent congestion would be a determinant of the residential location choice, it remains true that the driver-specific difference between minimum and maximum expected travel times is fairly similar across the drivers considered in this paper, since they all pass the C1–C2 link (which is the major source of congestion in this area) on their way to work. Hence, we speculate that a possible endogeneity bias is not a major concern here.

6.2. Use of the GWR coefficients

We use the GWR coefficients from Table 2 to approximate door-to-door travel times for both chosen and unchosen departure time intervals, as well as for the approximation of the departure

[17] A similar definition of expected travel times was used by Tseng et al. (2013), however, they do not apply the compound measure to all attributes of the utility function but to travel times only.

time from home. For all drivers, home and work locations are known, as well as the shortest network distances[18] between their home locations and C1, and between C2 and their work locations. It is not necessarily the case that these coincide with the start and end locations of the links that were considered in the original GWR models. In order to derive the location-specific coefficients for a specific location $\bar{u}$ ($\hat{\lambda}_0(\bar{u}), \hat{\lambda}_1(\bar{u})$), we therefore use spatial interpolation. Specifically, we apply the Gaussian[19] kernel function with the optimal bandwidths (see Table 2) to the distance between $\bar{u}$ and all locations $u(i)$, which are associated with the $i = 1, \ldots, N$ observations considered in the original GWR model. The GWR coefficients specific to location $\bar{u}$ are then equal to the original GWR coefficients weighted by the kernel weights with respect to $\bar{u}$.

$$\hat{\lambda}_p(\bar{u}) = \frac{1}{\sum_{i=1}^{N} w_i(\bar{u})} \sum_{i=1}^{N} w_i(\bar{u}) \hat{\lambda}_p(u(i)) \qquad \text{for } p = 0, 1 \tag{14}$$

We can use these coefficients and the shortest network distances to approximate driver-, day-, and time-of-day-specific door-to-door travel times, T_{zkj}, taking into account the relationship between $\bar{t}$ and $\tilde{t}$ as described in Eq. 6. Based on the resulting travel times, reward and schedule delay attributes are calculated. Finally, the departure time from home is defined as the departure time choice alternative j that results in a passage time at C1 closest to the observed one. In order to warrant the quality of the data, we do not consider those drivers in the departure time choice models who have a home (work) location that is more than 5 km away from the closest home (work) location considered in the original GWR estimation.

6.3. Data

Table 3 provides some descriptive statistics of the participants considered in the departure time choice models. It shows that under uncongested conditions a large majority of drivers prefer to arrive at work between 7:00 and 9:00 a.m. The incomes of the drivers considered in the analysis are rather high, compared to the Dutch average. The average length of the home–C1 link and C2–work link is smaller than the corresponding distances considered in the original GWR estimation (see Table 3). This might be due to the fact that we do not include those drivers who have a home or work location that is further than 5 km away from the closest home (work) location considered in the original GWR estimation, which is more likely to be the case for locations that are relatively far away from the C1–C2 link. As a result of this exclusion also average distances to the 'closest GWR neighbor' are small: 0.74 km for home locations and 0.23 km for work locations. The average values of the $\hat{\lambda}_p$ and their standard deviations correspond closely to those found in the original GWR estimation (see Table 2).

6.4. Model specifications

We define three models that differ in the computation of the speeds on the GPS-observed links, and that represent three ways in which a researcher might tackle the difficulties of observing only part of a trip. They have in common that speeds along the C1–C2 link are based on observed speeds.

[18]Retrieved from `openrouteservice.org`

[19]We tested alternative weighting functions, such as using the coefficients of the closest neighbor or the closest 5 neighbors, or applying a uniform weighting function. However, the results of the departure time choice models are barely affected by the specification of this weighting function.

Table 3: Descriptives of drivers considered in the choice model

Variables	Share or sample average	Std. dev.
Commuter characteristics		
Preferred arrival time at work		
<7:00 a.m	0.11	–
>9:00 a.m.	0.07	–
Monthly (net) income		
$<$ 3500 Euro	0.34	–
$>$ 5000 Euro	0.13	–
unknown	0.24	–
Commute-related variables		
Distance Home–C1 (in km)	22.61	11.26
Distance C2–Work (in km)	15.90	4.82
Distance to 'closest GWR neighbor': Home–C1 (in km)	0.74	0.95
Distance to 'closest GWR neighbor': C2–Work (in km)	0.23	0.40
$\hat{\lambda}_0^{Home-C1}$	49.68	16.06
$\hat{\lambda}_0^{C2-Work}$	46.00	10.54
$\hat{\lambda}_1^{Home-C1}$	0.32	0.10
$\hat{\lambda}_1^{C2-Work}$	0.25	0.08
Choices		
Nr. of choices per individual	21.05	12.93
Nr. of individuals	428	–
Total nr. RP choices	9010	–
Duration of the RP experiment (working days)	75	–

