How Effective are Marker Variables at Imputing Attitudes? An External Evaluation Using Vehicle Ownership Models

Jason Soria, Ph.D.
School of Civil and Environmental Engineering
Georgia Institute of Technology
790 Atlantic Drive, Atlanta, GA 30332
Email: jasonsoria@gatech.edu
ORCID: https://orcid.org/0000-0003-1319-9197

Patricia L. Mokhtarian, Ph.D.
School of Civil and Environmental Engineering
Georgia Institute of Technology
790 Atlantic Drive, Atlanta, GA 30332
Email: patmokh@gatech.edu
ORCID: https://orcid.org/0000-0001-7104-499X

TRB 103rd annual meeting, January 7-11, 2024
Submission date: August 1, 2023

Tables: 4 x 250 = 1000 words
Abstract: 250 words
Manuscript (including acknowledgements and citations): 6236 words
Total word count: 7486

Abstract

Information obtained from travel behavior surveys is used to inform major policy decisions affecting many people and industries. Such surveys are generally designed with minimizing respondent burden in mind, thereby reducing the number of topics a single survey can address. Attitudinal variables, although inarguably important to explaining behavior, are generally excluded from such surveys for that reason among others. We demonstrate the possibility of using a small set of attitudinal statements to adequately represent factor scores obtained from a much larger set. Specifically, we utilize a completed factor analysis to identify "marker" (important) attitudinal items, apply a Random Forest approach to those items to estimate ("impute") the factor scores assuming that the other items are unavailable, develop vehicle ownership models without and with the imputed attitudes, then compare the model estimation results and their predictive power. Attitudes provide more information than traditional modeling variables by themselves do, and particularly help predict zero-vehicle households better – likely because owning vehicles is the norm in the study area (Georgia, USA), and the attitudes help identify the exceptions. Comparing models with imputed attitudes to one using the original factor scores, those with imputed attitudes predicted vehicle ownership with less accuracy, but better than the model without any attitudes. This study shows that marker variables can successfully be utilized to add valuable attitudinal information to otherwise conventional travel behavior models. We discuss the implications for future survey development – specifically, the prospect that future surveys can include more attitudinal topics without worsening respondent fatigue.

Keywords: travel behavior surveys, vehicle ownership, Random Forest, factor analysis

1. Introduction

The explanation and prediction of travel demand is central to transportation planning, policy, and research. A variety of methods are available for estimating travel demand: at one extreme are simple, parsimonious models requiring little data, statistical knowledge, and computational power, while at the other extreme are complex, advanced, data-hungry models requiring substantial computational resources. Analysts must make tradeoffs to find the situationally-optimal balance between these two extremes. Particularly in the context of real-world travel demand forecasting, models are often simplified by limiting the number and kinds of variables that they account for. Such a decision has the benefits of reducing the data collection burden and involving more straightforward models, at the cost of loss of realism and lack of flexibility for dealing with a rapidly-changing landscape.

Attitudes comprise one type of variable that is commonly omitted from practice-ready travel demand models *(1)*. Conventional wisdom holds that attitudes are difficult to measure, and perhaps impossible to forecast, rendering them unsuitable for inclusion. And yet, attitudes have long been a staple in travel behavior *research*, invariably adding considerable explanatory power to models including them *(2)*. The present paper is the latest link in a chain of studies designed to address this dilemma *(3,4)*. Specifically, this research arc is exploring ways to make the inclusion of attitudinal variables into travel demand forecasting models more practical. Earlier stages established the value of using machine learning methods to impute ("transfer") missing attitudinal variables into the household travel survey datasets that supply the data used to estimate regional travel demand models. Particularly, in preliminary investigation Shaw *(4)* obtained very promising results by using just a few attitudinal statements, called "marker variables" (MVs), to represent the information contained in a far larger set of statements.

The present paper continues that line of inquiry. Specifically, the goals of this paper are twofold: first, to see how well we can predict factor scores (which were created from "numerous" attitudinal statements) from a small number of MVs ("internal evaluation"); and second, to see how well the predicted factor scores perform in a model of household vehicle ownership ("external evaluation").

The rest of this paper is organized as follows. The following section briefly reviews the literature on attitudinal MVs, as well as providing an overview of vehicle ownership (VO) models. Next, we explain the data obtained from a 2017 statewide travel behavior survey, provide descriptive statistics for variables used in the VO model, and describe the EFA results used in the data imputation process. The methodology section discusses MVs particular to this line of research, attitude imputation via Random Forest, and the multinomial logit (MNL) model applied to VO. Next, we present the results of the MNL VO model with and without attitudes, as well as the *VO* prediction accuracy associated with different methods for predicting attitudes. We then conclude by discussing the implications of *MV* prediction accuracy for practitioners and researchers alike.

2. Literature review

Marker variables, specifically abbreviated surveys focused on attitudes, are numerous in the psychology and health fields. A prominent example is using shortened versions of the Big Five personality inventory *(5)*. Within the field of transportation, there are fewer but pertinent examples. Cain et al. *(6)* took a 120-item survey about pedestrian landscapes and created a 54-item version to measure physical activity. Comparing the measures obtained from the full and abbreviated surveys, a strong correlation of 0.94 was found. Cerin et al. *(7)* developed an abbreviated version of the Neighborhood Walkability Scale using the correlation between the original scale and the Walk Score index *(8)*, which was later validated as an instrument for measuring neighborhood walkability *(9)*. Shaw *(4)* further explores the efficacy of MVs in several models of mode usage, finding that the models were significantly enhanced with the addition of attitudinal factor scores predicted from MVs. The attitudes improved the model fits for VO, ridesharing usage, vehicle-miles driven, public transit usage, and bicycle usage. Because the present study also models VO, we briefly review that subject.

To model VO, analysts typically turn to count models such as negative binomial and Poisson models *(10,11)*, or discrete choice models, whether ordered or unordered *(12)*. Ordinal models assume that the choices have a natural ordering, which certainly applies to different levels of VO, and are seen in many studies *(13-16)*. Unordered discrete choice models are often multinomial logit (MNL) or a variant *(17-20)*. Anowar et al. *(12)* gives a thorough review of several types of VO models. The variety of utilized models can be attributed to each formulation having its own pros and cons. Here, we follow Bhat and Pulugurta *(21)*, who found that MNL models of vehicle ownership were preferable to ordinal response models because they are less restrictive in nature.

Just as significant as the model form, the variables used in the model are also important. Traditionally, VO is modeled using socio-economic variables *(22-24)*. Some models have included built environment effects *(25,26)*. Though attitudes have been used to model VO in the past *(25,27)*, they have typically been quantified from factor-analyzing responses to a full set of attitudinal statements. When conducting an exploratory factor analysis (EFA), typically each attitudinal construct, or factor, is associated with three or more survey statements (items), thus entailing relatively larger data requirements than the MV approach we propose in this research. Here, we hypothesize that MVs can be used effectively in a VO model, and that the resulting model has better predictive accuracy than a model without attitudes altogether.

3. Data

The data consists of responses to the Georgia Department of Transportation (GDOT)-funded Emerging Technologies Survey conducted from September 2017 to January 2018. This section includes information relevant to the present study; more information can be found in Kim et al. *(28)*. The survey collected data on socio-economic characteristics of individuals and households, technology use, current and future expected travel behavior, and general attitudes and preferences.

Respondents were recruited through invitations mailed to two groups of individuals. Each invitation included a cover letter explaining the survey purpose as well as a paper copy of the survey, and provided individual-specific access codes to the online version of the survey. The first group of individuals was a geographically-stratified random address-based sample of 30,000 adults living in the 14 Metropolitan Planning Organization areas of the state of Georgia. Approximately 1,800 responses were collected from this group. The second group consisted of respondents to the Georgia subsample of the 2017 National Household Travel Survey (NHTS) who indicated that they were willing to be contacted about future surveys *(29)*. From this group, approximately 1,500 responses were collected. Data cleaning removed observations based on incomplete responses, surveys that were completed too quickly, flatlining (repeated answers within blocks of questions), and failed attention checks. The latter were questions that directed the survey taker to choose a specific response (e.g., “To confirm you’re really reading this, please select ‘strongly agree’ here”). After cleaning the data, we retained 3,178 observations for this analysis. Table 1 includes descriptive statistics of variables used in the VO model. The few missing responses within this dataset were imputed using a meticulously tuned Random Forest algorithm. For example, of the 46 attitude statements a respondent saw, a missing response to a single statement was imputed using that individual’s socio-economic variables and responses to other statements.

