VU Research Portal

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

Peer, S.; Knockaert, J.S.A.; Koster, P.R.; Tseng, Y.; Verhoef, E.T.

2011

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Peer, S., Knockaert, J. S. A., Koster, P. R., Tseng, Y., & Verhoef, E. T. (2011). *Door-to-door travel times in RP departure time choice models: An approximation method based on GPS data*. (TI Discussion Papers; No. 11-180/3). Tinbergen Institute (TI). <http://www.tinbergen.nl/discussionpapers/11180.pdf>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

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 is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

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 is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

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

Stefanie Peer (VU Amsterdam, Tinbergen Institute)
Jasper Knockaert (VU Amsterdam)
Paul Koster (VU Amsterdam)
Yin-Yen Tseng (VU Amsterdam)
Erik Verhoef (VU Amsterdam, Tinbergen Institute)

May 17, 2012

Abstract

A common way to determine values of travel time and schedule delay is by estimating departure time choice models using revealed preference (RP) data. The estimation of such models requires that (expected) travel times are known for both chosen as well as unchosen departure time alternatives. As the availability of such data is limited, most departure time choice studies 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. To approximate door-to-door travel times for which no complete measurement is possible, we develop a model that relates travel times on links with continuous speed measurements to travel times on links where relatively sparse 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 calculation 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 should be enriched with GPS data.

1 Introduction

Congestion tends to be strongly related to the time of the day, mainly because commuters start and end work at similar hours. Therefore, in order to evaluate policies that mitigate congestion, models are needed that take these time of day related patterns into account. The dominant modeling approach is to assume that travelers choose their optimal departure time, trading off the costs of travel time and the costs of arriving earlier or later than their preferred arrival time. If a monetary component that varies over the time of the day is included in the model, it is possible to compute the values of time and schedule delay, which can then be used as an input for the evaluation of transport policy options. The underlying idea of this so-called scheduling model goes back to the early work of Vickrey (1969) and Small (1982). Scheduling models are usually 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.

Scheduling models have been widely applied using stated preference (SP), revealed preference (RP), and combined SP–RP data. Overviews of empirical research can be found in Brownstone and Small (2005), Tseng et al. (2005) and Li et al. (2010). While SP data have the advantage that the researcher is in control of the choice set and the attributes of the alternatives, the results may depend strongly on the design of the experiment, and may be subject to hypothetical biases (Hensher, 2010). RP data, in contrast, are based on observed behavior, and therefore do not suffer from hypothetical biases. Nevertheless, they are used in relatively few studies, mainly due to requirements for detailed and high-quality data.

This paper focuses on the use of door-to-door travel time data in RP-based departure time choice models. So far, many studies that include RP-based departure time choice models do not take into account door-to-door travel times, or only very rough measures of door-to-door travel times. 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 been undertaken that explicitly and systematically surveyed the effects of using imprecise measures of (expected) travel times on the valuations of time and schedule delay.

We develop a model that allows for the (ex-post) approximation of driver-, day- and time-of-day-specific door-to-door travel times using GPS data. We make use of the travel times resulting from this model to determine the attributes of the departure time alternatives in a choice model. We show that less precise travel time definitions, that 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. We find that ignoring the spatial variation of speeds typically leads to measurement errors of the travel times, which in turn lead to a downward bias of the value of time. Not accounting for the temporal variation of speeds implies ignoring the fact that speeds tend to be positively correlated between the different parts of the door-to-door trip. For instance, it is likely that in the middle of the morning peak, speeds are relatively low on all parts of the trip. If these correlations are ignored, the differences between peak and non-peak travel times tend to be underestimated, and as a result, the (absolute) time coefficient tends to be overestimated.

The RP data used in this paper were gathered during a peak avoidance experiment ('Spitsmijden' in Dutch). Drivers were eligible for a monetary reward if they avoided traveling on a specific highway link (called '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. However, in this paper, we only consider a subgroup consisting of roughly one third of these participants for whom preferred arrival times at work are known. During that second phase only travel times along the C1-C2 link were measured. Figure 1 shows a schematic map of the experimental setup.

