A Bivariate Multinomial Probit Model for Trip Scheduling: Bayesian Analysis of the Work Tour

Jason D. Lemp
Cambridge Systematics, Inc.
9015 Mountain Ridge, Suite 210
Austin, TX 78759
jlemp@camsys.com

Kara M. Kockelman
(Corresponding author)
Professor and William J. Murray Jr. Fellow
Department of Civil, Architectural and Environmental Engineering
The University of Texas at Austin – 6.9 E. Cockrell Jr. Hall
Austin, TX 78712-1076
kkockelm@mail.utexas.edu
Phone: 512-471-0210 & FAX: 512-475-8744

Paul Damien
B.M. Rankin Jr. Professor of Business
McCombs School of Business
The University of Texas at Austin
paul.damien@mccombs.utexas.edu

The following paper is a pre-print and the final publication can be found in
Transportation Science 46 No. 3: 405-424, 2012
Presented at the 90th Annual Meeting of the Transportation Research Board

Key Words: Travel behavior modeling, tour scheduling, multinomial probit, Bayesian estimation

ABSTRACT

As tour-based methods for activity and travel participation patterns replaces trip-based methods, time-of-day (TOD) choice modeling remains problematic. In practice, most travel demand model systems handle tour scheduling via joint-choice multinomial logit (MNL) models, which suffer from the well-known independence of irrelevant alternatives (IIA) assumption. This paper introduces a random utility maximization (RUM) model of tour scheduling called the bivariate multinomial probit (BVMNP). This specification enables correlations across TOD alternatives, both outbound and return (on a tour) and over time slots (in a day). The model is estimated in a Bayesian setting on work-tour data from the San Francisco Bay Area (with 28 time slots). Empirical results suggest that a variety of individual, household, and tour characteristics have reasonable effects on scheduling behavior. For instance, older persons typically pursue work tours at earlier times of day, part-time workers pursue their work tours later, and those with additional activities and tours tend to arrive slightly later and leave much earlier than those

undertaking only a single tour, everything else constant. The model out-performs a comparable MNL, while offering reasonable implications under a variety of road-tolling scenarios.

1. INTRODUCTION

Activity scheduling is a key determinant of temporal variations in travel demand patterns. Yet this dimension of behavior is often greatly simplified in model specifications, particularly as compared to other choice dimensions, such as mode and destination (as noted in Vovsha et al. [2005] and TRB [2007]). As transportation policies become more focused on congestion and demand management (see, e.g., AASHTO 2007), behavioral variations across times of day are increasingly important. This is particularly the case when examining variable-pricing policies (Schofer 2005), which can shift travelers' time-of-day (TOD) choices to off-peak and shoulder periods.

Existing TOD models can be categorized into two broad groups: continuous and discrete. Continuous models generally rely on hazard-based specifications (see, e.g., Wang 1996 and Bhat and Steed 2002, among others) and allow for all times of day. As Bhat and Steed (2002) point out, discrete-choice methods rely on interval boundaries, usually set rather arbitrarily, and discretization always results in a loss in temporal resolution. Nonetheless, such methods typically are based in random utility maximization (RUM) theory, which provides a defensible and econometrically rigorous connection to microeconomic theories of behavior. Current travel demand model systems rely heavily on RUM for other travel choices (such as destination and mode), often integrating such choices in a behaviorally consistent fashion (via logsums, for example [see, e.g., PB Consult 2005]). Moreover, utility models offer a basis for calculating consumer surplus change (see, e.g., de Jong et al. 2007 or Kockelman and Lemp 2009), which is useful for policy and project evaluation (including, for example, environmental justice concerns). In addition, existing continuous methods do not appear capable of consistently incorporating the two-plus timing features of a tour (with possible exception of the continuous cross-nested logit model [Lemp et al. 2010]).

Most of the earliest TOD models used discrete choice methods. For example, in the context of work trip timing, Abkowitz (1981) and Small (1982) used the multinomial logit (MNL), while Chin (1990) turned to the nested logit and Small (1987) developed the ordered generalized extreme value (OGEV) model (to alleviate the independence of irrelevant alternatives (IIA) assumption). In each case, however, the choice spectrum was limited to the AM peak period, rather than the entire day. For large-scale demand systems, temporal variations across the entire day are needed.