Table 1 Descriptive statistics of the modeling variables (N = 3,178)

Variable type	Variable name	Additional description	Weighted average	Weighted std. deviation
Continuous variables	Case weight		1	1.383
	HH drivers	Number of people in household with a driver's license	1.838	2.833
	Activity density	Population plus employment per acre	5.076	16.195
	Number of stores	Number of stores within a mile of home	11.045	21.537
	Non-car alternatives	Attitude factor score	-0.008	1.760
	Pro-car owning	Attitude factor score	-0.032	1.901
	Pro-suburban	Attitude factor score	-0.028	1.737
	Modern urbanite	Attitude factor score	0.132	1.716
			Count (unweighted)	**Weighted share**
Categorical variables	MPO tier:			
	Non-MPO	Residential location is not within an MPO region	200	0.171
	Small	MPO region with population less than 200,000	801	0.135
	Midsize	Non-Atlanta MPO region with population greater than 200,000	1148	0.182
	Atlanta	Base category in MNL model	1029	0.512
	Annual HH income:			
	Less than $50K	Base category in MNL model	996	0.420
	$50K to $100K		1162	0.317
	More than $100K		1020	0.263
Dependent variable	Vehicle ownership category:			
	0	Base category in MNL model	78	0.054
	1		822	0.330
	2		1270	0.335
	3+		1008	0.281

An exploratory factor analysis (EFA) was used to identify underlying relationships among several attitudinal items. For this study, EFA condensed correlated responses to survey items on transportation mode preferences, residential location, technology, travel, time-use, and comfort with other people. The result of using this method on these items is shown in Table 2. Because the EFA results are readily available from an existing report and for brevity, we refer readers to Kim et al. *(28)* for details. The rest of this subsection is dedicated to explaining the essence of the factor analysis and how it is applied to this MV research. From the initial 46 statements, the analysis resulted in a 15-factor solution based on principal axis factoring with an oblimin rotation and using 38 statements (after discarding those not loading heavily on any factor, or comprising the only strongly-loading item on a factor). Bartlett scores were calculated for each factor. For this research, the four factors expected to be most relevant to VO were selected, and one MV for each factor is identified in Table 2.

Table 2 Exploratory factor analysis results

Factor	Statement	Pattern matrix loading[a,b]	Marker variable
Non-car alternatives	I like the idea of walking as a means of travel for me.	0.666	Yes
	I like the idea of bicycling as a means of travel for me.	0.628	
	I like the idea of public transit as a means of travel for me.	0.336	
Pro–car-owning	I definitely want to own a car.	0.748	Yes
	I am fine with not owning a car, as long as I can use/rent one any time I need it.	−0.576	
	I like the idea of driving as a means of travel for me.	0.535	
	As a general principle, I'd rather own things myself than rent or borrow them from someone else.	0.404	
Pro-suburban	I prefer to live in a spacious home, even if it’s farther from public transportation or many places I go to.	0.609	Yes
	I see myself living long-term in a suburban or rural setting.	0.387	
Modern urbanite	I like the idea of having stores, restaurants, and offices mixed among the homes in my neighborhood.	0.417	Yes
	My phone is so important to me, it’s almost part of my body.	0.350	
Tech-savvy	Learning how to use new technologies is often frustrating for me.	−0.866	
	I am confident in my ability to use modern technologies.	0.801	
Commute benefit	My commute is a useful transition between home and work (or school).	0.677	
	My travel to/from work (or school) is usually pleasant.	0.579	
	I wish I could instantly be at work (or school)—the trip itself is a waste of time.	−0.428	
Work-oriented	At this stage of my life, having fun is more important to me than working hard.	−0.572	
	I'm too busy to have as much leisure time as I'd like.	0.527	
	It’s very important to me to achieve success in my work.	0.298	
Materialistic	I usually go for the basic (“no-frills”) option rather than paying more money for extras.	−0.565	
	The functionality of a car is more important to me than the status of its brand.	−0.431	
	I would/do enjoy having a lot of luxury things.	0.426	
	I like to wait a while rather than being first to buy new products.	−0.357	
	I prefer to minimize the amount of things I own.	−0.341	
Polychronic	I prefer to do one thing at a time.	−0.834	
	I like to juggle two or more activities at the same time.	0.697	
Pro-environmental	Cost or convenience takes priority over environmental impacts (e.g., pollution) when I make my daily choices.	−0.914	
	I am committed to an environmentally friendly lifestyle.	0.481	
Pro-exercise	The importance of exercise is overrated.	−0.669	
	I am committed to exercising regularly.	0.663	
	Family/friends play a big role in how I schedule my time.	0.612	

Family/friends-oriented	It's okay to give up a lot of time with family and friends to achieve other worthy goals.	−0.468	
Waiting-tolerant	Having to wait is an annoying waste of time. Having to wait can be a useful pause in a busy day.	−0.831 0.533	
Travel liking	I generally enjoy the act of traveling itself. I like exploring new places.	0.618 0.593	
Sociable	I consider myself to be a sociable person. I'm uncomfortable being around people I don't know.	0.563 −0.507	

Note: Among the four factors of interest, the highest-magnitude correlation is 0.41 between non-car alternatives and pro-car owning

[a] Oblimin rotation; [b] Statements with loadings lower than 0.30 are suppressed (with an exception near that threshold for the work-oriented factor)

4. Methodology

4.1 Overview

Figure 1, adapted from Shaw's *(4)* Figure 4.1, depicts a schematic of our general approach to imputing missing (attitudinal) variables using machine learning. We use the superscripts C to denote variables that are common in both datasets and U to denote variables that are unique to each dataset. Suppose that in the "donor" dataset D_D there exists an outcome of interest O_D, which was calculated using variables Y_D^C and Y_D^U. Suppose there is another, "recipient", dataset D_R, having some variables in common with D_D, but not enough variables to calculate O_R because it is missing variables that are unique to D_D (i.e. Y_D^U). The goal of this approach is to develop a transfer process τ using f_D, a function trained using common variables X_D^C and Y_D^C to predict O_D, and then applying that function to the common variables X_R^C and Y_R^C to predict the outcome $\hat{O}_R$. The "enriched" dataset D_R' is the result of applying the transfer process $\tau_D = \{O_D, f_D(X_D^C, Y_D^C)\}$ and now containing a prediction $\hat{O}_R = f_D(X_R^C, Y_R^C)$.

Particularized to the present study, the donor dataset D_D contains socio-economic variables X_D^C and X_D^U, a list of attitudinal items Y_D^C and Y_D^U used in a factor analysis to quantify attitudes related to transportation, and the attitudinal factor scores O_D, calculated from Y_D^C and Y_D^U (see Section 3.1). We want the recipient dataset D_R to contain factor scores O_R, but this dataset does not contain all the attitudinal items needed to calculate them. So, we develop a transfer process τ, using the variables that are common to D_D and D_R to predict $\hat{O}_R$. In the transfer process for this study, our learning function f_D is a Random Forest model. Further, to avoid potential downstream endogeneity issues associated with using other variables (X_D^C) in the training function, we explore the accuracy of the transfer process supposing that *only* the common attitude items (MVs) Y_D^C and Y_R^C are available for f_D. We assess that accuracy through examining the closeness of the predicted scores to the original scores (internal evaluation), and then assess the efficacy of the predicted (imputed) factor scores in a vehicle ownership model (external evaluation).

In Section 4.2 we describe the transfer process with specific details. The description preserves the generality of Figure 1, but the present study departs from convention in one important way: for this proof-of-concept stage, *the donor, recipient, and (ultimately) enriched datasets are the same* ($D_D = D_R = D_R'$). In future research, the two samples will be entirely different.