Figure 1: Schematic Map of the Experimental Setup

In order to compute driver-, day and time-of-day-specific door-to-door travel times, we use geographically weighted regression (GWR). GWR is a spatial version of local (non-parametric) regression, which goes back to the seminal papers of Stone (1977) and Cleveland (1979). Our model explicitly takes into account travel time correlations across trip parts, in particular between the links for which only sparse GPS observations are available (home-C1, C2-work) and the link on which travel times are observed continuously (C1-C2). GWR allows for spatial heterogeneity in the correlation pattern. For a given speed on the C1-C2 link, predicted speeds along the home-C1 (C2-work) links are more similar between home (work) locations that are located closer to each other, but can differ more strongly between locations that are further apart. Figure 2 underlines the relevance of this approach. It plots the observed speeds leading from the home locations to camera C1 and from C2 to the work locations across space. Even without accounting for different times of the day, and therefore recurrent congestion levels, the spatial variation of speeds is substantial, with higher speeds occurring for locations closer to highways.

Figure 2: Speed Observations

While we use GPS data to approximate door-to-door travel times, also reported travel times or travel times generated from network models could be used for this purpose. However, both have in common that they are usually only available for a very limited number of departure times (e.g. peak vs. off-peak travel times), among which the researcher then needs to interpolate in order to approximate travel times for a large number of departure time alternatives. Furthermore, reported and network travel times are usually generic rather than day-specific. However, we will show in this paper that drivers have travel expectations that vary across days. 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 overestimated with a factor 1.5 (Knockaert et al., 2012).

The structure of the paper is as follows. Section 2 contains a review of past literature on RP-based departure time choice models. In the third section, 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 using door-to-door travel times, the data used for the choice models as well as alternative model specifications. Section 7 presents the results of the choice models, and Section 8 concludes.

2 Literature

A limited number of studies exists that estimate departure time choice models using RP data, either as only data source, or in combination with SP data¹. Most older RP-based studies assume that travel times do not vary between days. For instance, Abkowitz (1981), Cosslett (1977) and Small (1982, 1987) 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 calculated for 12 alternative

¹E.g. Börjesson (2008); Brownstone and Small (2005); Ghosh (2001); Small et al. (2005)

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 using RP-based departure time choice models 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 needs 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 the day. However, she also 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 general factors such as weather conditions.

Lam and Small (2001) only observe travel times directly along a 10 mile corridor for their joint route and departure time choice model. They estimate two models: The first model only takes into account travel times along this corridor, whereas the second model assumes that drivers face the same speeds as observed along the corridor also on the rest of their travel (except for a 5 mile access link, which is assumed to be uncongested). 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 less surprising. Besides the first model of Lam and Small (2001). Knockaert et al. (2009) and Tseng et al. (2011) only account for travel times on the specific highway link where the experiment took place (close by the one considered in this paper), without testing alternative travel time specifications.

3 Geographically Weighted Regression

3.1 Methodology

We use geographically weighted regression (GWR) to explain speeds on the links for which only GPS observations are available (home–C1, C2–work) by the speeds on the link on which continuous speed measurements are undertaken (C1–C2). 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. Thus, in the GWR model coefficients can differ over space (in our application across home or work locations, respectively). The location specific coefficients are then more similar for (home or work) locations situated close to each other, and can differ more strongly for locations further apart from each other. GWR has the advantage 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. Brunsdon et al., 1998; Charlton and

Fotheringham, 2009; Fan and Gijbels, 1996). The weight matrix is determined by the relative geographic locations. 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. Using spatial interpolation, we can also use them for predicting 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 then given by:

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

where $\hat{\lambda}(u)$ is the local parameter estimate, y is the $N \times 1$ column vector including the values of the dependent variable (the speeds on the GPS observed links) for all observations $i = 1, \dots, N$, and X is the $N \times p$ matrix of covariates (for p covariates, i.e. 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.

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

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 :

$$w_i(u) = f(d_i(u), h) \quad (3)$$

The functional form of $f(d_i(u), h)$ is assumed to be Gaussian².

