

Analysis of mode choice for intercity travel: application of a hybrid choice model to two distinct US corridors

Stephane Hess (corresponding author)

RSG, 55 Railroad Row, White River Junction, VT 05001

and

University of Leeds, Leeds LS2 9JT, United Kingdom

E-mail: s.hess@leeds.ac.uk & stephane.hess@rsginc.com

Greg Spitz

RSG, 55 Railroad Row, White River Junction, VT 05001

E-mail: gspitz@rsginc.com

Mark Bradley

RSG, 55 Railroad Row, White River Junction, VT 05001

E-mail: mark.bradley@rsginc.com

Matt Coogan

New England Transportation Institute, Ten Clay Road, White River Junction, VT 05001

E-mail: cooganmatt@aol.com

ABSTRACT

With growing concerns about greenhouse gas emissions and traffic congestion, there is an emphasis on encouraging shifts to public transport, for both short and long distance travel. Major differences exist across countries in how successful these efforts are, and the United States is often used as the key example of a country with a strong resistance to shifting away from private car use. Even within the United States however, there is strong heterogeneity across regions and across different types of travellers. This paper seeks to add empirical evidence to understand the drivers of mode choice for intercity travel, using stated choice data from two major US intercity corridors: the Northeast Corridor (NEC) and the Cascade Corridor. We develop a hybrid choice model that allows for deterministic and random variations across travellers in their preferences, some of which can be linked to underlying attitudinal constructs. Our results highlight extensive heterogeneity and provide interesting insights into the drivers of behaviour, and the relationship between attitudes and actual choices. As an example, we see that for some groups, notably West Coast respondents, a stronger anti-car attitude is counter-acted by a reduced utility for non-car modes when making choices, possibly due quality of public transport provision. Similarly, for other groups, such as older and female travellers, a reduced concern for privacy, which would benefit public transport, is counter-acted by a stronger pro-car attitude. These findings highlight the complex way in which attitudes can influence choices and provide insights for targeted policy interventions. Through scenario testing, we also show how future modal split might change depending on how these patterns of heterogeneity evolve over time, noting that the way this might happen is of course unknown at present.

Keywords: hybrid choice models; latent variables; attitudes; mode choice, Northeast Corridor, Cascade Corridor, NCRRP 03-02

1. INTRODUCTION

Understanding mode choice is of crucial interest for transport planning, and encouraging a shift from private car to public transport modes is an important component in efforts to reduce environmental impacts and ease congestion. While such efforts have been very successful especially for long distance travel in Europe,

the United States (US) is often used as an example of a country where people will drive long distance rather than rely on either ground or air based public transport¹. Even within the US, there is however extensive heterogeneity in modal preferences, both across population segments and across areas.

The NCRRP (2016) study on which the present paper is based aimed to gain a deeper understanding of mode choice by US travellers. We looked at two major US intercity corridors: the Northeast Corridor (NEC) and the Cascade Corridor on the West Coast.

The NEC is the biggest intercity corridor in the country and is in desperate need for infrastructure investment, which will cost many billions of dollars. It is of great importance to US policy makers due to its high demand and extreme infrastructure vulnerability. Nearly 56 million riders a year, approximately 200,000 riders on any given weekday, use the Hudson River rail tunnels, which are a critical link in the NEC system connecting New Jersey with New York City. These tunnels, over a hundred years old, were severely damaged in Hurricane Sandy and are the only heavy rail infrastructure that links New York City to points south and west. The tunnels are at maximum capacity throughout most hours of the day (they can carry up to 24 trains per hour). As happened in July 2015 for multiple days, problems with the tunnels significantly delayed and stranded tens of thousands of riders due to the vulnerability of this overtaxed infrastructure.

The Cascade Corridor is the name this study used for the corridor that goes along the Amtrak passenger train route called the “Cascades” in the Pacific Northwest. The route is operated by Amtrak and travels from Vancouver, British Columbia south to Eugene, Oregon, with major stops in Seattle and Portland and several smaller stops in between. It is named after the Cascade mountain range that the route parallels. Like the NEC, there are also auto and air alternatives for this corridor. The route from Vancouver to Eugene is roughly 466 miles, nearly the same mileage as the NEC. This route was used in the study as a “foil” to the NEC—to represent something other than the NEC. The NEC is unique to the US, with significant train services as might be found in Europe (headways for trains between NYC and Washington, DC of just 15 min). The Cascade corridor, meanwhile, is more typical of US intercity rail, with just four daily round trips between Portland and Seattle, with two daily round trips between Seattle and Vancouver, and two daily round trips between Eugene and Portland. Even with this relatively low frequency, the Cascade is Amtrak's eighth-busiest route, with a total annual ridership of 792,481 or roughly 2,000 riders per day.

There are also important differences between the two areas other than their rail service provision. While the NEC is quite urban and seen as “sophisticated”, the Pacific Northwest is more rugged and seen as “individualistic” (although it has a thriving urban and “grunge” scene in its own right). While more individualistic, the Pacific Northwest is often more environmentally aware and “greener” than other parts of the US, including the NEC. This presents important scope for heterogeneity in mode choice, only part of which is likely to be able to be explained by the analyst (Ortúzar & Willumsen, 2011). In addition to the influence of socio-demographic and trip characteristics as well as random variations, we tested whether some of this underlying heterogeneity, whether deterministic or random, could be linked to longer term attitudes of the traveller, and if this is different between NEC and Pacific Northwest, which would provide a core opportunity for interventions aimed at changing behaviour. Throughout, we were also keen to test for underlying geographical differences in behaviour – i.e., all else being equal, including level of service, does the behaviour differ between East Coast and West Coast travellers.

We conduct our analysis using a hybrid choice model or integrated choice and latent variable (ICLV) structure (Ben-Akiva et al., 1999a; Ashok et al., 2002; Ben-Akiva et al., 2002; Bolduc et al., 2005). Hybrid choice models have been used for a large variety of applications across different disciplines. In the context of attitudes in transport work, they have been used for example in the study of key decisions such as vehicle type (Glerum et al., 2013), mode choice (Atasoy et al., 2013; Kamargianni et al., 2014), route choice (Prato et al., 2012) and departure time choice (Thorhaug et al., 2016). They have been used to study the role of a

¹ Notwithstanding that the latter raises environmental concerns too.

wide variety of attitudes, ranging from privacy and security concerns (Daly et al., 2012) to environmental considerations (Kim et al., 2012). While an increasing number of applications have relied on stated choice (SC) data, a wide variety of applications also exist on revealed preference (RP) data, as highlighted by Kim et al. (2014).

Alongside numerous empirical applications, further refinements of the model framework have taken place, looking at the specification of the measurement model (Daly et al., 2012), how and where to incorporate the latent variables into the choice model (Bahamonde-Birke et al., 2017) and testing for non-linearity and distributional assumptions (Kim et al., 2016). Substantial efforts have also gone into improved estimation techniques for the model and proper identification (Bhat & Dubey, 2014; Daziano, 2015; Raveau et al., 2012; Vij & Walker, 2014). For a fuller overview of the development and applications of hybrid choice models, see Abou-Zeid and Ben-Akiva (2014).

The hype surrounding hybrid choice models has not come without concerns about their use. As pointed out by Vij & Walker (2016), many applications have oversold the empirical benefits of the model, and it is important to allow sufficient flexibility for the choice model component to also allow for heterogeneity not linked to the attitudinal constructs. This was a core consideration in our work, where we carefully attempted to study what part of the heterogeneity could in fact be linked to these attitudinal constructs rather than simply being misattributed heterogeneity that is driven by other factors.

In our work, we sampled respondents living in the larger metropolitan areas of Boston, New York, Philadelphia and Washington, DC (for NEC participants) or Portland, Seattle and Vancouver (for Cascade participants) and who made at least one intercity trip to other cities within their respective corridor. The sample for our study comprised roughly 5,500 respondents from the NEC recruited through an online sample and a previous study of auto users in the NEC, with just over 500 respondents obtained from an online sample for the Cascade corridor. We specified a flexible model with a large amount of deterministic and random heterogeneity, and also tested the actual impact of the latent constructs in a scenario testing context (see also Daziano & Bolduc, 2013). The findings show major differences across areas, trip types and traveller characteristics, where at least part of these differences can be linked to underlying attitudes. Interestingly from a policy perspective, we see that attitudes can differ from actual choices, and that different attitudes can have counter-acting effects. For example, West Coast travellers are less pro-car, but still have an increased utility for car in the choice model, potentially reflecting lower quality of service. Similarly, while the reduced concern for privacy for older and female travellers increases the appeal of public transport, this is counter-acted by a stronger pro-car attitude.

While the empirical data relates to a specific US setting, the findings should be of broader interest in that they highlight how differences in travel behaviour across regions may be attributed to factors going beyond the transport network itself.

The remainder of this paper is organised as follows. The following Section looks at the survey work conducted for this study. This is followed in Section 3 by an overview of the analysis of the attitudinal data, prior to the description of the hybrid choice model in Section 4. Model results are report in Section 5, before we turn to model application (Section 6) and finally present some conclusions (Section 7).

2. SURVEY WORK AND DATA PROCESSING

The survey incorporated a number of separate components, collecting background information on respondents' travel patterns, presenting them with a set of hypothetical choice scenarios, and finally collecting information on attitudes and socio-demographic characteristics.

At the start of the survey, respondents reported details on a recent trip for one of nine city pairs. For the NEC corridors, this concerned travel between Boston and either Philadelphia, New York City or Washington DC, between New York City and either Washington, DC or Philadelphia, and between

Philadelphia and Washington, DC. For the Cascade corridor, we looked at travel between any of the three combinations of Seattle, Portland and Vancouver.

The central part of the survey was formed by a set of stated choice (SC) scenarios. In each, respondents faced a choice between car (rental car for people without access to a car), bus, air and train, with the travel time and cost shown for each mode varying across the choice tasks. For car, cost was divided into parking cost, petrol costs and rental costs (if applicable). For the three non-car modes, travel time was divided into access time, on board time and egress time, while total time was also shown. Travel cost was presented as a total cost for the party as well as a per person cost. For the non-car options, no access cost or parking cost was included as the access mode choice was not specified. Both were shown for a one-way journey. The trip characteristics presented to respondents were pivoted around network levels for the specific origin and destinations indicated by the respondent, for all modes including the one chosen on the reference trip. Each respondent was faced with eight separate choice tasks (an example of SC scenario is shown in FIGURE 1).