Model 1 represents the model with the most sophisticated definition of door-to-door travel times, given the available data sources. Speeds on the home–C1 and C2–work links are then computed from the GWR coefficients, $\hat{\lambda}_0$ and $\hat{\lambda}_1$ (as derived in Eq. 14), and vary across space, days and time of the day. Model 2 represents the case when speeds outside C1–C2 differ between days and time of the day, however, not across space. Home–C1 and C2–work speeds are then independent from the home and work location. This setup is achieved by using the OLS coefficients λ_0 and λ_1 (see Table 2) rather than the GWR coefficients to determine home–C1 and C2–work travel times. Finally, Model 3 represents the case where travel times do not vary across space, days and time of the day. Speeds on all home–C1 and C2–work links are fixed to 70 km/h (i.e. $\lambda_0 = 70$, $\lambda_1 = 0$). Reasons for fixing the speed to exactly 70 km/h will be given below. While we do not explicitly consider a model where only C1–C2 travel times are accounted for (such a model would require a different specification of the departure time alternatives (at C1) as well as of the preferred arrival time (at C2)), Model 3 captures the main property of a model that is based on the C1–C2 link only, namely the effect of ignoring travel time correlations between the GPS-observed links and the C1–C2 link.

Table 4 provides an overview of the travel time attributes (for the 'current' travel times) associated with the 3 models. For the door-to-door travel times as well as the travel times along the sub-links, average speeds, average travel times and the average standard deviation of travel times (across departure time alternatives for a given choice situation) are reported. We find that the speeds and travel times are quite comparable across the models, which of course is desirable. If they were not comparable, the coefficients in the choice models might differ simply because travel times differ structurally across models (i.e. they are substantially longer in one model than in an other) rather than because of the different travel time specifications inherent to Models 1–3.

As expected, travel time variability is highest for Model 1 (GWR coefficients) and lowest for Model 3 (no variability outside C1–C2). Table 4 also substantiates our choice of fixing the speed to 70 km/h in Model 3: Speeds in Model 3 are close to those in Models 1 and 3, which are computed using GWR and OLS, respectively. While mean speeds are generally comparable across links (for all 3 models), the corresponding standard deviations differ substantially across links in Models 1 and 2. Not surprisingly, the highest standard deviation is evident along the C1–C2 link, which is an indication that there also most congestion occurs. So, the low travel time variabilities on the home–C1 and C2–work links reflect that on these links free-flow speed is close to the mean speed of ca. 70 km/h. The free-flow speed on the C1–C2 link, however, is close to 100 km/h (the speed-limit) and the expected speed can be shown to drop to ca. 60 km/h during peak hours, indicating a severe reduction in travel speeds. Finally, note that speeds along the C1–C2 link differ slightly across models. The reason is that for a given departure time alternative, the speed assumed on the home–C1 link determines the passage time of the C1–C2 link and as a consequence the speed along the same link.

Besides these three models, that differ with respect to how home–C1 and C2–work travel times are defined, we distinguish between two scenarios that differ in how departure times from home are defined, representing different amounts of information that a researcher might have. In the first scenario, the GWR coefficients are used to compute departure times in all models. We interpret these departure times as 'known' (observed) departure times, representing for instance the case when drivers are asked to fill in their departure times in a log-book. Since departure time choices are defined equally for all models in this first scenario, this also allows us to compare the models directly in terms of their fit. In the second scenario, departure times from home are calculated using