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

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

²The scale of the weights does not matter as $[X^T(\alpha W(u))X]^{-1} X^T(\alpha W(u))y = 1/\alpha [X^T W(u)X]^{-1} X^T \alpha W(u)y = [X^T W(u)X]^{-1} X^T W(u)y$.

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 :

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

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 over 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 tend to depart at different times of the day. Since we intend to use the coefficients for out-of-sample predictions, it is essential 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.³ 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

We use GWR to explain speeds on the GPS-observed links (home–C1 and C2–work) by speeds on the link with continuous speed measurements (C1–C2). Separate models are estimated for the home–C1 links and the C2–work links. The speeds on the home–C1 (C2–work) link for a specific home (work) location u are denoted by $v_t^{gps}(u)$, where $\tilde{t}$ is the departure time at home (C2). The speeds on the C1–C2 link when passing C1 at time t are denoted by v_t^{cont} . Note that, while $\tilde{t}$ is fixed for a given GPS observation, t needs to be defined relative to $\tilde{t}$. We define the relationship between $\tilde{t}$ and t differently for the home–C1 and C2–work case. In both cases, we denote the passage time at C1 that is defined as optimal for a given $\tilde{t}$ by $\bar{t}$.

In the home–C1 case, for an observed departure time from home, $\tilde{t}$, we use the C1–C2 speed also at time $\bar{t}$ 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 home–C1 speeds but at the same time the home–C1 speeds would determine the passage time at C1. 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

³These 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 interruptions, which might be due to stops the driver undertakes but also from missing GPS data.

time at C1 that results in the observed passage time at C2, $\tilde{t}$, given C1–C2 speeds v_t^{cont} :

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

The (location-specific) model that relates the GPS observed speeds to the C1–C2 speeds is assumed to have a local linear form:

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

Equation 7 establishes a simple direct relation between the speeds on the two links. It performs better in out-of-sample prediction than more complicated model structures. We tested models that include departure time as explanatory variable both in linear⁴ as well as nonlinear⁵ specifications. While the overall fit is comparable to the fit of the model based on Equation 7, these functions tend to yield large outliers at the begin and the end of the peak. Finally we also tested 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 models⁶, it has the disadvantage that a different set of coefficients is valid for different times of the day, making the applicability of the model less general. Equation 7, on the other hand, is applicable to all times of the day. Moreover, it does not require an assumption on the shape of the peak. Weekday- and weather-specific influences on door-to-door travel times are captured by the speeds on the continuously observed link.

4 Data

4.1 Experiment

The data used in this study were gathered from a peak avoidance (in Dutch: 'Spitsmijden') experiment taking place in the Netherlands. The participants could obtain a reward of 4 Euro if they did not use a specific highway link of 9.21 km (C1–C2) during the morning peak hours (6:30–9:30). Participants were not selected randomly but were invited for participation if they were observed to pass C1–C2 link multiple times a week. In the first phase of the experiment (11/2008–4/2009), 230 drivers were equipped with customized smartphones, also named Rabomobiël (Spitsmijden-Group, 2009). The 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 used to approximate door-to-door travel times for the second phase (09/2009–12/2009), where no GPS data were collected and therefore travel time measurements are only available for the C1–C2 link. The number of participants in the second

⁴E.g. $v_{\tilde{t}}^{gps}(u) = \lambda_0(u) + \lambda_1(u)v_{\tilde{t}}^{cont} + \beta_3(u)\bar{t} + \beta_4(u)\bar{t}^2$

⁵E.g. $v_{\tilde{t}}^{gps}(u) = \lambda_1(u) - \beta_2(u)(1 + \beta_3(u)v_{\tilde{t}}^{cont})e^{-(\beta_4(u)\bar{t} - \beta_5)^2} + \beta_6 v_{\tilde{t}}^{cont}$

⁶The 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.

phase was close to 2000, of which approximately one third are considered in this paper. These are the ones for whom preferred arrival times at work are known. An in-depth description of the Spitsmijden experiment can be found in Knockaert et al. (2012).