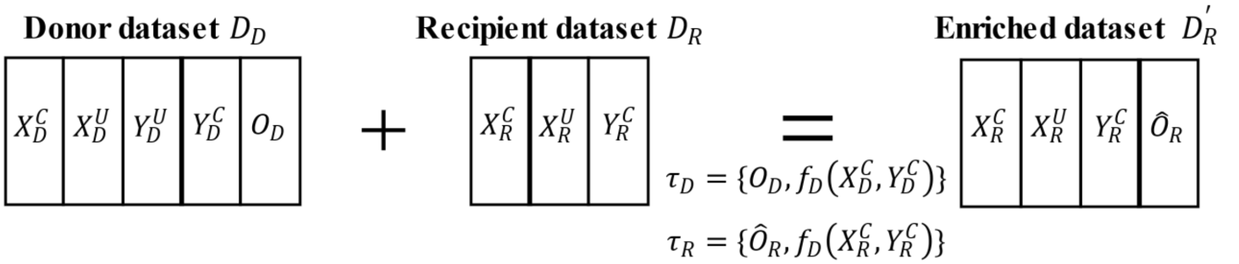

Figure 1 Schematic of the relationships among the various datasets used in this methodology, adapted from Shaw's *(4)* Figure 4.1

4.2 Developing the transfer process to impute attitude scores

The EFA described in Section 3 uses the most information in quantifying attitudes given a set of attitudinal statements, which are later used as explanatory variables in the MNL model described in the next subsection. However, the purpose of this research is to highlight the applicability of MVs for attitude score imputation. To investigate their efficacy, we use two methods of attitude score imputation in the transfer process. The first method simply standardizes the quantified MV responses and uses them directly to represent the factor scores. Utilizing the original coding of item responses (1 = strongly disagree, 2 = disagree, …, 5 = strongly agree), we transformed the four MVs to each have a mean of 0 and standard deviation of 1.

The second method is the machine learning-based transfer function approach sketched in Section 4.1, here operationalized with a non-parametric, data-driven Random Forest model. For brevity, we refer readers to Biau and Scornet *(30)* for the formulation and specifics of Random Forest. This model employs multiple decision trees, which use numeric cutoff points to predict an outcome. The steps below describe the development of the transfer process. The steps are also visually summarized in Figure 2.

Step 1. Factor analyze the attitudinal items in the entire pooled sample, with the resulting (Bartlett) scores treated as the "observed attitudes" O_D seen in Figure 1. Based on the factor analysis, identify the marker variables from among the full set of attitudinal items – generally the one or two highest-loading items on each factor. We selected the single highest-loading items from four specific factors, as shown in Table 2. These MVs correspond to Y_D^C.

Step 2. Train the transfer function f_D using the donor dataset D_D, in this application only with Y_D^C. In this study, Random Forest models are developed using the following substeps.

Step 2a. Partition the dataset into "training" (80% of the total sample) and "test" (the remaining 20%) datasets with the marker variables Y_D^C as explanatory variables (features) and attitude scores as the dependent variable (target). Importantly, each attitude score produces a different Random Forest model, each experiencing its own hyperparameter tuning process.

Step 2b. Using the training dataset, for a given attitude calibrate the Random Forest model f_D by completing a grid-search with 10-fold cross-validation for each point on the grid, using hyperparameter values "variables to try" $ntry \in \{1, 2, 3, 4\}$ and "number of forests" $nforests \in \{100, 200, 300, \dots, 1000\}$. The combination of grid search values that minimized root mean squared error (RMSE), averaged over the cross-validation process, was taken as the best set of hyperparameters. Beyond this combination, there is a greater possibility of overfitting the model on the training dataset and making it less accurate for out-of-sample prediction. RMSE is defined by Equation 1, where O_{Di} is the observed attitude score of person i in the donor set D_D, $\hat{O}_{Di}$ is the predicted attitude score, and N is the total number of observations within the training dataset. This is the *first internal evaluation*. If satisfied with the calibrated f_D model (i.e. relatively low RMSE), compute $cor(\hat{O}_{Di}, O_{Di})$ as another benchmark, and move on to the next substep. Otherwise, continue model calibration efforts until a satisfactory model is obtained.

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N}\left(\hat{O}_{Di} - O_{Di}\right)^2}{N}} \qquad \text{Equation 1}$$

Step 2c. Conduct the *second internal evaluation* by using the calibrated f_D to predict attitude scores for the test dataset, and measuring the correlation and RMSE between predicted and observed attitude scores to determine if overfitting is an issue. If it is evident (from substantial degradation in correlation and RMSE) that there is overfitting, re-evaluate the chosen hyperparameters until the model satisfactorily predicts test-sample data; otherwise, move to the next step.

Step 3. Now, we apply f_D to the recipient dataset D_R, using its marker variables Y_R^C to predict $\hat{O}_R$ and to create D_R'. In the present study D_D, D_R, and D_R' are the same data.

Step 4. Estimate the vehicle ownership models with MNL (see Section 4.3).

Step 4a. (On D_R') Estimate a vehicle ownership model with no attitudes and only X_R^C. This model represents the lower-bound benchmark, for comparison with the models in the next three substeps.

Step 4b. (On D_R') In addition to X_R^C, estimate a vehicle ownership model with standardized marker variable responses Y_R^C.

Step 4c. (On D_R') In addition to X_R^C, estimate a vehicle ownership model with predicted attitude scores $\hat{O}_R$.

Step 4d. (Only possible with D_D) Estimate a vehicle ownership model with observed attitude scores O_D and X_D^C. This represents the upper-bound benchmark, for comparison of the models in the previous two steps to a model containing the most informed quantification of attitudes.

Step 5. Complete *external evaluation* by comparing these models with respect to interpretability, goodness of fit (GOF), and predictive accuracy.

For each of the marker variables, the best number of predictors to try was three and best number of trees to "grow" was 100. With the hyperparameters decided, the remaining 20% of observations (approximately 650) are used as a test dataset to inspect if the model is overfitting. The results are shown in Figure 3. They indicate that the models do not have an overfitting problem, as the test dataset correlations and RMSE are close to the training dataset values, and past this hyperparameter combination the test-data RMSE gets worse. Figure 3 also reports Pearson's correlation and RMSE between the EFA Bartlett scores and standardized MV responses. This latter RMSE is again calculated with Equation 1, where $\hat{O}_R$ is now the standardized MV. As expected, MVs with higher-magnitude loadings (Table 2) generally have higher correlations with the calculated factor scores, and attitude pairs with higher correlations have lower RMSE. In general, the correlations are quite good, at about 0.8 or higher for three of the four, and 0.6 for the urbanite factor.