Table 4: Descriptives of the travel time attributes in the choice models

	Model 1 *GWR*	Model 2 *OLS*	Model 3 *constant*
Home–Work			
Mean speeds (in km/h)	70.16	70.58	67.77
Mean travel time (in min)	36.39	36.63	37.47
Std. dev. (in min)	6.25	5.95	3.99
Home–C1			
Mean speeds (in km/h)	71.05	75.48	70.00
Mean travel time (in min)	16.51	16.97	18.10
Std. dev. (in min)	1.82	1.47	0.00
C1–C2			
Mean speeds (in km/h)	71.19	70.91	70.27
Mean travel time (in min)	9.64	9.69	9.80
Std. dev. (in min)	3.97	3.99	4.00
C2–Work			
Mean speeds (in km/h)	64.37	68.26	70.00
Mean travel time (in min)	10.32	10.08	9.70
Std. dev. (in min)	0.90	0.93	0.00

the coefficients corresponding to the respective model. For instance, in Model 2, both (home–C1, C2–work) travel times as well as departure times are computed using the OLS coefficients. As a consequence, for a given driver and day, the departure time from home may differ across the models. We refer to this scenario as 'unknown departure time', as it is a proxy for the situation where departure times from home are not observed, and therefore need to be computed by the researcher. Clearly, for Model 1 the 'known' and the 'unknown' departure time scenario coincide.

7. Results: Departure time choice models

Table 5 shows the results of the GWR-based Model 1, as well as the results of Models 2 and 3 for the two scenarios defined in the previous section. We find all coefficients to be strongly significant. As mentioned above, it is not useful to compare the fit of the models in the scenario where departure times are 'unknown' to the ones where they are 'known'. However, we can compare the model fits for the scenario with 'known' departure times, where all models share the same dependent variable. We then obtain the reassuring result that we Model 1 yields the highest log-likelihood, followed by Model 2 and Model 3.

We obtain fairly high VOT estimates in all models (between 32 and 70 Euro/h). These can partly be explained by the rather high incomes of the participants (see Table 3), which tend to be correlated positively to the VOT (e.g. Small et al., 2005). Moreover, similarly high values have been obtained from an SP experiment conducted among the same set of drivers (Peer et al., 2013). In earlier estimations it has also be shown that some drivers seem to ignore the reward in their departure time decisions, meaning that their reward coefficient is close to 0. Since we do not account here for preference heterogeneity, the average reward coefficient is probably biased towards 0, causing the monetary valuations to be higher. Finally, the travel time coefficient might

Table 5: Estimation results: departure time choice models

		Departure time			
		'Known'		'Unknown'	
	Model 1	Model 2	Model 3	Model 2	Model 3
	GWR	*OLS*	*constant*	*OLS*	*constant*
β_R	0.22	0.21	0.19	0.17	0.19
t-stat.	(6.21)	(6.28)	(5.83)	(4.34)	(4.97)
β_T	-6.88	-7.31	-12.37	-8.06	-12.87
t-stat.	(-6.42)	(-5.56)	(-6.92)	(-6.70)	(-6.44)
β_E	-1.97	-1.96	-1.94	-2.11	-2.05
t-stat.	(-15.82)	(-16.01)	(-10.04)	(-16.40)	(-15.93)
β_L	-1.59	-1.60	-1.59	-1.56	-1.58
t-stat.	(-17.71)	(-17.80)	(-15.22)	(-17.61)	(-18.06)
θ	0.14	0.13	0.10	0.13	0.11
t-stat.	(5.23)	(5.09)	(4.60)	(5.26)	(5.17)
VOT (Euro/h)	31.85	34.32	65.61	47.69	69.73
t-stat. (different from 0)	(4.36)	(4.40)	(4.31)	(3.56)	(3.83)
t-stat. (different from Model 1)	–	(0.23)	(2.00)	(1.04)	(1.93)
VSDE (Euro/h)	9.12	9.20	10.26	12.49	11.08
t-stat. (different from 0)	(5.07)	(4.94)	(4.21)	(3.77)	(3.88)
t-stat. (different from Model 1)	–	(0.03)	(0.34)	(0.89)	(0.58)
VSDL (Euro/h)	7.36	7.51	8.47	9.23	8.54
t-stat. (different from 0)	(4.98)	(4.82)	(4.06)	(3.66)	(3.72)
t-stat. (different from Model 1)	–	(0.07)	(0.43)	(0.63)	(0.43)
Nr. Obs.	9010	9010	9010	9010	9010
LogLikelihood	-21526	-21599	-21678	-21441	-21441
Pseudo R-squared	0.157	0.154	0.151	0.155	0.154

also pick up disutility from travel time variability, which we were not able to estimate separately due to correlation between travel time and its variability.