4.2 Speed measurements

We use two types of speed measurements, GPS-observed speeds for the home–C1 and C2–links and speeds derived from loop detector data for the C1–C2 link.

In their raw form, GPS measurements are a series of data points with a location, speed and time stamp attached to them. From these data points, speed and distances can be calculated. We take into account GPS-observed trips that passed either through camera location C1, C2 or both. We construct separate datasets for travel times observed on the home–C1 and those observed on the C2–work links, in order to allow for using the GWR model to approximate travel times on a maximum number of start/end location combinations. A commuting trip (either along a home–C1 or C2–work link) 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. The distance is defined as the sum of the distances between the GPS observations belonging to the same trip.⁷ This approach yields very similar results compared to defining distance based on a 'fastest network distance'.⁸ We only consider trips along the home–C1 and C2–work links that were observed during the morning (between 5:00 and 11:00 a.m.) and on weekdays. We exclude trips that result in 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)⁹ are available. Also, the largest distance between 2 GPS measurements (for a given trip) must not exceed 2 km.

The GPS observed speeds are connected to C1–C2 speeds according to the temporal relations between the the starting time of the home–C1 (C2–work) trip and the passage times on the C1–C2 link established in Equation 6. The speeds on the C1–C2 link are available for any time of the day from a dense network of loop detectors. The speeds from these loop detectors are aggregated in time (15-minute intervals) and space (from the single detectors towards the entire link). The trajectory method is used for this purpose. An overview of this method can be found in Van Lint (2010), whereas the exact approach employed for the data used in this paper 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), increasing confidence that we can safely use the relation derived between the GPS-based speeds

⁷Since location measurements using GPS can be slightly imprecise, it is possible that the GPS device transmits different locations although the vehicle is standing still. If these distances are added to the overall trip distance, this overall distance might be overestimated, resulting in a structural downward bias of travel times. 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.

⁸Retrieved from openrouteservice.org

⁹Less than 10 observations within a perimeter of 5 km of the reference location.

and the camera observations during the first phase to make travel time predictions during the second phase.

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 as well as 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. Table 1 also 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. Moreover, Table 1 demonstrates the importance of using door-to-door travel times. While the length of the C1–C2 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 a very substantial part of the door-to-door trips.

Table 1: Descriptives GPS

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)	23.13 min	15.14 min
Mean distance (per trip)	30.91 km	16.34 km
Mean speed (per trip)	78.69 km/h	69.29 km/h
Mean number of GPS measurements (per trip)	1288	784
Mean time btw. 2 GPS measurements	1.09 s	1.15 s
Mean distance btw. 2 GPS measurements	23.36 m	21.61 m
Maximum distance btw. 2 GPS measurements	85.14 m	115.43 m
Overall st. dev. dep. (arr.) time at home/work	90.60 min	87.60 min
Median driver-specific st. dev. dep. (arr.) time at home/work	32.20 min	34.03 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

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 3: Optimal bandwidth choice

Figure 4 shows the regarding weights attached to the distance between reference location u and another observation i . It shows that the weights are subject to a steeper distance decay for the case of the C2-work links than for the home-C1 links. This might be the result of the relatively low dispersion of work locations in space, the higher number of observations for the C2-work case as well as a lower variance of local traffic conditions among work locations.

Figure 4: Kernel weights

A comparison of the optimal GWR model to the corresponding OLS model reveals that indeed the root mean squared error (RMSE) decreases and the R-squared increases substantially if one moves from OLS to GWR (Table 2). While these results hold for both the home-C1

and the C2–work links, they are stronger for the home–C1 links. This confirms that spatial heterogeneity plays a larger role for the home–C1 link.

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

Figure 5a shows a map of the area within which the home and work locations are situated. Figures 5b and 5c show the spatial pattern that results if speeds are predicted using the location-specific coefficients derived from the GWR model, assuming a C1–C2 speed of 50 or 100 km/h, respectively. These figures can be interpreted as a peak (50 km/h) and an off-peak (100 km/h) scenario. As expected, higher speeds are predicted for observations starting/ending along the highway, while lower speeds are predicted for observations starting/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 precise travel time definitions are used.