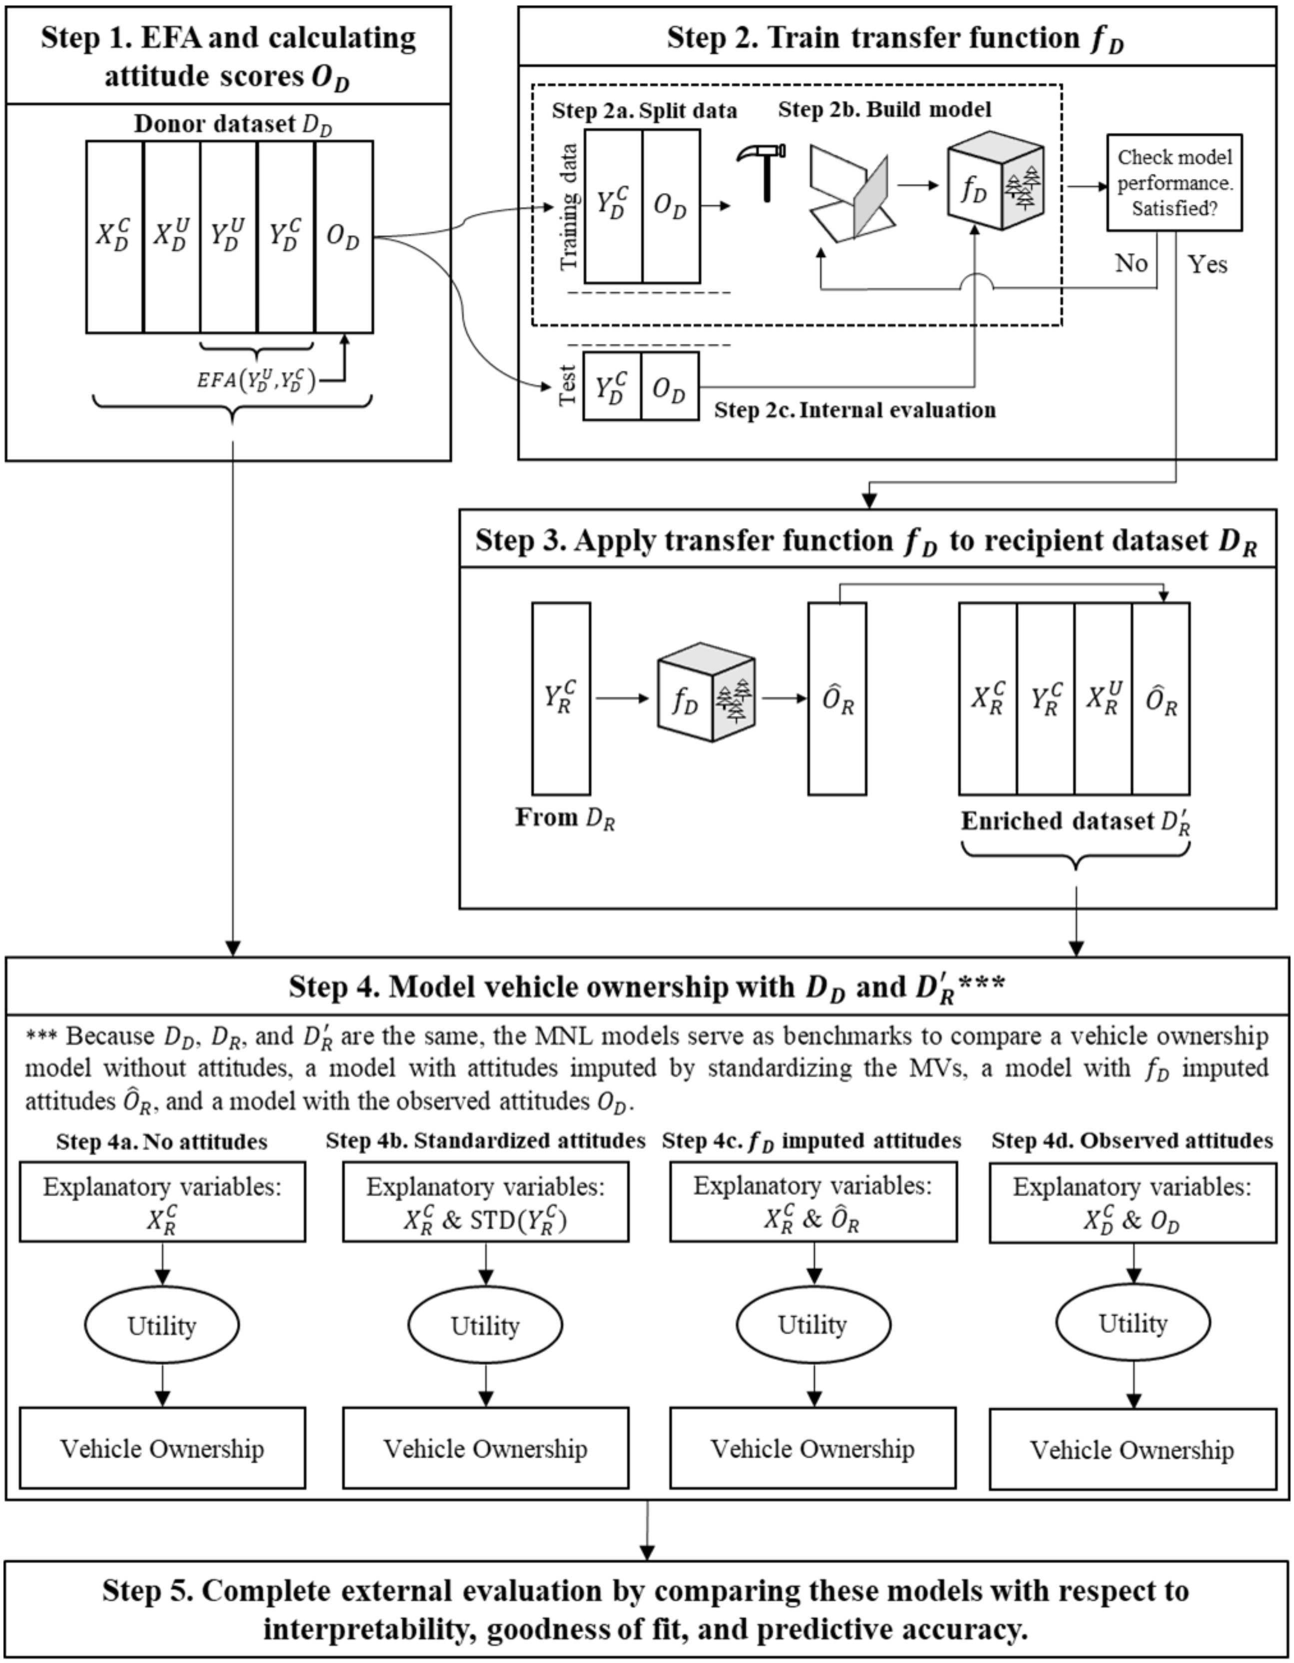

Figure 2 Transfer process specific to this study

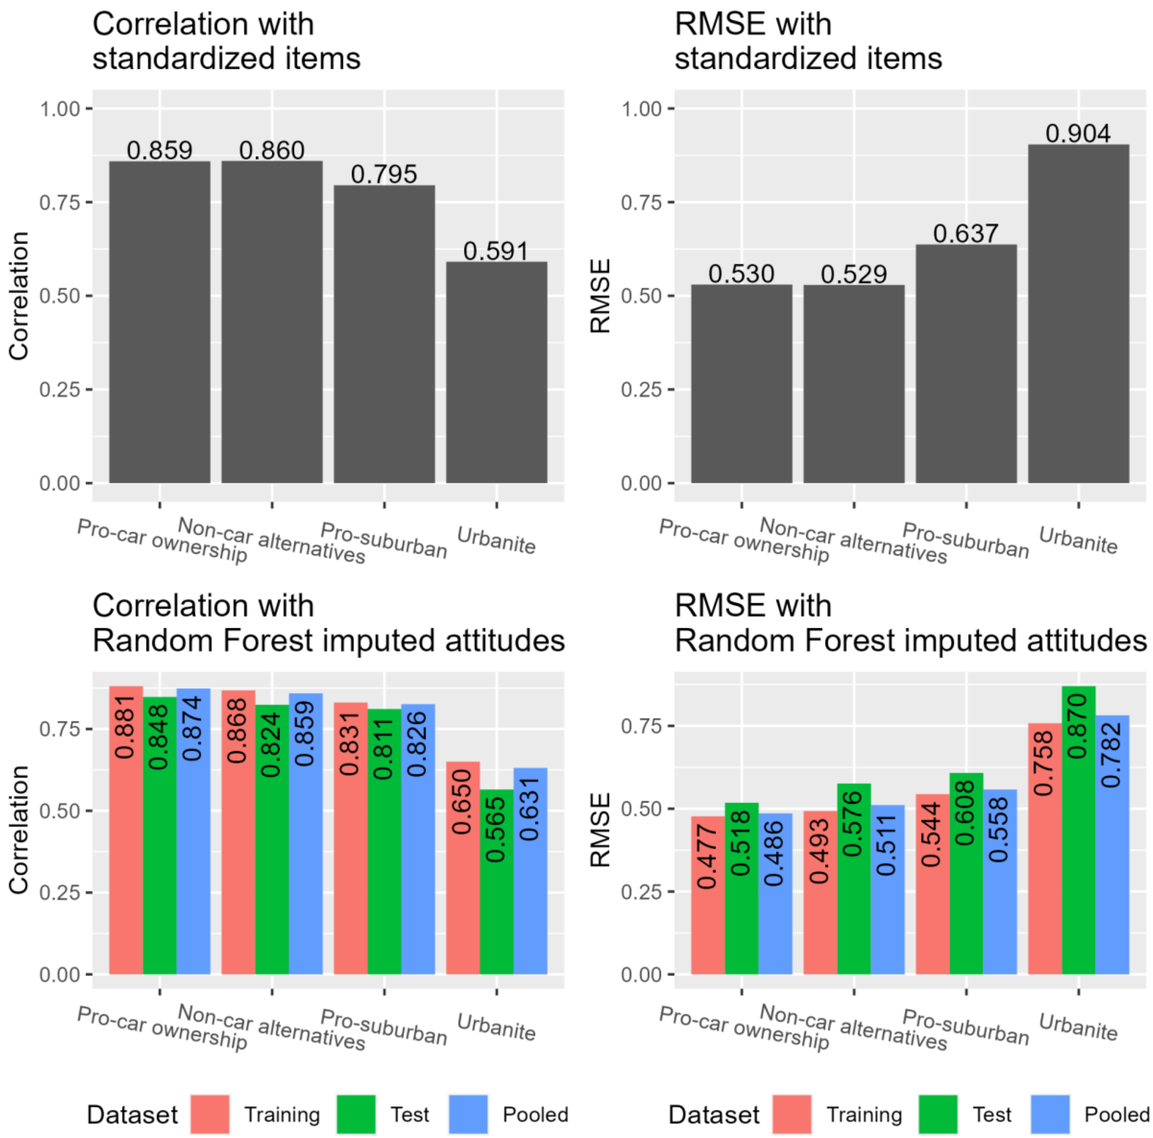

Figure 3 Correlations and root mean squared errors between observed and imputed factor scores