Below are 4 different travel options for your 2 day trip from your home to Boston. Assume that **none of the options require a transfer or connection**.
If the options below are the only options available for your trip, which would you prefer?
 Highlighted information may have changed.

Option 1: Train 	Option 2: Personal Car 	Option 3: Air 	Option 4: Bus 
Time driving to station & time at station: 0 hr 15 min		Time driving to airport, check-in & security: 1 hr 30 min	Time driving to station & time at station: 0 hr 15 min
On-board travel time: 3 hr 30 min	Time in car: 3 hr 40 min	Time in plane: 1 hr 3 min	On-board travel time: 3 hr 22 min
Destination station to final destination: 0 hr 38 min		Airport to final destination: 0 hr 51 min	Destination station to final destination: 0 hr 30 min
Total Travel Time: 4 hr 23 min	Total Travel Time: 3 hr 40 min	Total Travel Time: 3 hr 24 min	Total Travel Time: 4 hr 7 min
	Parking fees for total trip: \$68.00		
	One-way gas costs: \$32.00		
One-way cost per person: \$112.00	Implied one-way cost per person (½ of parking fees + one-way gas costs): \$34.00	One-way cost per person: \$125.00	One-way cost per person: \$60.00
One-way cost for entire party of 2: \$224	Implied one-way cost for entire party of 2: \$68.00	One-way cost for entire party of 2: \$250	One-way cost for entire party of 2: \$120
I prefer this option <input type="radio"/>	I prefer this option <input type="radio"/>	I prefer this option <input type="radio"/>	I prefer this option <input type="radio"/>

Figure 1: Stated Preference Experiment from the Survey

The actual combinations for the time and cost attributes shown in a given choice scenario were determined by a Bayesian D-efficient design (Rose and Bliemer, 2014). Different designs were produced across purposes (business vs non-business), for three different journey lengths (i.e. different numbers of days away), nine corridors and two types of car availability, leading to a total of 108 designs produced for this study. In each design, five levels were used for each attribute, with the middle level being the reference value (from the network data), with two decreases and two increases around this value. Most changes were between 25% reductions and 25% increases, with wider ranges used for bus and train costs.

The design process chooses the combinations of attribute levels that lead to lowest D-error (of those combinations tested) for the specific prior values assumed for the model parameters. The priors used in the design process were based on an extensive review of values obtained in past studies and also differed between business and non-business travellers. To allow for uncertainty in the priors, we used a Bayesian D-efficient design, where we worked with wide ranges, using normally distributed priors, with standard deviations that were 50% of the mean values.

After completing the SC component of the survey, respondents were asked to state their level of agreement with a set of 47 separate statements aimed at testing underlying attitudes. The specific statements used were

based on an extensive review of the literature and prior testing, with full details available in NCRRP (2016). While the survey did seek to elicit shorter attitudes towards the decision to take the train, it also aimed to obtain a set of four longer term values and preferences, which were hypothesized to be relevant to possible future scenarios, including orientation towards the private car, attitude towards urbanism/communal behaviour, preferences for personal privacy, and need to be productive and connected during travel.

In the final component of the survey, respondents were asked to provide data on a number of key socio-demographic and economic characteristics that would later be used to help explain the heterogeneity in mode choice behaviour, with key examples being age, income, education and gender.

Since the current project was research driven, the research team set out to obtain a sample that provided good coverage and enough sample of various income, age, and home location distributions rather than strictly being population-proportional. Data were collected through several different sampling strategies. For the NEC survey sampling, we used the e-mail addresses of respondents from the NEC’s Auto OD Study who indicated that they would be willing to participate in future research. Afterwards, we used commercial sample providers and recruited additional respondents. Supplemental purchased online sample was targeted to counterbalance some of the demographic skew in the NEC Auto OD Study sample (e.g., the sample was older, more likely to own car). For the Cascade Corridor, no pre-existing sample existed and the entire sample was purchased from a commercial sample provider.

Data cleaning removed respondents who failed to answer specific components of the survey or selected the same answer to each attitudinal question. In total, less than 10% of the sample were excluded. For the present study, we also excluded respondents who had indicated multiple purposes for their trip. We retained a final sample of 5,413 respondents, of which 503 were from the Cascade corridor.

3. ATTITUDINAL CONSTRUCTS

As mentioned in Section 2, answers to a large number of attitudinal questions were collected as part of the survey. Exploratory factor analysis was used to understand the links between these individual statements. An extensive discussion of this work is presented in NCRRP (2016). For the analysis in the present paper, we focus on four factors which are linked to 11 of the statements. These are summarised in TABLE 1, split into East Coast (E) and West Coast (W).

Table 1: Attitudinal Indicators

Statement	F1		F2		F3		F4	
	E	W	E	W	E	W	E	W
<i>“I enjoy being out and about and observing people”</i>	0.65	0.65						
<i>“I like to live in a neighbourhood where I can walk to a commercial or village centre”</i>	0.54	0.65						
<i>“If everyone works together, we could improve the environment and future for the earth”</i>	0.56	0.56						
<i>“Rather than owning a car, I would prefer to borrow, share, or rent a car just for when I need it”</i>			0.56	0.6				
<i>“I love the freedom and independence I get from owning one or more cars”</i>			-0.58	-0.65				
<i>“I feel I am less dependent on cars than my parents are/were”</i>			0.74	0.64				
<i>“The idea of being on a train or a bus with people I do not know is uncomfortable”</i>					0.91	0.68		
<i>“I don't mind traveling with people I do not know”</i>					-0.47	-0.73		
<i>“The thought of sharing a car with others for such a trip seems unpleasant to me”</i>						0.41		
<i>“It would be important to me to receive e-mail or text message updates about my bus or train trip”</i>							0.71	0.51
<i>“Being able to freely perform tasks, including using a laptop, tablet, or smartphone is important to me”</i>							0.59	0.79

4. MODEL SPECIFICATION

In our analysis, we grouped together data from the entire sample, and then allowed for differences in behaviour across different sample subsets, such as by trip purpose and corridor. This section discusses the specification of individual components of the hybrid choice model used in our application. We start with the definition of the latent attitudes before turning to the choice model component. Finally, we look at the joint estimation of the different model components. After careful consideration, we decided against including a “simple” or “base” model alongside the hybrid structure. It is well known that a model which is estimated only on the choice data alone will fit that data at least as well as the choice component of a hybrid structure if the same flexibility is used (Vij & Walker, 2016). Little insight can thus be gained from such a comparison. We instead focus on attempting to gain insights into what share of the heterogeneity in the hybrid model can in fact be linked to the attitudinal constructs.

4.1. Latent attitudes: structural model

Drawing on the factor analysis work in Section 3, we specified four latent attitudes, represented by latent variables α_l , with $l=1, \dots, L$, where $L=4$. These are hereafter referred to as:

- LV1: attitude toward sociability;
- LV2: attitude toward cars;
- LV3: attitude toward privacy; and
- LV4: attitude toward (information) technology.

After extensive specification testing, five person-level characteristics were used in the structural equations for the latent attitudes. These were gender (male used as the base), age (split into four categories, where 35-44 is the base), education (using those with a degree as the base and estimating an offset for those without a degree), employment status (using those in employment as the base and estimating an offset for those not in employment) and finally geography (using the NEC corridor as the base and estimating an offset for the West Coast). As we will see in Section 5, not each of these effects remains significant for every latent variable). Each of the latent attitudes is defined to have a deterministic and a random component, with latent attitude l for person n defined as:

$$\begin{aligned}\alpha_{n,l} = & \gamma_{l,female} \cdot Z_{n,female} \\ & + \gamma_{l,age\ under\ 35} \cdot Z_{n,age\ under\ 35} \\ & + \gamma_{l,age\ 45\ to\ 54} \cdot Z_{n,age\ 45\ to\ 54} \\ & + \gamma_{l,age\ 55\ to\ 64} \cdot Z_{n,age\ 55\ to\ 64} \\ & + \gamma_{l,age\ over\ 65} \cdot Z_{n,age\ over\ 65} \\ & + \gamma_{l,no\ graduate} \cdot Z_{n,no\ graduate} \\ & + \gamma_{l,no\ job} \cdot Z_{n,no\ job} \\ & + \gamma_{l,west\ coast} \cdot Z_{n,west\ coast} \\ & + \xi_{n,l}\end{aligned}$$

[1]

In Equation [1], $\xi_{n,l}$ is a standard Normal variate (mean of 0, standard deviation of 1), distributed across respondents, capturing the random element of the latent attitude.

4.2. Latent attitudes: measurement model

We have four latent attitudes in our model, and these are used in the measurement model component of our overall framework to explain the responses to the attitudinal indicators in Table 1.

All 11 questions use a 7-level Likert scale. With I_s used to refer to a given attitudinal question, and letting α_l be the associated latent attitude, we use an ordered logit model to explain the likelihood of the actual observed value of $I_{n,s}$ for respondent n as:

$$LI_{n,s} = \sum_{p=1}^7 x_{I_{n,s},p} \left(\frac{e^{t_{I_s,p} - \zeta_{l,s} \alpha_{n,l}}}{1 + e^{t_{I_s,p} - \zeta_{l,s} \alpha_{n,l}}} - \frac{e^{t_{I_s,p-1} - \zeta_{l,s} \alpha_{n,l}}}{1 + e^{t_{I_s,p-1} - \zeta_{l,s} \alpha_{n,l}}} \right) \quad [2]$$

where $x_{I_{n,s},p}=1$ if and only if respondent n chooses answer p for question s . The $t_{I_s,p}$ parameters are thresholds that are to be estimated, with the normalisation that $t_{I_s,0} = -\infty$ and $t_{I_s,7} = +\infty$. The estimated parameter $\zeta_{l,s}$ measure the impact of the latent variable α_l on I_s , where a significant estimate for $\zeta_{l,s}$ shows us that the latent attitude α_l has a statistically significant impact on the answers provided to the attitudinal question I_s .