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 a sequence of $k = 1, \dots, K$ choices among the alternatives $j = 1, \dots, J$, where $J = 18$, as departure time intervals of 15 minutes between 5:30 and 10:00 a.m. are considered. K is equal to the duration of the experiment of 75 weekdays. As most drivers

Figure 5: Speed Predictions

(a) Map

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

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

were not observed to pass the C1–C2 link on all 75 days, we only use the driver-specific subset of days, during which a driver has been observed to travel in the estimation.

The driver chooses the alternative that maximizes the following random utility function:

$$U_{zkj} = V_{zkj} + \epsilon_{zkj}, \quad (8)$$

The random utility function consists of a deterministic component V_{zkj} and a random component ϵ_{zkj} that follows a Gumble 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, 2010). The choice probability of alternative $\check{j}$ is given by

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

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 .

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

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 coefficients. They can be interpreted as the values attached to changes in expected travel times and schedule delays, respectively. The ratios are multiplied by (-1) in order to account for the fact that the monetary component is defined here as a reward.

$$\text{VOT} = -\beta_T/\beta_R \quad \text{VSDE} = -\beta_E/\beta_R \quad \text{VSDL} = -\beta_L/\beta_R \quad (11)$$

To determine the expected values of the attributes $A = \{R, T, SDE, SDL\}$, we have to take into account that the travel time realizations will become known to the traveller only after the departure time has been chosen. We model this by assuming that expectations are based on average travel times, and on specific information - unknown to the researcher - for that particular day. The expected values of the attributes are thus based on a compound measure of average travel time across the entire experiment, thus across days $k = 1, \dots, K$ ¹⁰ and the realized travel time on the day of travel, $\check{k}$, 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 extent of information available to drivers. The average travel times represent the long-run pattern of

¹⁰A similar definition of expected travel times was used by Tseng et al. (2011), however, they do not apply the compound measure to all attributes of the utility function but to travel times only.

travel times over the time of the day, which commuters 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:

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

If a driver z is eligible for a reward on day k^{11} , 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 preferred arrival time in the case no congestion would occur.

$$\begin{aligned} SDE_{zkj} &= \max(0, PAT_z - t_{zkj} - T_{zkj}) \\ SDL_{zkj} &= \max(0, t_{zkj} + T_{zkj} - PAT_z) \end{aligned} \quad (13)$$

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 time from home.

For all drivers, home and work locations are known, as well as the shortest network distances¹² 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 apply the Gaussian¹³ 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 $i = 1, \dots, N$ 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 $\bar{u}$.

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

¹¹Drivers 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.

¹²Retrieved from openrouteservice.org

¹³We 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.

We can use these coefficients and the shortest network distances to compute 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 Equation 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 improve 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 (17.73% of drivers). Moreover, we exclude drivers with a combination of $\hat{\lambda}_0(\bar{u}), \hat{\lambda}_1(\bar{u})$ that can lead to speed predictions larger than 110 km/h, given that the highest measured C1–C2 speed equals 107 km/h (1.77% of drivers).

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 average length of the home–C1 link and C2–work link is smaller than the regarding distances considered in the original GWR estimation (see Table 1). 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 go about 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. However, for a given driver, day and departure time, the C1–C2 speeds may still differ across the models as the passage time of the C1–C2 link and therefore also the speed along this link are determined by the speed that is assumed on the home–C1 link.