4.3 Multinomial logit for modeling vehicle ownership

Multinomial logit (MNL) is a model widely used across many disciplines. For brevity, we explain the key components of the estimation process for MNL and refer readers to Train *(31)* for more details. In the MNL model, each alternative receives a latent score referred to as "utility," based on a linear-in-parameters combination of explanatory variables. U_{ij}, the utility of person i for alternative $j \in \{0, 1, 2, 3+\}$ vehicles, is a function of a vector of explanatory variables $\boldsymbol{X_i}$, a vector of parameters (i.e. model coefficients, including an alternative specific constant) $\boldsymbol{\beta_j}$ (to be estimated), and error term ϵ_{ij}, assumed to be independently and identically Gumbel distributed (Equation 2). The coefficients are estimated to maximize the log-likelihood of all the observed choices, as described by Equation 3 and Equation 4. To uniquely identify the estimated parameters, we select the 0-vehicle alternative to be the reference category, for which all weights including the alternative specific constant are constrained to equal 0.

$$U_{ij} = \boldsymbol{X_i \beta_j} + \boldsymbol{\epsilon_{ij}} \qquad \text{Equation 2}$$

$$P(\text{person } i\text{'s choice} = j) = P_i(j) = \frac{exp(\boldsymbol{X_i \beta_j})}{\sum_{m=1}^{J} exp(\boldsymbol{X_i \beta_m})} \qquad \text{Equation 3}$$

$$\text{Log-likelihood } LL(\beta) = \sum_{i=1}^{N} \sum_{j=1}^{J} y_{ij} \ln(P_i(j)), \qquad \text{Equation 4}$$

where y_{ij} = 1 if person *i* chose alternative *j*, and 0 else. We then predict the probabilities of person i falling into a VO category and calculate the success index (used to compare the accuracy of predictions across models) shown in Equation 5. The success proportion is the sum of *predicted* probabilities of observations who *actually* choose that alternative divided by the sum of probabilities of an alternative being chosen over all observations, $\frac{\sum_i P_i(j) y_{ij}}{\sum_i P_i(j)}$. The success index is the success proportion normalized by the observed share of that alternative, $\left(\frac{\sum_i y_{ij}}{N}\right)$. Thus, a success index less than 1 indicates that the model does not predict this alternative's adoption even as well as a naïve market-share prediction would, and the greater than 1 the success index is, the better the model predicts adoption compared to the market-share prediction.

$$\text{Success index of alternative } j = \frac{\left(\frac{\sum_i P_i(j) y_{ij}}{\sum_i P_i(j)}\right)}{\left(\frac{\sum_i y_{ij}}{N}\right)} \qquad \text{Equation 5}$$

5. Results

Four MNL models are estimated on the full sample (see Section 4.2), and their results are shown in Table 3 and Table 4. All models contain the same household and residential location variables, with three of them containing four additional attitudinal explanatory variables. Each model uses a different measure of attitudes: the standardized single item approach, Random Forest-imputed scores based on four MVs, and the original EFA Bartlett scores. In this section we first compare the substantive results, and then the fit statistics and predictive power, of the four models.

In all models, consistent with prior research, household characteristics play an important role, highlighting the importance of the number of drivers and household income to determining vehicle ownership. The trend of coefficient values matches our intuition for these variables. Indicated by the positive and statistically significant coefficient values, the probability of owning *any* vehicles increases as the number of drivers increases. And as the number of drivers increases, the probability of owning *more* vehicles also increases. This is indicated by the increasing values of coefficients as the number of household vehicles increases. The same can be said for the household income categories. Relative to the base category of household income less than $50K, the probability of owning more vehicles increases as household income increases. Interestingly, for the no-attitudes MNL model the coefficient in the 1-car alternative for income greater than $100K has a lower value than the coefficient for incomes between $50K and $100K, while the same does not hold true for the model with attitudes. This may indicate that the attitudes are capturing some of the heterogeneity that is attributed to high income in the no-attitudes MNL.

Residential location variables are also included in these models. Relative to the Atlanta MPO, the probability of owning any vehicles decreases as the MPO region gets smaller, except for non-MPO (rural) areas. That is, all else equal, the probability of owning any vehicles versus *not* owning vehicles is not significantly different between Atlanta-region residents and rural residents. This may be because the household and land use characteristics capture the impact that non-MPO designation represents. The negative coefficient values for small and midsize coefficients may be explained by these areas having less activity diversity relative to the Atlanta MPO, thus a lower need for vehicles to visit fewer activities. The other residential location variables, activity density (i.e. population and employment per acre) and the number of stores within a mile, show intuitive results. The negative coefficients show that as both increase the probability of owning vehicles decreases, with the coefficient values becoming more negative as the vehicle category increases.

The inclusion of imputed attitudes in Models 2 and 3 appears to add substantial conceptual content. In these models, all attitudes are statistically significant at the 0.05 level or better except for urbanite (the least-well-predicted attitude among the four) in most instances, and the signs of the coefficients are as expected. In Model 4, using the "observed" attitudes, all four are statistically significant (although relatively weakly so for urbanite), and signs are as expected. In this model, as a person's positive attitude toward non-car alternatives and identification with being a modern urbanite increase, the probability of owning vehicles decreases. Conversely, as a respondent's positive attitudes toward car ownership and pro-suburban lifestyles increase, the probability of owning vehicles increases. In general for the observed attitudes, their coefficient magnitudes increase, meaning that their impact is amplified, as the number of vehicles increases. This pattern does not universally hold for Models 2 and 3 (with estimated attitudes), however. For example, in those two models, coefficients of the pro-car and pro-suburban attitudes do not monotonically increase with the number of vehicles.

Beyond providing a richer narrative, model fit statistics also show modest improvement once attitudes are included. Specifically, McFadden's adjusted ρ^2 (which penalizes models with more parameters) increases by about 4%, from 0.367 to 0.381, between Model 1 and Model 2. Interestingly, there is virtually no difference in the goodness of fit of Models 2 (MVs only) and 3 (Random Forest-imputed attitude scores). Even more interesting, however, is that there is only a negligible difference between those two models and Model 4. Thus, at least in this study, the implications are that the inclusion of a handful of well-chosen attitudinal items (four, in our case) can convey essentially the explanatory power embodied in (a) estimated factor scores obtained from the development and application of a complex machine learning algorithm; and (b) "observed" factor scores computed from a much larger (38, in our case) set of items.

Beyond McFadden's ρ^2 fit measures, the success indices in Figure 4 summarize the predictive power of the four models. Higher values mean more accurate predictions for vehicle ownership. The table shows that incorporating attitudes makes the greatest improvement in the prediction of owning zero vehicles, where the addition of even the standardized MV responses (Model 2) and machine learning-predicted attitude scores (Model 3) bring noticeable improvement over not having any attitudes at all. This seems quite natural, because 0-vehicle households are probably the most heterogeneous of the four VO categories, combining households who own zero vehicles by choice with those who own none by necessity. Accordingly, predicting 0-vehicle households more accurately may require the use of attitudes, since variables such as income and residential location characteristics do not capture this heterogeneity.

Similar to McFadden's adjusted ρ^2, the success indices also indicate that Model 2's predictions, using only the standardized Likert-type responses to the original MVs, did about as well as, and sometimes (notably for VO = 0) better than, Model 3's predictions, using the Random Forest-predicted attitude scores. Furthermore, we see again that on all success indices, Model 4 does only negligibly better than Models 2 and 3. These results reinforce the conclusion that neither elaborate imputation schemes nor lengthy batteries of attitudinal items in surveys may be needed so long as appropriate MVs are chosen.