4.3. Choice model component

In the choice model component, we explain the choice between the four modes of transport, i.e. car, bus, air and rail. The utility for mode i for person n in choice situation t is given by:

$$U_{n,i,t} = \delta_{n,i} + \tau_i \alpha_n + \beta_{n,i} x_{n,i,t} + \varepsilon_{n,i}, \quad [3]$$

where $\varepsilon_{n,i}$ is a type I extreme value error term, distributed identically and independently across alternatives and observations. This means that the choice model component of the overall model takes a Mixed Logit form. It is of course conceivable that there exist correlations between the individual public transport modes, for example, and a nested or error components structure would be an interesting area for future work.

We will now look separately at the four components of Equation [3], with a particular focus on the treatment of heterogeneity.

4.3.1 Mode specific constants

For the mode specific constants, we write:

$$\delta_{n,i} = \mu_{\delta_i} + \sigma_{\delta_i} \xi_{n,\delta_i} + \sigma_{\delta_i,2} \xi_{n,\delta_i}^2 + \lambda_i z_n + \varpi_i q_n \quad [4]$$

In this specification, we allow for deterministic and random heterogeneity in the utilities for different modes, where, for the random heterogeneity, we move away from purely parametric distributions by relying on the Fosgerau & Mabit (2013) approach.

In particular, we have that μ_{δ_i} is the estimated mean for the alternative specific constant for mode i . With ξ_{n,δ_i} being a standard Normal random variable, distributed independently across respondents, the addition of $\sigma_{\delta_i} \xi_{n,\delta_i}$ would imply that the alternative specific constant now follows a Normal distribution with a mean of μ_{δ_i} and a standard deviation of σ_{δ_i} . To move away from the restrictive shape of that distribution, we include an additional polynomial term $\sigma_{\delta_i,2} \xi_{n,\delta_i}^2$, where $\sigma_{\delta_i,2}$ is an estimated parameter that multiplies the square of the same standard Normal random variate used in $\sigma_{\delta_i} \xi_{n,\delta_i}$. If $\sigma_{\delta_i,2}$ tends to zero, then we revert to a Normal distribution. If $\sigma_{\delta_i,2}$ is positive, then we get a positive skew in the distribution, with the opposite applying with a negative $\sigma_{\delta_i,2}$. For normalisation, we set μ_{δ_i} , σ_{δ_i} and $\sigma_{\delta_i,2}$ to zero for bus, which was the

mode with lowest random heterogeneity in an overspecified model used for testing (the normalisation is not arbitrary in a mixed logit model).

Turning to the deterministic component of heterogeneity, we have that z_n presents a vector of respondent characteristics and q_n a vector of trip characteristics, with λ_i and ϖ_i measuring the impact of these two vectors on the value of the mode specific constants.

The characteristics included in the vector z_n for interactions with mode specific constants included:

- gender;
- age, using the same specification as in Section 4.1;
- education, using the same specification as in Section 4.1;
- employment, using the same specification as in Section 4.1;
- households with fewer cars than adults;
- households with more cars than licenses; and
- West Coast dummy.

It is worth acknowledging that a further increase in flexibility would be possible by making the socio-demographic effects area-specific, e.g. allowing for a different impact of gender for East Coast and West Coast travellers, and that our specification potentially means that the socio-demographic findings are primarily driven by East Coast respondents given the larger sample size for that group².

The parameters associated with these respondent characteristics measure the shift in the mode specific constants compared to a respondent with the base value for these characteristics, where once again bus was used as the base.

The vector q_n of trip specific characteristics included:

- journey purpose: with VFR as the base, shifts in the non-bus constants were tested for the four remaining purposes;
- party size: with single person as the base, shifts in the non-bus constants were tested for groups of two and groups of three or more;
- trip length in terms of overnight stays: with same day return as the base, shifts in the non-bus constants were tested for single overnight trips, two overnights and three or more nights away; and
- frequency: entered as the logarithm of daily service frequency and used in the constants for non-car modes. Note that this was included here as opposed to being listed as an explanatory variable below as it was not included as a variable in the survey, i.e. it was not explicitly shown to respondents.

The parameters associated with these trip characteristics again measure the shift in the mode specific constants compared to a trip with the base value for these characteristics. Bus was always used as the base except for frequency of service, given that this is an attribute that applies to the three non-car modes, but with different values.

² To test this possibility, some additional model runs were carried out which showed that out of 39 socio-demographic effects, a statistically significant difference could only be obtained between East Coast and West Coast respondents for five parameters. These related to a less strong preference for car when travelling in a group for West Coast respondents and a stronger air preference for female respondents. However, the effects themselves were weak and led to stability issues with other more behaviourally intuitive effects in the model.

For both traveller and trip characteristics, heterogeneity was only considered in the mean value of the mode specific constants, i.e. no additional impact on the random heterogeneity was incorporated (e.g. more variance for a given purpose).

4.3.2 Impact of latent attitudes

The second component in Equation [3], i.e. $\tau_i \alpha_n$, concerns the impact of the latent variables on the mode specific constants. Specifically, τ_i is a vector of parameters explaining the impact of the four latent attitudes on the utility of mode i , such that:

$$\tau_i \alpha_n = \sum_{l=1}^4 \tau_{i,l} \alpha_{n,l}, \quad [5]$$

where $\tau_{i,l}$ measures the impact of the latent variable l on the constant for mode i , where we use a normalisation setting $\tau_{i,l}$ to zero for $i=2$ (bus) and for all l .

Two points need mentioning here.

Firstly, in the present work, we consider the effect of the four LVs only on the ASCs, rather than also testing for an impact on the marginal sensitivities to level of service variables. This is an obvious simplification, but comes in the context of a model that is already highly complex to estimate and also in a study where the main interest is in the impact of attitudes on pure modal preferences. Nevertheless, we acknowledge that this potentially reduces the impact of the LVs on the model.

Secondly, the LVs follow a Normal distribution while the ASCs themselves are given additional distributional flexibility through the use of the polynomial transforms in Equation [4]. We acknowledge that this is a restriction and potentially reduces the share of the heterogeneity in the modal preferences that can be attributed to the LVs. Once again, this is a pragmatic choice in the context of an already complex model where the use of polynomial error terms for the LVs would also have complicated the overall normalisation of the model.

4.3.3 Explanatory variables

We finally look at the components of utility related to explanatory variables, i.e. $\beta_{n,i} x_{n,i,t}$ in Equation [3]. The attributes included in $x_{n,i,t}$ varied across modes but always included travel time (divided into access time, in vehicle time, and egress time), and travel cost. For car, access time and egress time were obviously set to zero.

We again allowed for extensive deterministic and random heterogeneity in the coefficients used to reflect the marginal utilities of these explanatory variables. With both time and cost being undesirable attributes, we used purely negative distributions. Mode specific time coefficients were used, along with generic cost and access time coefficients. We again made use the Fosgerau & Mabit (2013) polynomial specifications, and specified the value of the coefficient for attribute k for respondent n written as:

$$\begin{aligned} \beta_{n,k} = & -\exp[\mu_{\log(-\beta_{n,k}),work} \cdot q_{n,work} + \mu_{\log(-\beta_{n,k}),non-work} \cdot q_{n,non-work} \\ & + (\sigma_{\log(-\beta_{n,k}),work} \cdot q_{n,work} + \sigma_{\log(-\beta_{n,k}),non-work} \cdot q_{n,non-work}) \xi_{\beta_{n,k}} \\ & + (\sigma_{\log(-\beta_{n,k}),work,2} \cdot q_{n,work} + \sigma_{\log(-\beta_{n,k}),non-work,2} \cdot q_{n,non-work}) \xi_{\beta_{n,k}}^2 \\ & + \Delta_{\log(-\beta_{n,k}),vacation} q_{n,vacation} + \Delta_{\log(-\beta_{n,k}),mixed\ leisure} q_{n,mixed\ leisure} \\ & + \Delta_{\log(-\beta_{n,k}),other} q_{n,other} \\ & + \Delta_{\log(-\beta_{n,k}),west\ coast} z_{n,west\ coast}] \end{aligned} \quad [6]$$

This rather involved specification requires some additional explanations. The first line includes separate means for the log of the negative of the coefficient (remembering that we are using a negative exponential)

for respondents on work trips (where $q_{n,work}$ is equal to 1) and respondents on non-work trips, estimated as $\mu_{\log(-\beta_{n,k}),work}$ and $\mu_{\log(-\beta_{n,k}),non-work}$, respectively. Additional shifts in these means are incorporated for the three non-VFR leisure purposes (i.e. $\Delta_{\log(-\beta_{n,k}),vacation}$, $\Delta_{\log(-\beta_{n,k}),mixed\ leisure}$ and $\Delta_{\log(-\beta_{n,k}),other}$), meaning that $\mu_{\log(-\beta_{n,k}),non-work}$ relates to a VFR trip. Finally, a shift is also allowed for West coast respondents via $\Delta_{\log(-\beta_{n,k}),west\ coast}$.

Random heterogeneity is accommodated by two polynomial terms, multiplying $\sigma_{\log(-\beta_{n,k}),work}$ (for work trips) and $\sigma_{\log(-\beta_{n,k}),non-work}$ (for non-work trips) by a standard Normal variable $\xi_{\beta_{n,k}}$, and $\sigma_{\log(-\beta_{n,k}),work,2}$ (for work trips) and $\sigma_{\log(-\beta_{n,k}),non-work,2}$ (for non-work trips) by the square of that standard Normal, i.e. $\xi_{\beta_{n,k}}^2$. Noting that inside the negative exponential in Equation [6], we are working with Normal distributions, a negative value for the second polynomial term (i.e. $\sigma_{\log(-\beta_{n,k}),work,2}$ or $\sigma_{\log(-\beta_{n,k}),non-work,2}$) would lead to a negative skew in the Normal distribution which would in turn mean a less fat tail for the distribution after taking the exponential. This is a very useful way of allowing the data to push the model away from the very fat tail of a standard Lognormal distribution.

As with the mode specific constants, not every covariate had a significant impact on every coefficient, as discussed later, and for cost, the second polynomial term also dropped out.