Model 1 represents the standard model, where speeds on the home–C1 and C2–work links are computed from the GWR coefficients, $\hat{\lambda}_0$ and $\hat{\lambda}_1$ (as derived in Equation 14), and therefore vary across space, days and time of the day. Model 2 represents the case where speeds differ between days and time of the day, however, not across space. Home–C1 and C2–work speeds are independent from the home and work location. Thus, instead of the GWR coefficients, the OLS coefficients λ_0 and λ_1 from Table 2 are used to determine home–C1 and C2–work travel times. Finally, Model 3 represents the case where travel times do not vary across space, days

Table 3: Descriptives of Drivers considered in the Choice Model

Variables	Value or Sample Average	St. Dev.
<i>Commute-related variables</i>		
Preferred arrival time at work		
<7:00 a.m (in %)	11.21	–
>9:00 a.m. (in %)	12.31	–
Distance Home–C1 (in km)	22.38	15.69
Distance C2–Work (in km)	11.25	4.85
Distance to ‘closest GWR neighbor’: Home–C1 (in km)	0.74	0.95
Distance to ‘closest GWR neighbor’: C2–Work (in km)	0.24	0.41
$\hat{\lambda}_0^{Home-C1}$	49.75	16.10
$\hat{\lambda}_0^{C2-Work}$	46.10	10.57
$\hat{\lambda}_1^{Home-C1}$	0.32	0.10
$\hat{\lambda}_1^{C2-Work}$	0.25	0.08
<i>Choices</i>		
Nr. of choices per individual	20.95	12.98
Nr. of individuals	455	–
Total nr. RP choices	9530	–
Duration of the RP experiment (weekdays)	75	–

and time of the day. Speeds on all home–C1 and C2–work links are fixed to 80 km/h¹⁴ (i.e. $\lambda_0 = 80$, $\lambda_1 = 0$). 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 and the C1–C2 link.

Besides these three models, that differ in how home–C1 and C2–work travel times are defined, we also 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, departure times from home are calculated using the coefficients corresponding to the regarding model. For instance, in Model 2, both (home–C1, C2–work) travel times as well as departure times are then 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. In the second 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 the second scenario, this also allows

¹⁴This corresponds approximately to the mean C1–C2 speed.

us to compare the models directly in terms of their fit. Clearly, for Model 1 the 'known' and the 'unknown' departure time scenario coincide.

7 Results: Departure time choice models

Table 4 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. As the scale of the utility may differ between the models, we mainly compare the estimated VOT, VSDE and VSDL values, which are independent of the scale, rather than the values of the single coefficients (e.g. Train, 2003).

Table 4: Estimation results: Departure time choice models

	Model 1	Departure time			
		'Unknown'		'Known'	
		Model 2	Model 3	Model 2	Model 3
β_R	0.13	0.16	0.17	0.15	0.14
$t\text{-stat.}$	(3.57)	(4.06)	(4.49)	(3.90)	(3.92)
β_T	-5.56	-5.30	-8.62	-5.57	-9.21
$t\text{-stat.}$	(-4.81)	(-3.91)	(-3.68)	(-4.08)	(-3.98)
β_E	-1.73	-1.71	-1.67	-1.73	-1.66
$t\text{-stat.}$	(-10.29)	(-9.95)	(-10.04)	(-10.05)	(-10.10)
β_L	-1.44	-1.48	-1.51	-1.46	-1.50
$t\text{-stat.}$	(-15.72)	(-15.38)	(-15.22)	(-15.41)	(-15.27)
θ	0.15	0.17	0.14	0.15	0.12
$t\text{-stat.}$	(4.02)	(3.70)	(3.52)	(3.60)	(3.23)
VOT (Euro/h)	42.27	33.56	49.92	38.05	64.74
VSDE (Euro/h)	13.15	10.85	9.68	11.81	11.68
VSDL (Euro/h)	10.95	9.34	8.74	9.96	10.56
Nr. Obs.	9530	9530	9530	9530	9530
LogLikelihood	-23450	-23454	-23485	-23485	-23535
Pseudo R-squared	0.149	0.148	0.147	0.147	0.146

We obtain fairly high VOT estimates in all models (between 34 and 65 Euro/h). These can partly be explained by the rather high incomes of the participants (Knockaert et al., 2012, p.12), 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 (Knockaert et al., 2012).