Table 3 VO Models 1 and 2 results: without attitudes, and with standardized attitudinal items (MVs) (N = 3,178)

	1. NO ATTITUDES						**2. STANDARDIZED ATTITUDINAL ITEMS**					
	VO[a] = 1		VO = 2		VO = 3+		VO[a] = 1		VO = 2		VO = 3+	
Coefficient	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err
Alt. specific constant	0.615	0.45	-2.510***	0.49	-5.410***	0.53	0.571	0.47	-2.679***	0.52	-5.649***	0.56
No. of household drivers	1.980***	0.29	4.000***	0.31	5.140***	0.33	2.142***	0.32	4.206***	0.34	5.342***	0.35
Annual household income[b]: $50K to $100K	1.950***	0.35	2.860***	0.37	3.350***	0.38	2.021***	0.36	2.930***	0.38	3.436***	0.39
Annual household income: > 100K	1.440**	0.52	2.890***	0.52	3.420***	0.53	1.928***	0.46	3.549***	0.47	4.064***	0.48
MPO[c]: non-MPO	-0.016	0.28	0.042	0.31	0.624	0.33	-0.180	0.30	-0.134	0.33	0.404	0.35
MPO: small	-0.892***	0.26	-1.050***	0.30	-0.735*	0.32	-1.098***	0.28	-1.322***	0.32	-1.070**	0.33
MPO: midsize	-0.843***	0.24	-0.826**	0.27	-0.695*	0.30	-0.916***	0.26	-0.942**	0.29	-0.837**	0.31
Activity density	-0.033***	0.01	-0.050***	0.01	-0.079***	0.02	-0.029***	0.01	-0.043***	0.01	-0.062***	0.02
Number of stores	-0.047**	0.02	-0.098***	0.02	-0.122***	0.02	-0.038*	0.02	-0.086***	0.02	-0.108***	0.02
Attitude: non-car alternatives							-0.420***	0.10	-0.479***	0.12	-0.618***	0.12
Attitude: pro-car ownership							0.208**	0.07	0.574***	0.09	0.408***	0.11
Attitude: pro-suburban							0.365***	0.10	0.323**	0.11	0.536***	0.12
Attitude: urbanite							0.006	0.11	-0.083	0.12	-0.134	0.13
Log-likelihood of equally-likely (EL) alternatives model	-4405.643						-4405.643					
Final log-likelihood	-2763.534						-2687.875					
McFadden's ρ^2 (EL benchmark)	0.373						0.390					
McFadden's adjusted ρ^2 (EL benchmark)	0.367						0.381					

* p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001; [a] Reference alternative is 0-vehicle household; [b] Base category is annual household income < $50K; [c] Base category is Atlanta MPO region

Table 4 VO Models 3 and 4 results: Random Forest attitudes imputed with only MVs, and observed attitudes (N = 3,178)

	3. RANDOM FOREST IMPUTED ATTITUDES						4. OBSERVED ATTITUDES					
	VO[a] = 1		VO = 2		VO = 3+		VO[a] = 1		VO = 2		VO = 3+	
Coefficient	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err	Value	Rob. Std err
Alt. specific constant	0.627	0.47	-2.643***	0.51	-5.600***	0.55	0.612	0.47	-2.710***	0.51	-5.720***	0.56
No. of HH drivers	2.068***	0.31	4.153***	0.33	5.287***	0.34	1.990***	0.31	4.090***	0.33	5.240***	0.34
Annual household income[b]: $50K to $100K	1.958***	0.36	2.853***	0.38	3.340***	0.39	2.100***	0.39	2.980***	0.41	3.460***	0.42
Annual household income: > 100K	2.000***	0.48	3.627***	0.48	4.115***	0.49	2.180***	0.50	3.830***	0.50	4.370***	0.51
MPO[c]: non-MPO	-0.163	0.30	-0.109	0.34	0.419	0.35	-0.138	0.30	-0.077	0.33	0.459	0.34
MPO: small	-1.045***	0.28	-1.260***	0.32	-1.009**	0.34	-0.934***	0.28	-1.130***	0.32	-0.853*	0.34
MPO: midsize	-0.888***	0.25	-0.903**	0.29	-0.799*	0.31	-0.854***	0.26	-0.885**	0.29	-0.773*	0.31
Activity density	-0.028***	0.01	-0.042***	0.01	-0.058***	0.02	-0.027**	0.01	-0.037***	0.01	-0.053**	0.02
Number of stores	-0.038*	0.02	-0.087***	0.02	-0.109***	0.02	-0.028	0.02	-0.075***	0.02	-0.096***	0.0
Attitude: non-car alternatives	-0.504***	0.11	-0.538***	0.13	-0.655***	0.14	-0.296**	0.10	-0.267*	0.11	-0.381***	0.11
Attitude: pro-car ownership	0.246**	0.08	0.674***	0.10	0.580***	0.11	0.402***	0.07	0.806***	0.10	0.817***	0.10
Attitude: pro-suburban	0.365**	0.13	0.292***	0.14	0.577***	0.15	0.332**	0.11	0.396**	0.12	0.521***	0.13
Attitude: urbanite	0.126	0.12	0.003	0.14	-0.045	0.15	-0.209*	0.10	-0.286**	0.11	-0.396***	0.11
Log-likelihood of equally-likely (EL) alternatives model	-4405.643						-4405.643					
Final log-likelihood	-2688.466						-2684.081					
McFadden's ρ^2 (EL benchmark)	0.390						0.391					
McFadden's adjusted ρ^2 (EL benchmark)	0.381						0.382					

* p-value < 0.05, ** p-value < 0.01, *** p-value < 0.001; [a] Reference alternative is 0-vehicle household; [b] Base category is annual household income < $50K; [c] Base category is Atlanta MPO region