For the different time attributes, the contribution to the utility function is simply given by $\beta \cdot x$, i.e. a linear in attributes specification. For costs, a different approach was used. As shown in Section 2, costs were presented in the survey as total party cost. It is however clear that a respondent does not necessarily cover all the cost himself/herself. After extensive testing, we used a specification which, for air, bus and rail, used the per person cost in our models. For car, better performance was obtained by a specification which recognized that the driver often pays a larger share, and thus multiplied the total cost by $\frac{1}{1+\log(\text{party size})}$. This thus uses the presented cost if the party size is 1, but gradually dampens this with larger party sizes, but not in a linear way (e.g. the respondent pays 59% in the case of two travellers, instead of 50%, etc). In addition, we allowed for an income effect on the cost sensitivity, where we directly estimated the income elasticity. Using car as the example, the contribution of cost to the utility function in choice situation t would then be written as:

$$\beta_{n,cost} \cdot \frac{x_{n,cost,car,t}}{1+\log(\text{party size})} \cdot \left(\frac{inc_n}{\overline{inc}}\right)^{\lambda_{inc}}, \quad [7]$$

while for non-car modes, it would

$$\beta_{n,cost} \cdot \frac{x_{n,cost,j,t}}{\text{party size}} \cdot \left(\frac{inc_n}{\overline{inc}}\right)^{\lambda_{inc}}, \quad [8]$$

with j being air, rail or bus.

With this specification (Mackie et al., 2003), λ_{inc} is an estimated income elasticity on the cost sensitivity, where inc_n is the income of respondent n and \overline{inc} is the average income in the sample. Respondents with missing income were assigned the sample mean income after no significant differences in sensitivities were observed for them.

4.4. Joint model estimation

The combined utility specification now includes:

- the impacts of the explanatory variables, with randomly distributed time and cost coefficients, which vary across purpose and corridor, and where the cost coefficient is also interacted with income;

- the mode specific constants, which include a deterministic component as well as a random part; and
- an impact on the modal constants by the latent attitudes, which again include a deterministic and random component.

Two important points need to be made here.

Firstly, the deterministic heterogeneity terms included in the modal constants explained above relate to person as well as trip characteristics, while those terms mentioned earlier for the latent attitudes related only to person characteristics. This reflects the assumption that attitudes are stable for each person across different trips.

Secondly, all respondent characteristics included in the deterministic component of the latent attitude have also been included directly in the modal constant, thus avoiding a situation where a sociodemographic effect is erroneously captured as relating to attitudes when it may just relate to underlying modal preferences, or vice versa. As an example, it may well be the case that younger respondents travel less by car for reasons unrelated to their attitude toward cars. If age was included as a covariate only on the latent attitude toward cars but not separately on the modal constants, this inherent modal preference may erroneously be captured as an attitudinal difference. In very much the same way, the modal constants now include a random component that relates to the latent attitudes (through the inclusion of $\tau_i \alpha_n$ in Equation [3]) while separate random components ($\sigma_{\delta_i} \xi_{n,\delta_i} + \sigma_{\delta_i,2} \xi_{n,\delta_i}^2$ in Equation [4]) relate to random variations in preferences for modes which cannot be linked to latent attitudes, for example due to uncaptured journey-specific effects, where we again acknowledge that the treatment for the ASC-specific heterogeneity is more flexible than for that linked to the LVs.

The above discussion brings us to an additional important point. The specification of an alternative specific constant now includes two separate random terms, both with deterministic interactions on the mean, some of which relate to the same underlying sociodemographic variables. In a standard choice model, this would be an over specification, with two parameters capturing the same effect. What allows us to separately identify the two components is that one of them, namely the latent variable component, is also used in a separate measurement model.

With $i_{n,t}$ being the alternative chosen by respondent n in task t (out of $T=8$), we have that the likelihood of the observed choices and answers to attitudinal questions for respondent n is given by:

$$L_n = \int_{\alpha} \int_{\beta} \int_{\delta} \prod_{t=1}^T \frac{e^{v_{i_{n,t}}}}{\sum_{j=1}^4 e^{v_{j,t}}} \prod_{s=1}^{11} LI_{s,n} f(\alpha) f(\beta) f(\delta) d\delta d\beta d\alpha \quad [9]$$

where we use a Logit kernel for the choice model component, and where $LI_{s,n}$ is defined as above as an ordered logit model. Both the component relating to the choices (i.e. the Logit kernel) and the component relating to the attitudinal questions are a function of the vector of latent variables α , while the choice model component is also a function of the random components used in the marginal utility coefficients (β) and the random components used in the alternative specific constants (δ). This is why the entire likelihood function is integrated over the distribution of α , β and δ . This integration is carried out at the level of an individual respondent, i.e. recognising the repeated choice nature of the data (in terms of the 8 choice tasks) as well as the simultaneous collection of data on 11 attitudinal indicators. This thus allows for correlation across these 19 observed outcomes for a given individual.

The resulting model structure is rather complex, and is estimated on a large sample, with 43,304 observed choices from 5,413 respondents, while at the same time explaining the answers to 59,543 attitudinal questions (11 per respondent). Our model incorporates four latent variables as well as three randomly distributed alternative specific constants and seven randomly distributed marginal utility coefficients, where these were split between work and non-work, leading to 21 random components, with a full covariance

matrix. For the constants and marginal utility coefficients, we additionally moved away from standard parametric distributions. A total of 224 parameters were estimated in the final specification, which includes the large number of non-random terms. The resulting flexibility meant that classical estimation of the model became computationally intractable and we instead relied on Bayesian estimation, using the implementation of Hierarchical Bayes (see Train, 2001 for a comparison with classical estimation) in RSGHB (Dumont & Keller, 2015). While classical estimation is impractical for a model of this size, Bayesian estimation relies on priors which could also influence the results. We used the RSGHB default assumption of a $N(0, \Omega)$ multivariate prior for the random coefficients, where Ω is a diagonal matrix with zero off-diagonals and diagonals set to 2. The zero means for the prior would only push the means of the random parameters towards zero; this however depends on the degree of variance in relation to the size of the estimated parameters. With the overall small values for the random terms, the impact of the assumptions about Ω should be minimal.

5. MODEL RESULTS

The modelling effort undertaken for this work was substantial, and the results are very detailed and are presented across a number of different tables³. We will now look at the different parts of the results in turn.

5.1. Measurement model for latent variables

TABLE 2 presents the results for the measurement model component, using the grouping and ordering already used in the factor analysis results in TABLE 1. We show the means of the posteriors, which in classical terms are equivalent to maximum likelihood estimates, and standard deviations of the posteriors, which in classical terms are like classical standard errors. We will focus our discussion on the ζ parameters as the threshold parameters simply reflect the ordered level of the indicators and the nature of the model.

Looking at the statements used as dependent variables for the first latent variable (LV1), the consistent negative signs for all three statements indicate that a higher value for the first latent variable leads to a lower level of agreement with the three statements. This means that respondents with a higher value for LV1 are less sociable. For the second latent variable, we see the expected opposite sign for the second statement, where the overall pattern of signs for the three ζ parameters means that respondents with a higher value for LV2 have a more positive attitude towards cars than other respondents. For the third latent variable, we again see the expected opposite sign for the second statement, where the overall pattern of signs for the three ζ parameters means that respondents with a higher value for LV3 are more concerned about privacy than other respondents. Finally, for the fourth latent variable, the positive signs for both ζ parameters mean that respondents with a higher value for LV4 are more pro technology than other respondents.

5.2. Structural model for latent variables

TABLE 3 presents the results for the structural model for the four latent variables. These need to be interpreted alongside the sign of the ζ parameters in TABLE 2. For age, no difference was observed between the 45-54 group and the base age group of 35-44, so that for the structural equations, the base age group used was in effect 35-44. We see that female respondents and those aged under 35 have a less positive value for LV1, meaning that they are more sociable. On the other hand, the opposite applies for respondents who are less educated and those not in employment. West coast respondents also have a higher value for LV1, i.e. are less sociable, but this effect is less important than for other characteristics, especially gender.

³ The final log-likelihood of the model is rather unimportant in the absence of model comparisons. We obtain a log-likelihood for the overall model of -124,073.5, with -24,082.63 for the choice model component alone (equating to a ρ^2 for the choice model of 0.59).

Table 2: Results for Measurement Model

Statement	associated LV	ζ		t_1		t_2		t_3		t_4		t_5		t_6	
		post μ	post σ												
I enjoy being out and about and observing people	LV1: low sociability	-1.659	0.044	-5.524	0.060	-4.498	0.061	-3.450	0.050	-1.987	0.025	-0.375	0.013	1.921	0.048
I like to live in a neighbourhood where I can walk to a commercial or village centre		-1.458	0.031	-4.326	0.074	-3.227	0.056	-2.490	0.042	-1.400	0.030	-0.212	0.005	1.559	0.026
If everyone works together, we could improve the environment and future for the earth		-1.335	0.051	-4.694	0.092	-4.126	0.054	-3.615	0.066	-2.487	0.059	-1.101	0.025	0.648	0.023
Rather than owning a car, I would prefer to borrow, share, or rent a car just for when I need it	LV2: pro-car	-1.358	0.029	-2.282	0.024	-0.724	0.014	0.109	0.002	1.074	0.015	1.763	0.022	2.909	0.042
I love the freedom and independence I get from owning one or more cars		1.554	0.051	-4.520	0.073	-3.678	0.059	-3.143	0.041	-2.095	0.025	-0.885	0.017	1.029	0.020
I feel I am less dependent on cars than my parents are/were		-2.256	0.066	-1.348	0.037	0.615	0.023	1.636	0.026	2.590	0.038	3.662	0.060	5.058	0.082
The idea of being on a train or a bus with people I do not know is uncomfortable	LV3: concerned about privacy	2.744	0.048	-3.505	0.035	-0.446	0.006	0.931	0.014	2.591	0.051	4.377	0.047	6.132	0.086
I don't mind traveling with people I do not know		-1.411	0.025	-4.013	0.071	-2.668	0.045	-1.592	0.025	-0.459	0.016	0.747	0.021	3.082	0.041
The thought of sharing a car with others for such a trip seems unpleasant to me		0.626	0.010	-4.721	0.065	-3.508	0.035	-2.871	0.034	-1.437	0.037	0.089	0.003	2.143	0.033
It would be important to me to receive email or text message updates about my bus or train trip	LV4: pro-tech	1.515	0.026	-6.308	0.134	-4.599	0.067	-3.609	0.062	-1.903	0.039	-0.099	0.003	2.238	0.070
Being able to freely perform tasks, including using a laptop, tablet, or smartphone is important to me		2.099	0.065	-2.561	0.040	-1.088	0.023	-0.382	0.010	0.546	0.022	1.420	0.038	2.774	0.056

For the second latent variable, we see that female respondents are more pro-car, as are older respondents (with a monotonic trend across age groups, with no difference between the base group and the 45-54 group), where the age effect is stronger than the gender effect. The same applies to respondents who are less educated. On the other hand, those not in employment are less pro-car, as are West coast respondents compared to East coast respondents.