When comparing the VOTs between Models 1 and 2, we find that they are very close to each other for the 'known'-departure-time scenario. Nevertheless, the loglikelihood of Model 2 ('known') is worse than that of Model 1. We find a slightly smaller reward coefficient, and somewhat higher (absolute) time coefficient in Model 2 ('known'). This finding is even more pronounced for the 'unknown'-departure-time scenario. While we had expected a downward bias on all coefficients due to measurement errors (Bhatta and Larsen, 2011; Train, 2003) as the OLS coefficients predict travel times less precisely than the GWR coefficients (see Table 2), we only find a lower reward coefficient, which might indeed be an indication of a measurement error. A possible explanation for the higher time coefficient (in absolute terms) is related to the setup of the peak avoidance experiment: Since travel times tend to be (negatively) correlated to rewards (due to off-peak rewarding), the travel time coefficient might compensate for the more imprecisely measured reward, which translates into a (slightly) lower reward coefficient.

Comparing the VOT in Models 1 and 3, we find substantial overestimation of the VOT in Model 3 (66 Euro/h if departure times are defined as 'known' and 70 Euro if departure times are defined as 'unknown'), meaning that the point estimates of the VOT between Models 1 and 3 differ by more than a factor 2. The difference is significant at the 5% level for 'known' departure times and almost significant at the 5% level for 'unknown' departure times. The main underlying mechanism is straightforward: Not accounting for the correlations in travel times across links (as the GWR- and OLS-based models do) leads to an underestimation of the difference between peak and off-peak door-to-door travel times as only variations on the C1–C2 link are considered. In turn this results in an overestimation of the absolute time coefficient and the VOT.

Unlike the VOT, the VSDE and the VSDL are similar across all models, ranging between 9 and 12 Euro/h and between 7 and 9 Euro/h, respectively, with the differences to the estimates of Model 1 being insignificant. Different travel time (mis-)specifications are thus mainly reflected in different travel time coefficients, whereas the remaining coefficients remain surprisingly stable. The latter results is another indication that the assumption of a fixed speed of 70 km/h is a good representation of the average speed on the home–C1 and C2–work links, as it does not induce a structural overestimation of schedule delay early (in the case of too high speed) or schedule delay late (in case of too low speed). Also the relative weight attached to the current travel times relative to average travel times, θ, is quite stable across the models, ranging between 0.10 and 0.14. Departure time decisions are thus more affected by past travel times (and the resulting attributes) than by travel times on the day of travel. This is plausible, since travel times on the day of travel are unknown in advance, and, moreover, drivers may find it difficult to change their routines from day to day. However, the fact that θ is significantly different from 0 in all models shows that travel time expectations differ across days. This finding proves the relevance of using models that are able to approximate day-specific travel times on those links for which no direct travel time measurements are available, rather than using reported travel times or travel times derived from network models, both of which are usually not day-specific.

8. Conclusions

We use geographically weighted regression (GWR) to estimate the location-specific relationship between speeds on links for which only sparse GPS data are available and speeds on a main link for which continuous speed measurements from loop detectors are available. We are able to show that speeds can be predicted significantly better using the GWR methodology rather than OLS, indicating that the correlation pattern between speeds on the link with continuous speed measurements and the links with GPS measurements varies over space. The estimated location-specific coefficients are then used to approximate driver-, day- and time-of-day-specific door-to-door travel times, which in turn are used as input for computing the travel time, scheduling and reward attribute of a (discrete) departure time choice model.

The choice model is based on data of more than 400 Dutch drivers who participated in a peak avoidance experiment over a period of around 4 months. They could gain a reward of 4 Euro per weekday for not passing a certain link, C1–C2, during morning peak hours. Due to the reward that varies by time of the day, monetary valuations of time and schedule delays can be derived. The data requirements for estimating choice models based on real (revealed preference: RP) data are high: In contrast to choice data gathered in a hypothetical setting (stated preference: (SP) data) the possible choice alternatives as well as the attributes of both chosen and unchosen alternatives must be defined by the researcher. This is frequently an ambiguous task, as it is not clear which information one had when making the choice.