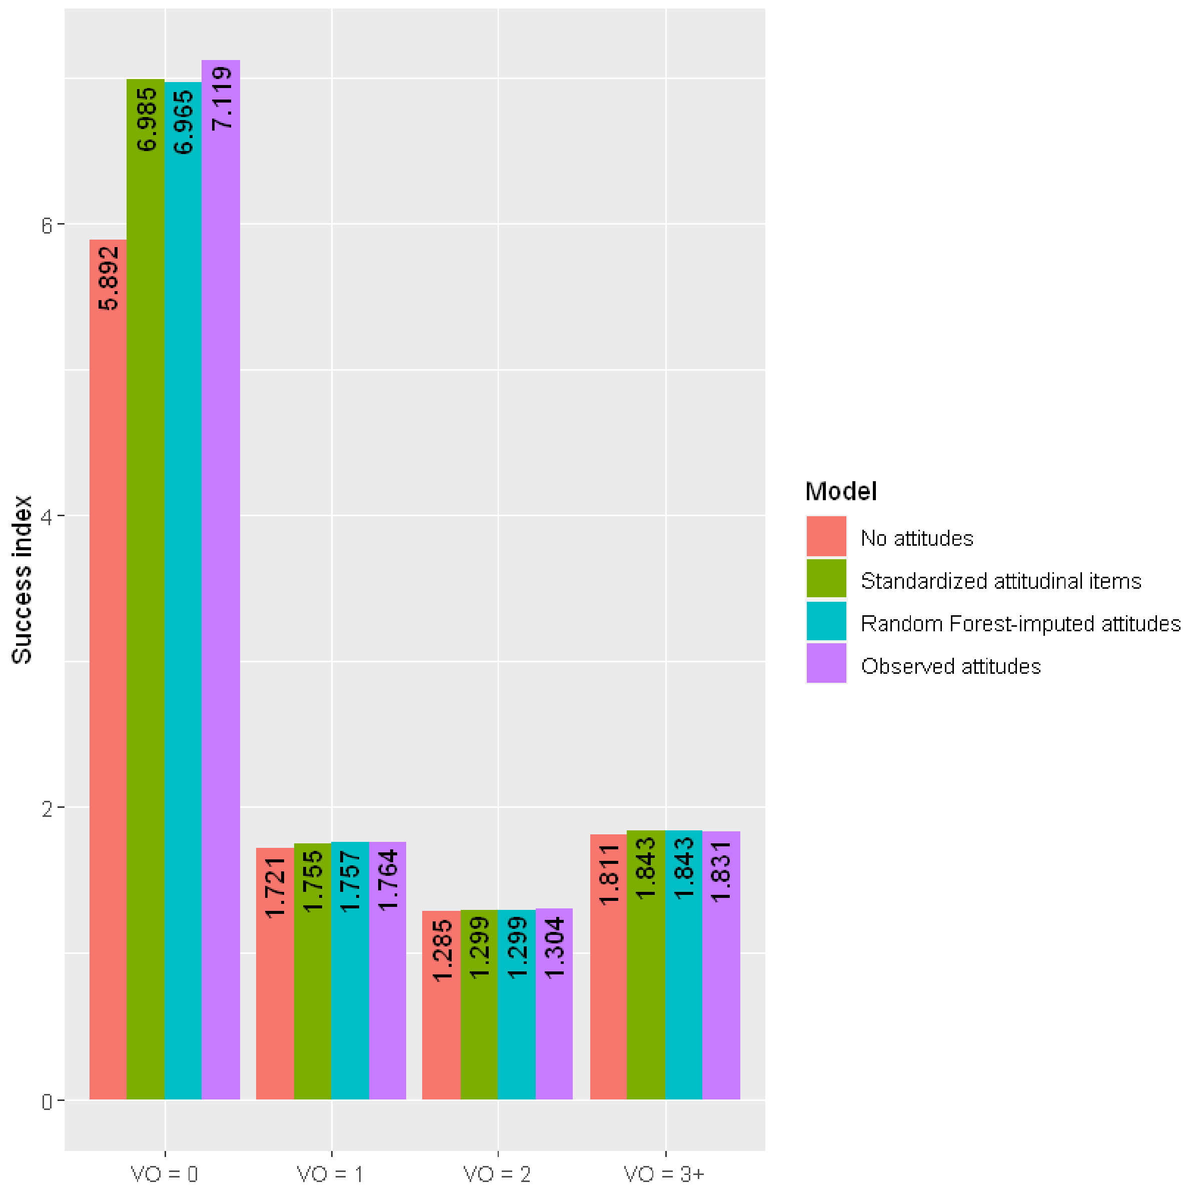

Figure 4 Success indices of the various models containing different levels of attitude inclusion

6. Discussion

Consistent with numerous previous studies *(2)*, our results support the usefulness of attitudes for explaining travel behavior. From Tables 4-5 and Figure 4, containing the model and predictive results, the addition of four attitude variables generates notable improvements in model fit statistics and predictive accuracy. Even when attitude scores are only imputed rather than observed, and even when they are imputed with only a few MVs, vehicle ownership predictions are still more accurate than a similar model without them. We discuss some implications of these results below.

Using MVs, there is less of a need to measure a comprehensive set of attitudinal items, when scores on the underlying constructs can be estimated with fewer items than were needed to derive them. The first phase of this research required data on 38 attitudinal items, which, using EFA, were condensed into scores on 15 factors. Selecting only four items allowed us to predict VO with improved accuracy over a model that included no attitudes (see Figure 4 results). Further, the methodology demonstrated in this study can be extended to identify multiple sets of MVs from different donor surveys, each set designed to predict different behaviors (see Figure 5). For example, the results of VO (behavior A) and mode choice (behavior B) models, which each include measured attitudes from fully informed EFA solutions, could be used to determine which MVs to use when collecting new data, without the need for an excessively large section of attitudinal items in the recipient sample survey. Over time, the profession may identify a more or less standard set of MVs having versatile utility with respect to a number of common travel behaviors, but we would consider standardization to be premature without substantially more experimentation taking place first.

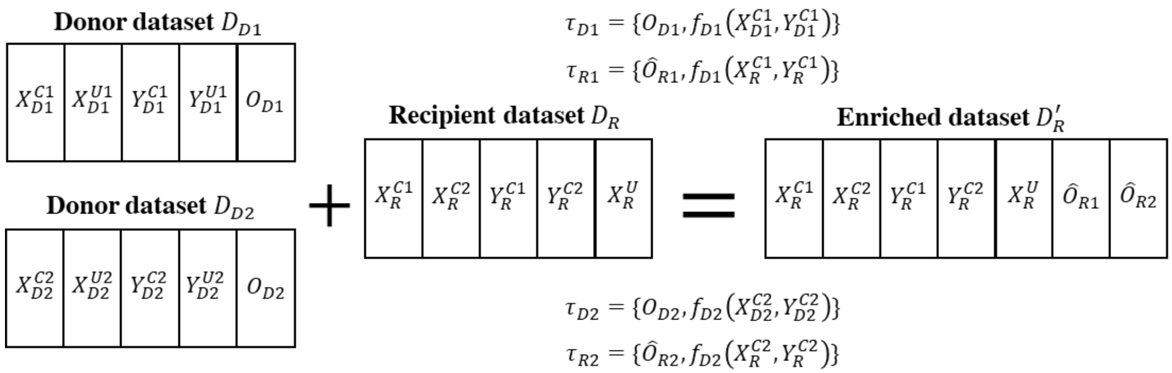

Figure 5 Obtaining MVs from multiple surveys, to predict multiple behaviors

In addition to the possibility of combining several abbreviated surveys for future data collection efforts, this MV analysis also demonstrates the efficacy of using standalone items in travel behavior models to represent attitudes that are usually described with *several* items. The standalone item responses performed generally as well as, and in the 0-vehicle household predictions, better than, the Random Forest-predicted attitudes did. A benefit of the MV selection process is identifying the items that best represent the underlying attitudes, analogous to machine learning algorithms capable of providing feature importance information. Therefore, a strong relationship between item and attitude can likely be used to make more accurate VO predictions (in lieu of unavailable fully informed attitude scores) while maintaining a clear model interpretation.

An important caveat to this line of argument, however, is to beware the reification fallacy, that is, "believing our own labels". Even with a fully-informed factor analysis identifying, say, 10 factors, with the literature-advised *(32)* 3 – 6 statements strongly associated with each factor but with all (say) 40 items contributing richness and texture to the resulting scores, it is all too easy to imbue a factor with the meaning that the analyst wants/expects it to have, beyond the meaning that can be attributed to its most strongly-loading items. It is even easier to do so when a factor is imputed from a set of (say) 10 MVs (only one or two of which will be strongly associated with it, but all of which are, again, contributing to the imputed score), and easiest of all to do so when only a single standardized item is used to represent a factor. In the latter case, the "factor" can technically not be assumed to represent anything beyond the specific content of the item itself. Nevertheless, well-designed items should not be dismissed solely because they are singular *(33,34)* – after all, we include *non-attitudinal* singular variables in models all the time, and in many cases they are just as much "markers" of some more complex causal influence (and interpreted as such) as attitudinal items would be. Surely, including a small set of carefully worded single attitudinal items in a model can usefully explain more variation in behavior than excluding them altogether would do!

Though these results yield novel insights, more research is needed to demonstrate its effectiveness in different contexts. A limitation of this study is that the MV process was used with a single sample rather than applied across altogether different donor and recipient samples. Future research can explore the transferability of MVs across two important dimensions: spatially and temporally. In the spatial sense, the formulas generating attitudinal factor scores, or even the highest-loading items informing those scores, may not be constant across different geographic areas. For example, the order of items that load most strongly onto the pro-car owning attitude may not be the same between highly urbanized and highly rural areas. Similarly, they may not be the same over time. On the other hand, although we may not expect strong spatial or temporal transferability of factor *structures* (i.e. factor *loadings*, and thence factor *score coefficients*), we might well expect stronger stability of the *marker variables themselves*. In other words, the same set of MVs could be efficacious in a large number of contexts, suggesting yet again that simply including MVs directly in the model may be preferable to undergoing elaborate procedures to simulate scores that may not be nearly as reliable across settings. Even so, however, we should not expect a "standard" set of MVs to apply across very different cultures, at least not without considerably more experimentation than has occurred to date.

Continuing the investigation of the machine learning approach to attitude imputation, however, another avenue of future research is investigating how large a donor dataset needs to be to estimate training functions that effectively generate out-of-sample, MV-imputed attitudes. Another line of inquiry is to expand the set of MVs beyond the minimal four used in the present study – e.g. choosing one MV for each of the 15 factors in Table 2 – to allow MVs for other factors to inform the imputation of scores on the target attitudes, as they already do in a standard EFA.

7. Conclusions

By estimating multiple MNL models using varying "intensities" of attitude measurement, we demonstrate that (a) a model with attitudes is a better predictor of vehicle ownership than a model without them and (b) even when the attitudes are imputed using only responses on four items, models with imputed attitudes predict reasonably well compared to the model using the observed attitude scores. This shows that the approach can be an effective tool to approximate attitudes that can be used later in the modeling process. It also suggests that MVs from several donor datasets can be used to develop a survey that has a fraction of the attitudinal items needed to develop a factor analysis from scratch and yet produce a similar result with respect to explanatory power. As this research continues to expand, we acknowledge the limitations of the present study and suggest directions for future research. Specifically, we recommend applying the MV framework to different donor and recipient datasets, and investigating the spatial and temporal transferability of attitudes.

8. Acknowledgements

This study was funded by TOMNET (Teaching Old Models NEw Tricks), a University Transportation Center supported by the U.S. Department of Transportation through Grant No. 69A3551747116. Xinyi Wang, Ilsu Kim, and Seung Eun Choi provided helpful comments that informed the direction of this research. Sung Hoo Kim provided information crucial to understanding the data.