Turning to the concern for privacy (LV3), we see that female respondents are less concerned about privacy, as are older respondents, with a marked difference for the oldest age group. On the other hand, concern for privacy is higher for less educated respondents, those not in employment, and West coast respondents.

Finally, the results for the pro-tech latent variable (LV4) indicate that women and younger respondents are more pro-tech, while less educated respondents, those not in employment, and West coast respondents are less pro-tech. Again, age shows the strongest effect.

Before moving on, it is worth briefly discussing the overall picture and see how these results compare with expectations. The gender effects are potentially surprising for LV2 and LV4, showing women to be more pro-car and pro-tech. This could suggest that the actual choices made (with women being less heavy users of cars and technology) are influenced by circumstances rather than desire. The findings for age are largely in line with expectations, showing higher sociability for younger people, and a less pro-car attitude, along with increased concerns for privacy and a greater interest in technology. Reduced education and employment leads to reduced sociability but also reduced concern for privacy. The reduced pro-car and pro-tech attitudes for those not in employment could be related to financial constraints. Finally, for West Coast vs East Coast respondents, the reduced pro-car attitude for the former could again suggest a disconnect between desire (not to use car) and actual choices (heavy reliance on car) as a result of circumstances (transport network).

Table 3: Results for Structural Equations Model

	γ	LV1: low sociability		LV2: pro-car		LV3: concerned about privacy		LV4: pro-tech	
		post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ
Female (vs male)		-0.288	0.008	0.081	0.002	-0.072	0.001	0.190	0.008
Aged under 35 (vs 35-54)		-0.097	0.004	-0.395	0.013			0.289	0.003
Aged 55-64 (vs 35-54)				0.163	0.010	-0.175	0.003	-0.178	0.005
Aged 65 and over (vs 35-54)				0.307	0.011	-0.425	0.009	-0.392	0.021
No graduate (vs graduate)		0.141	0.006	0.144	0.006	0.295	0.007	-0.133	0.005
Not employed (vs employed)		0.091	0.004	-0.114	0.004	0.203	0.005	-0.122	0.001
West coast (vs East)		0.028	0.001	-0.111	0.002	0.270	0.008	-0.129	0.006

5.3. Mode specific constants in choice model

For the mode specific constants, we first look at the impact of the latent variables in the utility functions, i.e. the $\tau_i \alpha_n$ component in Equation [3]. The estimates for the individual τ parameters are reported in TABLE 4. For the first latent variable (reduced sociability), we see that, with bus as the base, respondents with a higher value for this latent variable have a greater utility for car and (less so) air, with no significant impact on rail. The pro-car latent variable (LV2) shows the expected positive impact on the utility for car (which is the strongest impact of any of the latent variables) with a smaller positive impact on the utility for air. Turning to the reduced concern for privacy attitude (LV3), we see a reduction in utility for car and (less so) air for those respondents with a higher value for this latent attitude compared to bus, with a smaller increase in the utility for rail. Finally, respondents who are more pro-tech (LV4) have a reduced utility for car compared to all other modes. Overall, these findings make intuitive sense, with car and air being seen as more private modes, and driving a car being less conducive to using mobile technology.

Table 4: Impact of Latent Variables on Mode Constants

τ	car		air		rail	
	post μ	post σ	post μ	post σ	post μ	post σ
LV1: low sociability	0.100	0.003	0.055	0.002		
LV2: pro-car	1.194	0.025	0.355	0.006		
LV3: concerned about privacy	0.760	0.017	0.566	0.011	-0.235	0.004
LV4: pro-tech	-0.374	0.005				

TABLE 5 shows the part of the utility function related to the mode specific constants net of attitudinal impacts, i.e. the individual components of Equation [4]. With bus as the base, we first see negative mean values for the other three mode specific constants, along with sizeable standard deviations (σ), indicating variations across respondents in their baseline preferences for the different modes. However, the positive signs and large values for the additional polynomial term (σ_2) which multiplies the square of the standard Normal variate, means that the distributions are all positively skewed, showing a large share of respondents with a baseline preference for non-bus modes.

The baseline preferences also change as a function of trip and traveller characteristics, which we now look at in turn. With VFR as the base purpose, we see a positive shift in the utility for car and rail for work trips compared to bus, with a smaller reduction in the baseline utility for air. For vacation, we see a shift away from bus in the baseline preferences for all modes, least so for car, while for mixed leisure, this only benefits air and marginally rail travel. Finally, for other purposes, there is a reduction in the baseline preference for air and rail. For city pairs with higher frequency of air service, there is an increase in the baseline preference for air, reflecting improved scheduling and seat availability. The same applies for bus, but with a much smaller effect, possibly reflecting less concern about seat availability. Travelling with additional people increases the preference for all non-bus modes, especially car and rail. However, when travelling with two or more people, there is a reduction in the baseline utility of non-bus modes, especially for air and rail. As the number of nights the traveller stays away from home increases, the utility of non-bus modes increases compared to bus, where this effect is uniform across non-bus modes for two nights and also for three of more nights, while, for overnight trips, the effect is strongest for rail, ahead of car and air.

We next turn to respondent characteristics. We see that, compared to East Coast travellers, West Coast respondents have a substantially increased utility for car and rail compared to bus. This finding is directly contrary to the findings in the structural equations, a point we return to below. Female respondents have an increased utility for air compared to bus, but a reduced utility for car and rail (same effect). For respondents aged under 35, the utility is reduced for all modes compared to bus, where the effect is stronger (and the same) for car and rail. For the two older age groups, a less clear picture emerges and some effects are constrained to be equal across groups after earlier estimations found no differences. We see an increased (equal) utility for all non-bus modes for the 45-54 group. In the 55-64 age group, we also see an increase in the utility for all non-bus modes, but this is stronger for rail. Lastly, in the highest age group, small reductions are observed for air and rail compared to bus, but an increased utility for car. Respondents without a degree have a reduced utility for all non-bus modes, especially for rail and car, while for those not in employment, this is also the case but the reduction is strongest for air, ahead of car and rail, potentially linked to financial constraints. Finally, the number of cars in a household has the expected effect; respondents from households with fewer cars than adults have a reduced utility for car compared to all other modes, while those from households with more than one vehicle per license holder have an increased utility for car.


Table 5: Estimation Results for Mode Specific Constants

	car		air		rail		bus		
	post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ	
Distribution parameters	μ	-0.544	0.013	-2.106	0.027	-0.123	0.003		
	σ	0.551	0.012	0.162	0.005	0.931	0.034		
	σ_2	2.557	0.054	1.224	0.017	0.966	0.020		
	work (vs VFR)	0.503	0.012	-0.207	0.007	1.009	0.029		
	vacation (vs VFR)	0.175	0.006	0.513	0.008	0.383	0.009		
	mixed leisure (vs VFR)			0.635	0.013	0.076	0.006		
	other (vs VFR)			-0.325	0.020	-0.376	0.007		
Trip characteristics (ϖ)	effect of log freq			0.664	0.020			0.105	0.002
	one other person (vs alone)	0.397	0.006	0.195	0.007	0.330	0.008		
	two or more other people (vs alone)	-0.167	0.002	-0.673	0.016	-0.426	0.006		
	one night (vs day return)	0.200	0.005	0.076	0.005	0.361	0.007		
	two nights (vs day return)	0.200	0.005	0.200	0.005	0.200	0.005		
	3 plus nights (vs day return)	0.407	0.011	0.407	0.011	0.407	0.011		
	West coast (vs East)	0.895	0.021	0.036	0.021	0.801	0.020		
	female (vs male)	-0.181	0.004	0.311	0.023	-0.181	0.004		
	aged under 35 (vs 35-44)	-0.278	0.010	-0.135	0.011	-0.278	0.010		
	aged 45-54 (vs 35-44)	0.463	0.022	0.463	0.022	0.463	0.022		
Person characteristics (λ)	aged 55-64 (vs 35-44)	0.148	0.004	0.148	0.004	0.517	0.010		
	aged 65 and over (vs 35-44)	0.148	0.004	-0.065	0.005	-0.013	0.005		
	no graduate (vs graduate)	-0.788	0.012	-0.381	0.009	-1.233	0.016		
	not employed (vs employed)	-0.405	0.012	-0.504	0.011	-0.212	0.014		
	fewer cars than adults	-0.161	0.005						
	vehicles per license greater than 1	0.055	0.001						


As an illustration of the role of heterogeneity in the mode specific constants, FIGURE 2 shows the cumulative distribution function (CDF) of the constant for car in a number of stages. The first panel shows the impact that the polynomial terms have on the shape of the distribution, with the positive value for σ_2 leading to a positive skew in the distribution and a cutting off of the left tail. The contrast with a corresponding Normal distribution (with the same mean and standard deviation) is clear to see. In the second panel, we add in the socio-demographic and trip characteristics, i.e. λ_i and ϖ_i . This again leads to a clear change in the shape of the distribution. On the other hand, the inclusion of the impacts of the latent attitudes in the third panel only makes a small difference to the shape of the distribution, a point we will return to later in the paper.