In this paper, we focus specifically on the difficulties related to defining door-to-door travel times when travel times are only available on a specific sub-link (C1–C2), whereas in reality departure time choices are clearly determined by door-to-door travel times. We apply GWR combining the continuously measured travel times on the C1–C2 link with sparse GPS data in the areas where home and work locations of the participants are situated. Comparing the results of choice models whose attribute values are computed using GWR with more naïve methods, we find that the value of travel time (VOT) is strongly biased if the (usually positive) travel time correlations between the C1–C2 link and the remaining parts of the door-to-door trip are ignored. More specifically, ignoring the travel time correlations leads to an underestimation of the differences in travel times across choice alternatives, and consequently to an overestimated VOT. As a result, the point estimates of the VOTs can differ by up to a factor 2.

Our paper also demonstrates how an RP model can be set up without using reported data – besides the preferred arrival time – neither for travel nor for departure times, thus avoiding biases that might result from the use of self-reported data. We also show that the estimates obtained from RP models strongly depend on how (expected) travel times are defined, thus, potentially explaining some of the variation in the values of time and schedule delay found in the literature. While the choice modeling literature often emphasizes that one of the main advantage of RP data over SP data is the fact that they constitute 'real choices', we show that RP-based estimates can heavily depend on how the attributes of the choice model are specified. RP estimates may not suffer from hypothetical biases, but they can still be substantially biased if the choice attributes are defined in a way that is inconsistent with the attribute values that individuals actually consider when deciding on their departure times.

In this paper, we find that among the alternative travel time measures, the travel times derived using GWR are closest to what drivers take into account when making their choices, since these can best explain the observed departure time choices. Nevertheless, one might still argue that the GWR-based travel times could be very different from the travel times that travelers actually take into account, also because drivers might not be familiar with all departure time alternatives, or perceive frequently chosen alternatives as more favorable (a common phenomenon that is often referred to as 'choice-supportive bias'). While an accompanying study by Peer et al. (2013) finds that drivers strongly overreport travel times, and that they do so slightly more for departure times alternatives they are barely familiar with, it finds no evidence that such over-reporting actually affects behavior: the degree of over-reporting does not explain the choices made. It is probably not possible (due to obvious endogeneity) to test whether drivers overestimate travel times more for relatively unfamiliar alternatives also in their real-life choice situations, however, even if it was, our main interest would still be the reaction of drivers to 'objective' rather than 'perceived' travel times, as also for policy evaluations typically only 'objective' travel time measurements are available.

While here we focused on the application of the GWR method to calculate door-to-door travel times as an input for departure time choice experiments, the method developed in this paper is generally applicable for travel time predictions in situations where continuous speed measurements are only available on some links and GPS data are available on other links. We expect such travel time predictions to gain relevance in the near future as more and more GPS data become available. Another feasible application is to use the travel times derived from the GWR model as pivots in SP experiments, and thus to create choice experiments that seem more realistic to respondents and are thus less prone to be affected by hypothetical biases. Given the difficulties of setting up an RP experiment and to avoid the correlations in the attribute levels inherent to many RP experiments, such a procedure can often be useful. Hybrid models that enrich SP experiments with real (e.g. GPS) data (as well as their counterparts that enrich RP data with hypothetical choice alternatives (e.g. Washington et al., 2012)) are expected to gain momentum in the future. The GWR methodology developed in this paper can thus serve as an important method for defining travel time attributes in both RP and SP settings.

Acknowledgements

We would like to thank two anonymous referees for their helpful comments. The Dutch Ministry of Infrastructure and the Environment as well as by the TRANSUMO foundation are thanked for their financial support. Also funding from the ERC (Advanced Grant OPTION #246969) for the research of Paul Koster and Erik Verhoef, and from IPRISM for the research of Jasper Knockaert is gratefully acknowledged. We thank Pascal Neis (University of Heidelberg) for providing us with access to the OpenRouteService platform and Chris Jacobs for assistance with GIS.

Abkowitz, M., 1981. An analysis of the commuter departure time decision. Transportation 10 (3), 283–297.

Aultman-Hall, L., Du, J., 2006. Using spatial analysis to estimate link travel times on local roads. In: Transportation Research Board 85th Annual Meeting. Washington, DC.

Bhat, C., 1998a. Accommodating flexible substitution patterns in multi-dimensional choice modeling: formulation and application to travel mode and departure time choice. Transportation Research Part B 32 (7), 455–466.

Bhat, C., 1998b. Analysis of travel mode and departure time choice for urban shopping trips. Transportation Research Part B 32 (6), 361–371.

Bhatta, B., Larsen, O., 2011. Errors in variables in multinomial choice modeling: A simulation study applied to a multinomial logit model of travel mode choice. Transport Policy 18 (2), 326–335.