9. References

1. Ewing, R., Sabouri, S., Park, K., Lyons, T., & Tian, G. Key Enhancements to the WFRC/MAG Four-Step Travel Demand Model. NITC-RR-1086. Portland, OR: Transportation Research and Education Center (TREC), 2019. https://dx.doi.org/10.15760/trec.246.
2. Bhagat-Conway, Matthew Wigginton, Laura Mirtich, Deborah Salon, Nathan Harness, Alexis Consalvo, and Shuyao Hong (2022) Subjective variables in travel behavior models: a critical

review and Standardized Transport Attitude Measurement Protocol (STAMP). *Transportation*. https://doi.org/10.1007/s11116-022-10323-7.
3. Malokin, Aliaksandr, Patricia L Mokhtarian, and Giovanni Circella (2019) A Transfer Learning-Based Framework for Enriching National Household Travel Survey Data with Attitudinal Variables. https://doi.org/10.35090/gatech/66813.
4. Shaw, Faaiqa Atiyya. *Methods for Enriching Transportation Survey Datasets: With Sample Applications using Psychometric Variables.* PhD dissertation, Civil and Environmental Engineering, Georgia Institute of Technology, 2021. http://hdl.handle.net/1853/64640.
5. Rammstedt, Beatrice, and Oliver P. John (2007) Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of Research in Personality* **41(1)**, 203-212. https://doi.org/https://doi.org/10.1016/j.jrp.2006.02.001. https://www.sciencedirect.com/science/article/pii/S0092656606000195.
6. Cain, Kelli L., Kavita A. Gavand, Terry L. Conway, Carrie M. Geremia, Rachel A. Millstein, Lawrence D. Frank, Brian E. Saelens, Marc A. Adams, Karen Glanz, and Abby C. King (2017) Developing and validating an abbreviated version of the Microscale Audit for Pedestrian Streetscapes (MAPS-Abbreviated). *Journal of Transport & Health* **5**, 84-96.
7. Cerin, Ester, Terry L Conway, Brian E Saelens, Lawrence D Frank, and James F Sallis (2009) Cross-validation of the factorial structure of the Neighborhood Environment Walkability Scale (NEWS) and its abbreviated form (NEWS-A). *International Journal of Behavioral Nutrition and Physical Activity* **6(1)**, 1-10.
8. Hall, C Michael, and Yael Ram (2018) Walk score® and its potential contribution to the study of active transport and walkability: A critical and systematic review. *Transportation Research Part D: Transport and Environment* **61**, 310-324.
9. Silveira, Stephanie L., and Robert W. Motl (2020) Abbreviated Neighborhood Environment Walkability scale in persons with multiple sclerosis: Initial validation of score inferences. *Journal of Transport & Health* **19**, 100952. https://doi.org/https://doi.org/10.1016/j.jth.2020.100952. https://www.sciencedirect.com/science/article/pii/S2214140520301560.
10. Karlaftis, Matthew, and John Golias (2002) Automobile ownership, households without automobiles, and urban traffic parameters: are they related? *Transportation Research Record* **1792**, 29-35.
11. Shay, Elizabeth, and Asad J Khattak (2007) Automobiles, trips, and neighborhood type: Comparing environmental measures. *Transportation Research Record* **2010**, 73-82.
12. Anowar, Sabreena, Naveen Eluru, and Luis F Miranda-Moreno (2014) Alternative modeling approaches used for examining automobile ownership: a comprehensive review. *Transport Reviews* **34(4)**, 441-473.
13. Baldwin Hess, Daniel, and Paul M Ong (2002) Traditional neighborhoods and automobile ownership. *Transportation Research Record* **1805**, 35-44.
14. Kim, Hong Sok, and Eungcheol Kim (2004) Effects of public transit on automobile ownership and use in households of the USA. *Review of Urban & Regional Development Studies: Journal of the Applied Regional Science Conference*. https://doi.org/10.1111/j.1467-940X.2005.00090.x.
15. Cao, Xinyu, Patricia L Mokhtarian, and Susan L Handy (2007) Do changes in neighborhood characteristics lead to changes in travel behavior? A structural equations modeling approach. *Transportation* **34**, 535-556.
16. Ma, Lu, and Sivaramakrishnan Srinivasan (2010) Impact of individuals' immigrant status on household auto ownership. *Transportation Research Record* **2156**, 36-46.
17. Guo, Zhan (2013) Does residential parking supply affect household car ownership? The case of New York City. *Journal of Transport Geography* **26**, 18-28.
18. Dissanayake, Dilum, and Takayuki Morikawa (2002) Household travel behavior in developing countries: Nested logit model of vehicle ownership, mode choice, and trip chaining. *Transportation Research Record* **1805**, 45-52.

19. Ryan, James M, and Gregory Han (1999) Vehicle-ownership model using family structure and accessibility application to Honolulu, Hawaii. *Transportation Research Record* **1676**, 1-10.
20. Anowar, Sabreena, Shamsunnahar Yasmin, Naveen Eluru, and Luis F Miranda-Moreno (2014) Analyzing car ownership in Quebec City: a comparison of traditional and latent class ordered and unordered models. *Transportation* **41**, 1013-1039.
21. Bhat, Chandra R, and Vamsi Pulugurta (1998) A comparison of two alternative behavioral choice mechanisms for household auto ownership decisions. *Transportation Research Part B: Methodological* **32(1)**, 61-75.
22. Dargay, Joyce, Dermot Gately, and Martin Sommer (2007) Vehicle ownership and income growth, worldwide: 1960-2030. *The Energy Journal* **28(4)**.
23. Dargay, Joyce, and Dermot Gately (1997) Vehicle ownership to 2015: implications for energy use and emissions. *Energy Policy* **25(14-15)**, 1121-1127.
24. Bhat, Chandra R, Rajesh Paleti, Ram M Pendyala, Karen Lorenzini, and Karthik C Konduri (2013) Accommodating immigration status and self-selection effects in a joint model of household auto ownership and residential location choice. *Transportation Research Record* **2382**, 142-150.
25. Kim, Sung Hoo, and Patricia L. Mokhtarian (2018) Taste heterogeneity as an alternative form of endogeneity bias: Investigating the attitude-moderated effects of built environment and socio-demographics on vehicle ownership using latent class modeling. *Transportation Research Part A: Policy and Practice* **116**, 130-150. https://doi.org/https://doi.org/10.1016/j.tra.2018.05.020. https://www.sciencedirect.com/science/article/pii/S0965856417313721.
26. Li, Jieping, Joan L Walker, Sumeeta Srinivasan, and William P Anderson (2010) Modeling private car ownership in China: investigation of urban form impact across megacities. *Transportation Research Record* **2193**, 76-84.
27. Mokhtarian, Patricia L., and Xinyu Cao (2008) Examining the impacts of residential self-selection on travel behavior: A focus on methodologies. *Transportation Research Part B: Methodological* **42(3)**, 204-228. https://doi.org/https://doi.org/10.1016/j.trb.2007.07.006. https://www.sciencedirect.com/science/article/pii/S0191261507000744.
28. Kim, Sung Hoo, Patricia L Mokhtarian, and Giovanni Circella (2019) *The Impact of Emerging Technologies and Trends on Travel Demand in Georgia*. Report FHWA-GA-19-1631, School of Civil and Environmental Engineering, Georgia Institute of Technology. https://rosap.ntl.bts.gov/view/dot/56095.
29. Federal Highway Administration (2018) *Summary of Travel Trends 2017 National Household Travel Survey*. https://nhts.ornl.gov/assets/2017_nhts_summary_travel_trends.pdf.
30. Biau, Gérard, and Erwan Scornet (2016) A random forest guided tour. *TEST* **25(2)**, 197-227. https://doi.org/10.1007/s11749-016-0481-7. https://doi.org/10.1007/s11749-016-0481-7.
31. Train, Kenneth E. *Discrete choice methods with simulation*. Cambridge university press, 2009.
32. Fabrigar, Leandre R, Duane T Wegener, Robert C MacCallum, and Erin J Strahan (1999) Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods* **4(3)**, 272.
33. Gardner, Donald G, Larry L Cummings, Randall B Dunham, and Jon L Pierce (1998) Single-item versus multiple-item measurement scales: An empirical comparison. *Educational and psychological measurement* **58(6)**, 898-915.
34. Diener, Ed (1984) Subjective well-being. *Psychological Bulletin* **95(3)**, 542.