5.4. Level of service variables in choice model


TABLE 6 reports the estimates of the parameters for the distribution of the marginal utility parameters. As shown in Equation [6], these relate to the parameters of the log of the negative value of the coefficient. We estimate separate baseline distribution parameters (μ , σ and σ_2) for work and non-work trips, with additional shifts in the means for the different non-work purposes as well as for West Coast respondents. The implications of the estimates are best understood in the context of the monetary valuations in Section 5.5, but for now, a number of observations can already be made.



a) Distribution with polynomials vs Normal with same mean and standard deviation



b) Distribution with polynomials vs distribution including person and trip characteristics



c) Distribution with person and trip characteristics vs full distribution including LV effects

Figure 2: Distribution of ASC for Air

Table 6: Estimation Results for Level of Service Attributes

Base distribution
parameters

	car travel time		bus travel time		air travel time		rail travel time		access time		egress time		cost	
	post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ	post μ	post σ
$\mu_{\log(-\beta_{k,n}),work}$	-4.118	0.050	-3.803	0.047	-5.941	0.241	-4.200	0.026	-4.080	0.081	-4.136	0.071	-3.384	0.037
$\sigma_{\log(-\beta_{k,n}),work}$	0.480	0.009	1.028	0.017	0.664	0.044	0.266	0.004	0.478	0.016	-0.381	0.020	0.828	0.028
$\sigma_{\log(-\beta_{k,n}),work,2}$	0.139	0.007	0.310	0.008	0.552	0.014	0.291	0.010	0.060	0.005	0.190	0.004		
$\mu_{\log(-\beta_{k,n}),non-work}$	-4.670	0.035	-3.711	0.026	-5.061	0.063	-4.402	0.025	-4.123	0.044	-4.232	0.113	-2.887	0.021
$\sigma_{\log(-\beta_{k,n}),non-work}$	0.193	0.008	0.686	0.026	1.039	0.035	0.298	0.008	0.491	0.016	0.774	0.009	0.788	0.025
$\sigma_{\log(-\beta_{k,n}),non-work,2}$	0.387	0.010	0.100	0.003	0.497	0.018	0.332	0.006	0.259	0.009	0.206	0.008		
shift in mean for vacation ($\Delta_{\log(-\beta_{k,n}),vacation}$)	0.118	0.002	-0.218	0.002	-0.114	0.003	-0.140	0.002					-0.163	0.008
shift in mean for mixed leisure ($\Delta_{\log(-\beta_{k,n}),mixed\ leisure}$)	-0.148	0.010	-0.350	0.016	0.227	0.009	-0.215	0.007	-0.154	0.004	-0.449	0.009	-0.243	0.006
shift in mean for other ($\Delta_{\log(-\beta_{k,n}),other}$)	0.239	0.013			-0.073	0.003			-0.163	0.004	-0.192	0.005		
shift in mean for West Coast ($\Delta_{\log(-\beta_{k,n}),west\ coast}$)	-0.466	0.013	0.078	0.002	0.270	0.011			-0.283	0.007	-0.087	0.001	0.133	0.007
income elasticity (λ_{inc})													-0.152	0.004

The estimation process showed that no additional gains could be made from including a second polynomial term (i.e. σ_2) for the cost coefficient while the same is not the case for all the other marginal utility coefficients, for both work and non-work. We see a number of impacts of trip purpose too. A negative shift in the mean for the underlying Normal will imply a less negative sensitivity (after taking the negative exponential) for the resulting marginal utility coefficient. This shows for example reduced cost sensitivity (compared to VFR) for vacation and mixed leisure trips, with increased sensitivity to car travel time for *other* purposes. Differences also arise for West Coast respondents, with reduced sensitivity to car travel time (possibly linked to lower levels of congestion and hence a more enjoyable driving experience), access time and egress time, but increased sensitivity to bus travel time, air travel time and travel cost. Finally, we see a negative income elasticity on cost, showing that for a 10% increase in income, we would observe a 1.5% drop in cost sensitivity.

5.5. Implied monetary valuations

TABLE 7 reports implied monetary valuations for the different travel time components as well as willingness to pay measures for avoiding bus travel. We report only the mean valuations here.

Table 7: Implied Monetary Valuations

SAMPLE LEVEL MEASURES (\$/hr)							
	car time	air time	rail time	bus time	access time	egress time	
all	28.07	32.31	30.85		56.74	35.85	36.49
work	45.44	21.70	48.49		101.27	44.21	47.70
vacation	26.87	34.75	26.80		45.00	37.24	38.26
mixed leisure	22.43	45.78	26.57		41.53	34.49	27.39
VFR	21.57	33.31	26.67		47.62	32.99	33.70
other	26.35	32.10	26.70		48.32	28.25	28.22
West coast	16.69	37.05	26.94		52.85	25.62	30.74
East coast	29.24	31.82	31.25		57.14	36.91	37.08

SAMPLE LEVEL BUT AT INCOME OF 125K (\$/hr)							
	car time	air time	rail time	bus time	access time	egress time	
all	29.59	34.32	32.53		59.75	37.87	38.52
work	47.03	22.49	50.19		105.00	45.73	49.39
vacation	28.50	36.95	28.44		47.72	39.53	40.58
mixed leisure	24.20	49.42	28.63		44.80	37.17	29.53
VFR	22.86	35.43	28.28		50.41	34.98	35.73
other	28.05	34.25	28.42		51.34	30.08	30.06
West coast	17.66	39.34	28.52		55.99	27.13	32.46
East coast	30.81	33.80	32.95		60.14	38.97	39.14

MODE PREFERENCES (EXPRESSED IN \$ TO AVOID BUS)						
car vs bus	air vs bus	rail vs bus		car vs bus	air vs bus	rail vs bus

all	62.50	17.81	29.98	aged under 35	22.56	5.61	-1.02
work	101.16	17.90	72.59	aged 35 to 44	55.18	15.53	24.31
vacation	61.39	27.09	28.17	aged 45 to 54	69.89	24.82	35.72
mixed leisure	58.96	36.72	23.80	aged 55 to 64	64.95	17.95	39.56
VFR	48.81	13.20	18.56	aged 65 and over	60.98	8.98	20.89
other	45.46	-3.21	1.80	no degree	52.35	6.93	2.14
female	55.87	20.77	23.83	graduate	66.83	22.45	41.87
male	70.05	14.43	37.00	employed	67.45	22.74	35.76
West Coast	81.95	33.31	46.30	not employed	51.20	6.56	16.79
East Coast	60.49	16.21	28.30				

For the value of travel time measures, we report results both at the sample level (i.e. taking into account the actual income of respondents) and at an income of \$125K (to avoid confounding between income and purpose effects, especially). Starting with the first set of results, we see that, overall, in vehicle time is valued relatively evenly across car, air and rail, while time spent travelling by bus has a much higher disutility. The value of access and egress time falls in between these two groups. However, notable differences arise across subsegments. For example, respondents on work trips have a lower valuation for air travel time while the gap between bus travel time and other valuations is the strongest across all segments for these respondents. On the other hand, respondents on mixed leisure trips value air travel time the highest (i.e. greatest disutility). Big differences also arise between East Coast and West Coast respondents. While the valuations of in vehicle time are higher for East Coast respondents for car, rail and bus, as well as for access and egress time, West Coast respondents have a higher value of travel time for air. The gap in valuations between the two corridors is most obvious for car, while it is less strong for rail and bus; this is in line with an overall stronger preference for car travel on the West Coast, possibly reflecting differences in typical congestion levels. The second set of results in TABLE 7 are overall very consistent with the sample level results, suggesting that the core differences across segments reported earlier are indeed linked to trip characteristics and geography rather than income differences.

We finally turn to the trade-offs between the mode specific constants and the cost sensitivity, i.e. expressing a monetary valuation for travelling by car, air or rail, as opposed to bus, all else being equal. These valuations are obtained for the final values for the mode specific constants, thus combining δ_i and $\tau_i\alpha_n$. We see that overall, this willingness to pay is highest for car, followed by rail and air. Major differences again arise across purpose segments (e.g. a much higher valuation for car for work trips, and a much lower valuation for air and rail for *other* non-work trips). There are differences by gender (e.g. a stronger male preference for car) as well as by corridor (overall higher willingness to pay to avoid bus on the West Coast). The findings for age, education and employment are in line with the earlier discussions in Section 4.3.

5.6. Role of latent variables in model

As discussed in Sections 5.1 to 5.5, our hybrid model provides a wealth of results and highlights interesting distinctions across socio-demographic groups as well as travellers in different corridors. We already alluded to some examples where the impact in the structural model for the latent variables is different from that in the choice model itself, with a notable example being the more pro-car attitude for female respondents, but their reduced utility for car itself in the choice model (net of the impact of the latent attitudes). This motivates an additional investigation into the sources of heterogeneity in the modal preferences.

Remember that the overall preferences for a given mode i (net of the level of service variables) is given by $\delta_{n,i} + \tau_i\alpha_n$ in Equation [3]. Both $\delta_{n,i}$ and α_n incorporate deterministic and random heterogeneity, as shown in Equations [4] and [1], respectively. We can then contrast the *pure* random heterogeneity introduced into the utility for mode i through $\sigma_{\delta_i}\xi_{n,\delta_i} + \sigma_{\delta_{i,2}}\xi_{n,\delta_i}^2$ with that introduced via the latent variables, given by $\tau_{i,l}\xi_{l,n}$. The results of this are shown in TABLE 8. We first see that the highest level of heterogeneity across respondents occurs for the mode specific constant for car. More importantly, the vast majority of the heterogeneity is in each case attributed to $\delta_{n,i}$, i.e. the component not linked to the attitudinal constructs. This is also in line with the observations from the third panel of FIGURE 2. This finding is in itself not completely surprising. There are many factors that could drive heterogeneity in modal preferences that are not linked to underlying attitudes, including person-specific ease of access to different modes that cannot be completely captured by the attributes in our survey. In addition, it is worth remembering again that the random component in Equation [4] incorporates a higher level of distributional flexibility than the latent variable in Equation [1].

Table 8: Sources of Random Heterogeneity in Mode Constants

	car	air	rail
$var(\delta_{n,i} + \tau_i \alpha_n)$	15.292	3.580	2.728
$var(\delta_{n,i})$	13.224	3.182	2.678
$var(\tau_{i,1} \xi_{n,1})$	0.010	0.003	0
$var(\tau_{i,2} \xi_{n,2})$	1.425	0.126	0
$var(\tau_{i,3} \xi_{n,3})$	0.578	0.320	0.055
$var(\tau_{i,4} \xi_{n,4})$	0.140	0	0

A similar and arguably more interesting investigation is possible for the part of heterogeneity linked to traveller characteristics (remembering that trip characteristics do not affect the latent variables). In particular, we have that $\alpha_{i,n}$ has a deterministic component $\gamma_l z_n$, meaning that the respondent characteristics used in the structural equation for the latent variables (Equation [1]) affect the utility for a given mode i as $\sum_{l=1}^4 \tau_{i,l} \gamma_l z_n$. The same respondent characteristics also influence the utility for mode i directly through $\lambda_i z_n$.