Compared to the VOT of 42 Euro/h obtained in Model 1, VOTs obtained in Model 2 are relatively low (34 Euro/h if departure times are defined as 'unknown', and 38 Euro if departure times are defined as 'known'). One would expect that the finding that OLS coefficients predict travel times less precisely than GWR coefficients (see Table 2) results in a measurement error of the travel time, which biases downwards the time coefficient (Bhatta and Larsen, 2011;

Train, 2003). However, probably as a consequence of the experimental setup, we find that the lower VOTs are mainly the result of a higher reward coefficient rather than lower time coefficients. In the peak avoidance experiment, from which the data used in this paper are drawn, rewards are relatively insensitive towards changes in travel times. Since rewards tend to be (negatively) correlated to travel times (as only off-peak peak trips are rewarded, which are typically characterized by relatively low travel times), the reward coefficient might pick up also utility gains from lower travel times in the case that travel times measured imprecisely. This hypothesis is supported by the finding that when Model 1 and Model 2 are compared for the case of 'known' and therefore equal departure times, the t-statistic of the time coefficient is higher for Model 1 and the t-statistic for the reward coefficient is higher in Model 2.

Comparing the VOT in Models 1 and 3, we find substantial overestimation of the VOT in Model 3 (50 Euro/h if departure times are defined as 'unknown', and 65 Euro for Model 2 if departure times are defined as 'known'). In the scenario of 'known' departure times, the VOT of Model 1 is overestimated by 53%. This finding can be attributed to the fact that Model 3 ignores travel time correlations between the GPS-observed links and the C1–C2 link, which have been proven to exist (see Table 2). Not accounting for the correlations leads to an underestimation of the difference between peak and off-peak travel times, and therefore to 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 10 and 13 Euro/h and between 9 and 11 Euro/h, respectively. This is also an indication that the assumption of a fixed speed of 80 km/h is a good representation of the average speed on the home–C1 and C2–work links, as it does not lead to a structural overestimation of schedule delay early (in the case of a too high speed) or schedule delay late (in case of a 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 0.12 and 0.17. 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 relevancy 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.

As mentioned above, it is not useful to compare the fit of the models in the scenario where departure times are 'unknown', and therefore based on the coefficients of the regarding model. If we compare the model fit for the scenario where departure times are 'known', and therefore all models have the same dependent variable, we find that, as expected, Model 1 yields the highest log-likelihood, followed by Model 2 and Model 3.

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 are available. We find that the predictive power of these models improves considerably if one uses GWR rather than OLS. We use the coefficients resulting from the GWR model to approximate driver-, day- and time-of-day-specific door-to-door travel times. These are then used in the estimation of departure time choice models with attributes based on door-to-door travel times. We are able to show that not accounting for spatial variation of speeds leads to a downward bias of the value of time (VOT), whereas not accounting for the temporal variation in speeds results in an upward bias of the VOT. These findings seem to be generally valid in settings where speeds depend on the start and end locations of the trips and are positively correlated across trip parts.

Our paper also demonstrates how an RP model can be set up without using reported data neither for travel nor for departure times, thus avoiding any biases that might result from the use for reported data. We also show that the estimates obtained from RP models strongly depend on how (expected) travel times are defined, thus, explaining some of the variation in the values of time and schedule delay found in the literature. This finding contrasts existing literature on the comparison between SP and RP data, which usually emphasizes the effect of the SP design and the laboratory setting on the regarding estimates, whereas the coefficients obtained from the RP models are considered as 'true' estimates. Here, we show that also RP estimates are heavily dependent on how the attributes of the choice model are specified.

We focused on the application of the GWR method to calculating door-to-door travel times as an input for the departure time choice experiments, however, the method 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.

Acknowledgements

This study is financially supported by the Dutch Ministry of Infrastructure and the Environment, as well as by the TRANSUMO foundation. Also the financial support from the ERC (Advanced Grant OPTION #246969) for the research of Paul Koster and Erik Verhoef 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.