Bierlaire, M., Chen, J., Newman, J., 2010. Modeling route choice behavior from smartphone GPS data. Report TRANSP-OR, Transport and Mobility Laboratory, Ecole Polytechnique Fédérale de Lausanne (101016).

Börjesson, M., 2008. Joint RP-SP data in a mixed logit analysis of trip timing decisions. Transportation Research Part E 44 (6), 1025–1038.

Bowman, A., Hall, P., Prvan, T., 1998. Bandwidth selection for the smoothing of distribution functions. Biometrika 85 (4), 799.

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

Brunsdon, C., Fotheringham, S., Charlton, M., 1998. Geographically weighted regression. Journal of the Royal Statistical Society: Series D (The Statistician) 47 (3), 431–443.

Carrion, C., Levinson, D., 2013. Valuation of travel time reliability from a GPS-based experimental design. Transportation Research Part C 35, 305–323.

Castro, P., Zhang, D., Li, S., 2012. Urban traffic modelling and prediction using large scale taxi GPS traces. Pervasive Computing, Lecture Notes in Computer Science 7319, 57–72.

Charlton, M., Fotheringham, A., 2009. Geographically weighted regression (White Paper). National Centre for Geocomputation, National University of Ireland Maynooth

Chen, W., Li, Z., Yu, M., Chen, Y., 2005. Effects of sensor errors on the performance of map matching. Journal of Navigation 58 (2), 273–282.

Choi, K., Chung, Y., 2002. A data fusion algorithm for estimating link travel time. ITS Journal 7 (3-4), 235–260.

Chu, L., Oh, S., Recker, W., 2005. Adaptive Kalman filter based freeway travel time estimation. In: Transportation Research Board 84th Annual Meeting. Washington, DC.

Cleveland, W., 1979. Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association, 74 (368), 829–836.

Cosslett, S., 1977. The trip timing decision for travel to work by automobile. The Institute of Transportation Studies, The University of California, Berkely and Irvine, CA.

Daly, A., Hess, S., 2011. Simple approaches for random utility modeling with panel data. In: Transportation Research Board 90th Annual Meeting. Washington, DC.

De Fabritiis, C., Ragona, R., Valenti, G., 2008. Traffic estimation and prediction based on real time floating car data. In: Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems. Madeira, Portugal.

El Esawey, M., Sayed, T., 2012. A framework for neighbour links travel time estimation in an urban network. Transportation Planning and Technology 35 (3), 281–301.

Fan, J., Gijbels, I., 1996. Local polynomial modelling and its applications. Chapman & Hall/CRC, London,UK.

Fifer, S., Greaves, S., Rose, J., Ellison, R., 2011. A combined GPS/stated choice experiment to estimate values of crash-risk reduction. Journal of Choice Modelling 4 (1), 44–61.

Ghosh, A., 2001. Valuing time and reliability: Commuters' mode choice from a real time congestion pricing experiment. Doctoral dissertation, University of California, Irvine, CA.

Hackney, J. K., Bernard, M., Bindra, S., Axhausen, K. W., 2007. Predicting road system speeds using spatial structure variables and network characteristics. Journal of Geographical Systems 9 (4), 397–417.

Harris, P., Fotheringham, A., Crespo, R., Charlton, M., 2010. The use of geographically weighted regression for spatial prediction: An evaluation of models using simulated data sets. Mathematical Geosciences 42 (6), 657–680.

Hensher, D., 2010. Hypothetical bias, choice experiments and willingness to pay. Transportation Research Part B 44 (6), 735–752.

Herring, R., Hofleitner, A., Abbeel, P., Bayen, A., 2010. Estimating arterial traffic conditions using sparse probe data. In: Proceedings of the 13th International IEEE Conference on Intelligent Transportation Systems. Madeira, Portugal.

Idé, T., Kato, S., 2009. Travel-time prediction using Gaussian process regression: A trajectory-based approach. In: Proceedings of the 9th SIAM International Conference on Data Mining (SDM). Sparks, NV.

Jenelius, E., Koutsopoulos, H., 2013. Travel time estimation for urban road networks using low frequency probe vehicle data. Transportation Research Part B 53, 64–81.

Kamarianakis, Y., Prastacos, P., 2005. Space–time modeling of traffic flow. Computers & Geosciences 31 (2), 119–133.