The results of these comparisons are summarised in TABLE 9. We can make a number of specific and interesting observations. For geography (West vs East), there is a bigger direct effect on car and rail, while the effect is bigger through the LVs for air. For car and air, the directionality of the impact is the same through direct effects and through the LVs, but the opposite happens for rail, where we see that the higher baseline preference for rail is reduced by an increased concern for privacy. For car and air, where there are impacts by multiple LVs (unlike for rail), an even more interesting picture emerges. Indeed, the reduced pro-car attitude for West Coast respondents helps reduce the otherwise positive impact on the car utility stemming from increased concern for privacy and reduced interest in technology for West Coast respondents. Similarly, for air, we see that the reduced pro-car attitude for West Coast respondents helps mitigate the increased utility for air resulting from the increased concern for privacy. These opposite effects provide important scope for focussed nudging of attitudes and thus behaviour.

Similarly, interesting findings arise for gender, education and employment. Overall, the direct effect is stronger than the effect through the LVs, across all modes, for these three socio-demographics. However, there are again subtle distinctions to be made between the two. For gender, we see for example that the increased pro-car attitude for women is counteracted by increased sociability, reduced concern for privacy and a more pro-tech attitude, leading to an overall negative impact on the modal utility through the LVs. The same happens for air, where the negative impact through increased sociability and reduced concern for privacy outweighs the pro-car attitude to lead to a negative combined effect, albeit a weak one compared to the direct effect. Similarly, for rail, the reduced concern for privacy at least reduces, albeit lightly, the negative baseline effect for female travellers. For those without a degree, all four attitudes lead to an increase in the baseline utility for car, where this goes some way to countering the reduced direct effect on the utility that could in itself reflect reduced access to cars. The same is the case for air, while the increased concern for privacy for those without a degree serves to further reduce the utility for rail. For respondents not in employment, the reduced pro-car attitude mitigates the positive influence on the car utility caused by greater concern for privacy and reduced interest in technology, but the resulting positive combined effect is outweighed by a negative baseline shift, likely at least in part linked to financial constraints. The same is the case for air, while for rail, the increased concern for privacy further reduces that utility.

Turning finally to age, a number of findings are worth highlighting, such as reduced concern with privacy for older respondents counter-acting the increased utility for car resulting from their more pro-car attitude and reduced pro-tech attitude. Addressing the pro-car attitude of such respondents thus has clear potential to influence mode choice. The same effect can be observed for air in the highest age category, where the

more important role for the privacy concern attitude now outweighs the pro-car attitude to lead to a negative effect.

The richness of these results highlights that simply looking at the overall impact of all LVs and comparing it to the direct effect would only provide a limited picture of what is happening in the models. Furthermore, the opposite directions across the four LVs in some cases lead to clear scope for targeted interventions.

Table 9: Sources of Deterministic Heterogeneity in Mode Constants

	direct	LV1: low sociability	LV2: pro-car	car LV3: concerned about privacy	LV4: pro-tech	combined LV effect	overall effect
West coast (vs East)	0.895	0.003	-0.132	0.205	0.048	0.124	1.019
Female (vs male)	-0.181	-0.029	0.097	-0.055	-0.071	-0.058	-0.239
Aged under 35 (vs 35-44)	-0.278	-0.010	-0.471		-0.108	-0.589	-0.867
Aged 45 to 54 (vs 35-44)	0.463						0.463
Aged 55 to 64 (vs 35-44)	0.148		0.195	-0.133	0.067	0.129	0.277
Aged 65 and over (vs 35-44)	0.148		0.367	-0.323	0.146	0.191	0.339
No graduate (vs graduate)	-0.788	0.014	0.171	0.225	0.050	0.460	-0.328
Not employed (vs employed)	-0.405	0.009	-0.136	0.154	0.045	0.073	-0.332

	direct	LV1: low sociability	LV2: pro-car	air LV3: concerned about privacy	LV4: pro-tech	combined LV effect	
West coast (vs East)	0.036	0.002	-0.039	0.153		0.115	0.151
Female (vs male)	0.311	-0.016	0.029	-0.041		-0.028	0.283
Aged under 35 (vs 35-44)	-0.135	-0.005	-0.140			-0.145	-0.280
Aged 45 to 54 (vs 35-44)	0.463						0.463
Aged 55 to 64 (vs 35-44)	0.148		0.058	-0.099		-0.041	0.107
Aged 65 and over (vs 35-44)	-0.065		0.109	-0.240		-0.131	-0.196
No graduate (vs graduate)	-0.381	0.008	0.051	0.167		0.226	-0.155
Not employed (vs employed)	-0.504	0.005	-0.040	0.115		0.079	-0.424

	direct	LV1: low sociability	LV2: pro-car	rail LV3: concerned about privacy	LV4: pro-tech	combined LV effect	
West coast (vs East)	0.801			-0.063		-0.063	0.737
Female (vs male)	-0.181			0.017		0.017	-0.164
Aged under 35 (vs 35-44)	-0.278						-0.278
Aged 45 to 54 (vs 35-44)	0.463			0.000			0.463

Aged 55 to 64 (vs 35-44)	0.517	0.041	0.041	0.558
Aged 65 and over (vs 35-44)	-0.013	0.100	0.100	0.087
No graduate (vs graduate)	-1.233	-0.069	-0.069	-1.303
Not employed (vs employed)	-0.212	-0.048	-0.048	-0.260

6. Model Application and Scenario Tests

As a final step in testing the impact of the attitudinal constructs on the model, we conducted a number of scenario tests. To apply the model, we used sample enumeration, where we took the 5,413 respondents in our sample and applied the model to each of them, under various what-if scenarios. In particular, we looked at what would happen if attitudes were to change in the future. Rather than arbitrarily looking at a percentage change in a given attitude (which is meaningless given the scale of the constructs), we instead tested what would happen if everyone’s attitudes were like those of a given segment of the population. No changes were made to the remaining parts of the model, i.e. the direct impact of socio-demographic and trip characteristics or the level of service measures.

The results of this exercise are illustrated in FIGURE 3, where it is important to note that the scale differs across the four panels given the large differences in the role of the four latent variables. A total of 12 scenario tests were run for each of the four latent variables. These relate to geography (2 possible attitudes), gender (2 possible attitudes), age (4 possible attitudes given that the base group and 45 to 54 can be combined), education (2 possible attitudes) and employment status (2 possible attitudes). To be more precise, if for example we look at the geography scenario, then in the West Coast run, we would assign everyone the attitudes of a West Coast respondent, but still respecting their other socio-demographics (e.g. age) in the calculation of their attitudes (thus not leading to a homogenous attitude).

We see from the first panel that changes in the sociability attitude have a completely negligible impact on modal split, with very small shifts only, which is caused by the low values for τ_1 (cf. TABLE 4). From the second panel, we can see that for the pro-car attitude (LV2), the biggest changes in mode split can be obtained by adopting either the attitude of the oldest age group (leading to an increase in the mode share for car) or the youngest age group (leading to an increase in the share of public transport). For the privacy attitude (LV3), we see that a shift towards the attitude of the oldest age group would benefit bus and rail, while a shift to West Coast attitudes would reduce the share for these modes. Finally, for the pro-tech attitude (LV4), age again plays the biggest role, with a shift to older attitudes benefiting car, and a shift to younger attitudes benefiting all the public transport modes. A shift to West Coast attitudes would also have a negative impact on the public transport mode shares. Finally, FIGURE 4 looks at a shift in all four latent attitudes at the same time, highlighting the role of age, education and also a subtle East coast – West coast difference.




Figure 3: Scenario Tests (Note Differences In Scale)




Figure 4: Scenario Test For Combined LVs

7. SUMMARY AND CONCLUSIONS

This paper has presented a detailed investigation of intercity mode choice for travellers in two major corridors in the US. Using a flexible hybrid choice model with attitudinal constructs, we allow for extensive heterogeneity in sensitivities across individual travellers. Some of this heterogeneity can be linked to characteristics of the traveller and/or trip, while a remaining part is random variation.

Through the use of the hybrid choice model, a share of this heterogeneity can be attributed to four attitudinal constructs, which we describe as “low sociability”, “pro-car”, “concern for privacy” and “pro-tech”. We in particular find that the “pro-car” attitude and “concern for privacy” attitude play a non-trivial role in the mode choice process. Travellers who have a more favourable attitude towards cars are much more likely to choose car, and also somewhat more likely to choose air, compared to bus and rail, while the same applies to travellers who have a greater concern for privacy. Those more concerned about privacy are also less likely to choose rail than bus, all else being equal. The impact of the remaining two latent attitudes is much weaker, though we do see that more “pro-tech” travellers are less likely to choose car.

A diverse and interesting picture emerges when studying the drivers of these attitudes, with a number of traveller characteristics playing a role. Crucially, we see some differences between the drivers of attitudes and the drivers of choices. For example, women are more pro-car in their attitudes but this aside are less likely to choose car. On the other hand, West Coast respondent are less pro-car in their attitudes, but are

more likely to choose car in the actual choice scenarios. This shows the scope for a disconnect between attitudes/desires and actual choices. A traveller may well have a negative attitude towards car, which will reduce the appeal of car in the choices too, but this can be outweighed by other factors. A core example comes in comparing East Coast and West Coast respondents – even after accounting for differences in travel time, cost and frequency of service, we see that West Coast travellers have a reduced utility for the three non-car options compared to East Coast travellers, net of the impact of the attitudinal constructs. This could be due to a perceived or actually experienced quality of the non-car options on the West Coast. It clearly opens up the possibility of a mode shift for West Coast respondents through improvements to the transport network, albeit that we also need to be mindful that the increased concern for privacy by West Coast travellers outweighs their reduced pro-car attitude.

As highlighted by Chorus & Kroesen (2014), a key issue with hybrid choice models comes in how to use them to inform policy making. Given the arbitrary scale of the latent attitudinal constructs, it is meaningless to look at a percentage shift in attitudes. Rather, we perform a number of “scenario tests” where we investigate how modal splits are likely to change if specific population segments adopted the attitudes of other groups. The most optimistic scenario for public transport is that people will keep their current attitudes toward car orientation (broadly anti-car) and ICT/productivity (broadly pro-tech) as they age (attitudes “stay together with the cohort”) and that “Generation Z” (the next cohort coming after the “Millennials”) will have the same attitudes as current Millennials. Education can also play a role. However, if younger generations maintain their concern for privacy (LV3), then this will counteract the benefits of the other three LVs. The most pessimistic scenario for public transport is that each age cohort will adopt the attitudes toward auto orientation and technology that the previous cohort had at that same age.