References

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

- 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: Methodological*, 32(7), 455–466.
- Bhat, C. (1998b). Analysis of travel mode and departure time choice for urban shopping trips. *Transportation Research Part B: Methodological*, 32(6), 361–371.
- Bhatta, B., and 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.
- Börjesson, M. (2008). Joint RP-SP data in a mixed logit analysis of trip timing decisions. *Transportation Research Part E: Logistics and Transportation Review*, 44(6), 1025–1038.
- Bowman, A., Hall, P., and Prvan, T. (1998). Bandwidth selection for the smoothing of distribution functions. *Biometrika*, 85(4), 799.
- Brownstone, D., and 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.
- Brunsdon, C., Fotheringham, S., and Charlton, M. (1998). Geographically weighted regression. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 47(3), 431–443.
- Charlton, M., and Fotheringham, A. (2009). Geographically weighted regression (White Paper). *National Centre for Geocomputation, National University of Ireland Maynooth Maynooth, Co Kildare, Ireland*, 1–17.
- Cleveland, W. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American statistical association*, 829–836.
- Cosslett, S. (1977). The trip timing decision for travel to work by automobile. In D. McFadden et al. (Eds.), *Demand model estimation and validation, the urban travel demand forecasting project*. The University of California, Berkely and Irvine: The Institute of Transportation Studies.
- Daly, A., and Hess, S. (2010). Simple approaches for random utility modelling with panel data. *European Transport Conference 2010*.
- Fan, J., and Gijbels, I. (1996). *Local polynomial modelling and its applications*. London,UK: Chapman & Hall/CRC.
- Ghosh, A. (2001). *Valuing time and reliability: Commuters mode choice from a real time congestion pricing experiment*. Unpublished doctoral dissertation, University of California.
- Hensher, D. (2010). Hypothetical bias, choice experiments and willingness to pay. *Transportation Research Part B: Methodological*, 44(6), 735–752.
- Knockaert, J., Ettema, D., Verhoef, E., Koster, P., Peer, S., and Tseng, Y.-Y. (2012). *Spitsmijden Gouda-Zoetermeer* [Scientific Report].
- Knockaert, J., Verhoef, E., and Rouwendal, J. (2009). Spitsmijden experiment: Reward to battle congestion. *Transportation Research Board 88th Annual Meeting*.
- Lam, T., and 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, Z., Hensher, D., and Rose, J. (2010). Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. *Transportation Research Part E: Logistics and Transportation Review*, 46(3), 384–403.

- McCafferty, D., and Hall, F. (1982). The use of multinomial logit analysis to model the choice of time to travel. *Economic Geography*, 236–246.
- Modelit. (2009). *TRIP Handleiding* [Technical Report]. Culemborg, Netherlands. Available from <http://www.modelit.nl/modelit/pdf/MVManual.pdf>
- Small, K. (1982). The scheduling of consumer activities: Work trips. *The American Economic Review*, 467–479.
- Small, K. (1987). A discrete choice model for ordered alternatives. *Econometrica: Journal of the Econometric Society*, 409–424.
- Small, K., Winston, C., and Yan, J. (2005). Uncovering the distribution of motorists' preferences for travel time and reliability. *Econometrica*, 73(4), 1367–1382.
- Spitsmijden-Group. (2009). *The effects of rewards in Spitsmijden 2: How can drivers be persuaded to avoid peak periods?* Netherlands. Available from www.spitsmijden.nl
- Stone, C. (1977). Consistent nonparametric regression. *The Annals of Statistics*, 5(4), 595–620.
- Train, K. (2003). *Discrete choice models with simulation*. Cambridge, UK: Cambridge University Press.
- Tseng, Y., Knockaert, J., and Verhoef, E. (2011). A revealed-preference study of behavioural impacts of real-time traffic information. *Transportation Research Part C: Emerging Technologies*.
- Tseng, Y., Rietveld, P., and Verhoef, E. (2005). *A meta-analysis of valuation of travel time reliability*. Available from <http://www.transumofootprint.nl>
- Van Lint, J. (2010). Empirical evaluation of new robust travel time estimation algorithms. *Transportation Research Record: Journal of the Transportation Research Board*, 2160, 50–59.
- Vickrey, W. (1969). Congestion theory and transport investment. *The American Economic Review*, 251–260.
- Wheeler, D., and Páez, A. (2010). Geographically weighted regression. *Handbook of Applied Spatial Analysis*, 461–486.