Knockaert, J., Ettema, D., Verhoef, E., Koster, P., Peer, S., Tseng, Y.-Y., 2012. Spitsmijden Gouda-Zoetermeer. URL `http://www.spitsmijden.nl/downloads/SM2D_scientific_report_final.pdf`

Knockaert, J., Tseng, Y.-Y., Verhoef, E., Rouwendal, J., 2012. The Spitsmijden experiment: A reward to battle congestion. Transport Policy 24, 260–272.

Krause, A., Horvitz, E., Kansal, A., Zhao, F., 2008. Toward community sensing. In: Proceedings of the 7th International Conference on Information Processing in Sensor Networks. IEEE Computer Society, St. Louis, MO.

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

Leduc, G., 2008. Road traffic data: Collection methods and applications. Working Papers on Energy, Transport and Climate Change: European Commission Joint Research Centre, Institute for Prospective Technological Studies.

Li, Z., Hensher, D., Rose, J., 2010. Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. Transportation Research Part E 46 (3), 384–403.

McCafferty, D., Hall, F., 1982. The use of multinomial logit analysis to model the choice of time to travel. Economic Geography, 236–246.

Miura, H., 2010. A study of travel time prediction using universal kriging. Top 18 (1), 257–270.

Modelit, 2009. TRIP Handleiding.
URL http://www.modelit.nl/modelit/pdf/MVManual.pdf

Peer, S., Knockaert, J., Koster, P., Verhoef, E., 2013. Overreporting vs. Overreacting: Commuters' Perceptions of Travel Times. Tinbergen Institute Discussion Paper TI 2013-123

Peer, S., Koopmans, C., Verhoef, E., 2012. Prediction of travel time variability for cost-benefit analysis. Transportation Research Part A 46, 79–90.

Quddus, M. A., Ochieng, W. Y., Noland, R. B., 2007. Current map-matching algorithms for transport applications: State-of-the art and future research directions. Transportation Research Part C 15 (5), 312–328.

Ramos, G. M., Frejinger, E., Daamen, W., Hoogendoorn, S., 2012. A revealed preference study on route choices in a congested network with real-time information. In: Proceedings of the 13th International Conference on Travel Behaviour Research. Toronto, Canada, pp. 15–20.

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

Small, K., 1987. A discrete choice model for ordered alternatives. Econometrica: Journal of the Econometric Society 55 (2), 409–424.

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

Stone, C., 1977. Consistent nonparametric regression. The Annals of Statistics 5 (4), 595–620.

Train, K., 2003. Discrete choice models with simulation. Cambridge University Press, Cambridge, UK.

Tseng, Y., Knockaert, J., Verhoef, E., 2013. A revealed-preference study of behavioural impacts of real-time traffic information. Transportation Research Part C, 30, 196–209 .

Tseng, Y., Rietveld, P., Verhoef, E., 2005. A meta-analysis of valuation of travel time reliability.
URL http://www.transumofootprint.nl

Van Lint, J., 2010. Empirical evaluation of new robust travel time estimation algorithms. Transportation Research Record 2160, 50–59.

Vickrey, W., 1969. Congestion theory and transport investment. The American Economic Review 59 (2), 251–260.

Wang, H., Yue, Y., Li, Q., 2013. How many probe vehicles are enough for identifying traffic congestion?–A study from a streaming data perspective. Frontiers of Earth Science 7 (1), 34–42.

Washington, S., Ravulaparthy, S., Rose, J., Hensher, D., Pendyala, R., 2012. Bayesian imputation of non-chosen attribute values in revealed preference surveys. Journal of Advanced Transportation (doi: 10.1002/atr.201).

Westgate, B. S., Woodard, D. B., Matteson, D. S., Henderson, S. G., 2011. Travel time estimation for ambulances using Bayesian data augmentation. School of Operations Research and Information Engineering, Cornell University.

Wheeler, D., Páez, A., 2010. Geographically weighted regression. Handbook of Applied Spatial Analysis, 461–486.

Zito, R., d'Este, G., Taylor, M. A., 1995. Global positioning systems in the time domain: How useful a tool for intelligent vehicle-highway systems? Transportation Research Part C 3 (4), 193–209.

Zou, H., Yue, Y., Li, Q., Yeh, A. G., 2012. An improved distance metric for the interpolation of link-based traffic data using kriging: a case study of a large-scale urban road network. International Journal of Geographical Information Science 26 (4), 667–689.