Another concern in the literature concerns the empirical benefit of including attitudinal constructs rather than just relying on a flexible Mixed Logit model, a point discussed at length by Vij & Walker (2016). In our application, we take care to avoid misattributing sources of heterogeneity to the latent variables by ensuring that the base utility specification has at least the same level of flexibility in terms of random heterogeneity as well as socio-demographics variables. The analysis in Section 5 shows that, especially for the random component, only a small share of the heterogeneity in modal preferences can actually be linked to the latent constructs, where it is again important to acknowledge that part of this could be due to the more flexible treatment of random heterogeneity in the utility functions through the polynomial specification. Similarly, while some socio-demographic characteristics (especially age for car) have a strong influence on mode choice through the latent attitudes, overall, the direct influence of these characteristics on mode choice through inclusion of socio-demographics in the utility function is bigger than through the latent variables. However, the impact through the latent variables is certainly not negligible. What is more, when looking at individual latent attitudes, there is a very diverse set of strong impacts of socio-demographic characteristics. Some of these go in opposite directions (e.g. women being more pro-car but less concerned about privacy) which explains the reduced total impact through the latent constructs, but which clearly opens up possibilities for targeted interventions to change attitudes and hence behaviour.

We have estimated a very detailed specification of a hybrid choice model, incorporating large amounts of deterministic and random heterogeneity in sensitivities, with parts linked to the attitudinal constructs as well as parts independent of them. For the random heterogeneity in the choice model part, we added additional flexibility by moving away from pure parametric distributions. As discussed in Section 1, hybrid choice models have become a widely applied tool in choice modelling, almost analogous to when “standard” Mixed Logit first became computationally feasible twenty years back. However, it has also become clear from theoretical discussions (especially by Vij & Walker, 2016) that many studies have oversold the benefits of the model and provided flawed comparisons with “base” models. While the results on attitudinal components presented in Section 5 provide additional insights on both the formation and role of the latent attitudes, let us also remember that the complexity of the model used here meant the estimation of 224 parameters with a total of 14 random component in the model. We had to resort to Bayesian estimation and the specification search for the model components took an extensive amount of time, much

more so than would have been the case for a model without the attitudinal constructs. In the context of large scale applications aimed at providing transport policy guidance, a case by case decision about the benefits of hybrid choice models for incorporating attitudinal constructs still needs to be made, and practitioners should not underestimate the complexity of the model or overestimate its benefits. Similarly, there should be no feeling that the application of hybrid choice models is now a requirement.

As is the case with any research study, there is always scope for further developments. Our specification of a hybrid choice model focussed solely on attitudinal construct while there is clearly also scope to look at attribute perceptions and other latent components. Additionally, as mentioned already, we limit the role of the latent variables in the choice model by testing for their impact only on the alternative specific constants, rather than also testing for effects on the sensitivities to level of service characteristics. Further insights could also be gained by increasing the flexibility of the random component of the latent variables, through more complex univariate distributions as well as correlation between the individual latent constructs. With either of these departures, special care would be required to find appropriate normalisations for the model. Finally, our study used the entire sample for estimation rather than working with a hold out sample for validation. This is primarily motivated by the fact that we are dealing with a rather uniform sample in terms of the choices that respondents faced, limiting the risk of overfitting and hence the insights that could be gained from out of sample prediction. Such tests become far more interesting in the presence of data collected at different points in time, which also opens up other benefits, not least in terms of longitudinal measurement of attitudinal indicators, as suggested by Chorus & Kroesen (2014).

ACKNOWLEDGEMENTS

The original study on which this work is based was funded by TRB's NCRP program. Stephane Hess acknowledge the support by the European Research Council through the consolidator grant 615596-DECISIONS for additional modelling work. We also wish to acknowledge the suggestions of three unknown referees who helped us in streamlining the paper.

REFERENCES

- Abou-Zeid, M., Ben-Akiva, M., Bierlaire, M., Choudhury, C., Hess, S. (2010), Attitudes and Value of Time Heterogeneity. In: Van de Voorde, E., Vanellander, T. (Eds.), *Applied Transport Economics - A Management and Policy Perspective*. De Boeck Publishing, Brussels, pp. 523-545.
- Abou-Zeid, M., Ben-Akiva, M. (2014), Hybrid choice models. In: Hess, S., Daly, A. (Eds.), *Handbook of Choice Modelling*. Edward Elgar, Cheltenham.
- Ashok, K, Dillon, W.R., Yuan, S. (2009), Extending discrete choice models to incorporate attitudinal and other latent variables. *Journal of Marketing Research* 39: 31–46.
- Atasoy, D., Glerum, A., Bierlaire, M. (2013), Attitudes towards mode choice in Switzerland. *disP – The Planning Review* 49: 101-117.
- Bahamonde-Birke, F.J., Kunert, U., Link, H., Ortúzar, J. de D. (2017), About attitudes and perceptions: finding the proper way to consider latent variables in discrete choice models, *Transportation*, 44: 475-493.
- Ben-Akiva, M., Walker, J., McFadden, D., Gärling, T., Gopinath, D., Bolduc, D., et al. (1999), Extended framework for modelling choice behaviour. *Marketing Letters* 10: 187–203.
- Ben-Akiva, M., Walker, J., Bernardino, A.T., Gopinath, D.A., Morikawa, T., Polydoropoulou, A. (1999), *Integration of Choice and Latent Variable Models*. Massachusetts Institute of Technology, Cambridge.
- Ben-Akiva, M., McFadden, D., Train, K., Walker, J., Bhat, C., Bierlaire, M. (2002), Hybrid choice models: progress and challenges. *Marketing Letters* 13: 163–75.
- Bhat, C.R., Dubey, S.K. (2014), A new estimation approach to integrate latent psychological constructs in choice modelling. *Transportation Research Part B: Methodological* 67:68–85.

- Bolduc, D., Ben-Akiva, M., Walker, J., Michaud, A. (2005), Hybrid choice models with logit kernel: applicability to large scale models. In: Lee-Gosselin, M., Doherty, S. (Eds.), *Integrated Land-Use and Transportation Models: Behavioural Foundations*. Elsevier, Oxford, pp. 275–302.
- Chorus, C., Kroesen, G. (2014), On the (im-)possibility of deriving transport policy implications from hybrid choice models. *Transport Policy* 36: 217-222.
- Daly, A.J., Hess, S., Patruni, B., Potoglou, D., Rohr, C. (2012), Using ordered attitudinal indicators in a latent variable choice model: a study of the impact of security on rail travel behaviour. *Transportation* 39: 267-297.
- Daziano, R.A. (2015), Inference on mode preferences, vehicle purchases, and the energy paradox using a Bayesian structural choice model. *Transportation Research Part B: Methodological* 76: 1–26.
- Daziano, R.A., Bolduc, D. (2013), Incorporating pro-environmental preferences towards green automobile technologies through a Bayesian hybrid choice model. *Transportmetrica A: Transport Science* 9: 74–106.
- Dumont, J., Keller, J. (2015), RSGHB: functions for hierarchical Bayesian estimation: a flexible approach. R package version 1.1.2 (<https://CRAN.R-project.org/package=RSGHB>).
- Glerum, A., Stankovikj, L., Thémans, M., Bierlaire, M. (2013), Forecasting the demand for electric vehicles: accounting for attitudes and perceptions. *Transport Science* 48: 483-499.
- Glerum, A., Atasoy, B., Bierlaire, M. (2014), Using semi-open questions to integrate perceptions in choice models. *The Journal of Choice Modelling* 10: 11–33.
- Fosgerau, M., Mabit, S. (2013). Easy and flexible mixing distributions. *Economic Letters* 120: 206–210.
- Kamargianni, M., Ben-Akiva, M., Polydoropoulou, A. (2014), Incorporating social interaction into hybrid choice models. *Transportation* 41: 1263-1285.
- Kim, J., Bae, Y.K., Chung, J.H. (2012), Effects of personal pro-environmental attitudes on mode choice behaviour: new eco-friendly water transit system in Seoul, South Korea. *Transportation Research Record* 2274:175-183.
- Kim, J., Rasouli, S., Timmermans, H. (2014), Hybrid choice models: principles and recent progress incorporating social influence and nonlinear utility functions. *Procedia Environmental Sciences* 22: 20-34.
- Kim, J., Rasouli, S., Timmermans, H. (2016), A hybrid choice model with a nonlinear utility function and bounded distribution for latent variables: application to purchase intention decisions of electric cars. *Transportmetrica A: Transport Science* 12: 909-932.
- Mackie, P., Wardman, M., Fowkes, A.S., Whelan, G., Nellthorp, J., Bates, J.J. (2003), Values of travel time savings in the UK. Report to Department of Transport. Institute for Transport Studies, University of Leeds and John Bates Services, Leeds and Abingdon.
- NCRRP (2016), Intercity Passenger Rail in the Context of Dynamic Travel Markets, NCRRP Report 4, Washington, DC.
- Ortúzar, J. de D., Willumsen, L.G. (2011), *Modelling Transport*. John Wiley & Sons, Chichester.
- Prato, C.G., Bekhor, S., Pronello, C. (2012), Latent variables and route choice behaviour. *Transportation* 39: 299-319.
- Raveau, S., Yáñez, M.F., Ortúzar, J. de D. (2012), Practical and empirical identifiability of hybrid discrete choice models. *Transportation Research Part B: Methodological* 46: 1374–1383.
- Rose, J.M., Bliemer, M.C.J. (2014), Stated choice experimental design theory: the who, the what and the why. In Hess, S. and Daly, A. (Eds.), *Handbook of Choice Modelling*. Edward Elgar, Cheltenham, pp. 152-177.
- Thorhauge, M., Haustein, S., Cherchi, E. (2016), Accounting for the theory of planned behaviour in departure time choice. *Transportation Research Part F: Psychology and Behaviour* 38, 94-105.

- Train, K.E. (2001), A Comparison of Hierarchical Bayes and Maximum Simulated Likelihood for Mixed Logit. Working Paper, Department of Economics, University of California at Berkeley.
- Vij, A., Walker, J. (2014), Hybrid choice models: the identification problem. In: Hess, S., Daly, A. (Eds.), Handbook of Choice Modelling. Edward Elgar, Cheltenham, pp. 519–564.
- Vij, A., Walker, J. (2016), How, when and why integrated choice and latent variable models are latently useful. Transportation Research Part B: Methodological, 90: 192-217