Several researchers have modeled TOD choice for the entire day (with MNL and OGEV model specifications), using broad alternative intervals of 3 or more hours (see, e.g., Bhat 1998 and Steed and Bhat 2000). Recent advances in activity-based modeling (where the unit of travel is the tour, rather than the trip) have led to the application of several two-dimensional TOD choice models (since a tour has at least two timing components: outbound and return legs). Vovsha and Bradley (2004), Abou Zeid et al. (2006), and Popuri et al. (2008) modeled tour timing in this way, using joint MNL specifications, each with relatively short (30-minute or 1-hour) alternatives.

While discrete choice methods offer several advantages over continuous methods, most applications of two-dimensional TOD choice models in large-scale travel demand model systems use MNL specifications (see, e.g., PB Consult 2005). This is partly because it offers closed-form choice-probability expressions, but also because it can be estimated with relative ease, even with large numbers of alternatives. When relatively short TOD choice intervals are considered (e.g., 1 hour or 30 minutes), the number of joint choice alternatives grows quickly in two dimensions. However, one would expect error term correlations to exist between alternatives close in time, and the MNL cannot accommodate such correlations.

To address the various issues described above, this paper describes a type of two-way autoregressive error term correlation structure for a multinomial probit (MNP) model of tour TOD choices. The model is estimated on work tour data from the 2000 San Francisco Bay Area Travel Survey (BATS). To accommodate the open-form probability expressions of the MNP, Bayesian estimation techniques are employed. And, to avoid the large number of alternatives that emerge with short time intervals (of 30 minutes, as used here) across two timing dimensions, a bivariate MNP (BVMNP) specification is used, where each commuter chooses exactly two (rather than one) alternatives, one from each timing dimension.

The next section specifies the BVMNP model, while Section 3 discusses the Bayesian estimation procedures used. Section 4 introduces the data set, Section 5 presents analytical results, and Section 6 offers some concluding remarks.

2. BIVARIATE MULTINOMIAL PROBIT FORMULATION

Like the MNL, the multinomial probit (MNP) relies on a latent random utility specification. However, the random error terms follow a normal distribution (rather than a type I extreme value, or Gumbel, distribution). The normality results in open-form expressions for alternative probabilities (unlike the MNL), which is why the MNP has not been utilized to a greater extent in the literature. Thanks to Bayesian and other sophisticated statistical methods, one need not assume error terms are independent and identically distributed with the MNP. In this section, a bivariate MNP (BVMNP) model specification for tour TOD choice is formulated, where the twin variables of interest are a tour's home-to-work arrival time and work-to-home departure time.

2.1 Random Utility Framework and Model Specification

Each alternative in an MNP model has a (latent) random utility, and the decision-maker always chooses the alternative offering the greatest underlying utility value. Since the MNP model developed here seeks to reflect two-dimensional travel-timing, the choice context needs special attention. One reasonable way to approach the problem is to consider it in a single dimension. Instead of choosing tour arrival times and tour return times, one may assume that individuals jointly choose tour arrival and return times, and the analyst need only consider a single choice dimension. For instance, consider the following joint utility specification:

$$U(t_a, t_r) = V_1(t_a) + V_2(t_r) + V_3(t_r - t_a) + \varepsilon_{ar} \tag{1}$$

Here, V_1 is the systematic utility component related to arrival time t_a, V_2 is the component related to return time t_r, and V_3 is the component related to duration $t_r - t_a$. A key difficulty with this approach is that one is usually interested in rather small time intervals as alternatives; and, in two dimensions, the number of alternatives can become quite large. For instance, if 30-minute intervals are used, one has 1,176 alternatives. For an MNP model, this produces a covariance matrix of size 1,176 x 1,176, presenting a number of computational difficulties in model estimation. With this in mind, a bivariate multinomial probit (BVMNP) model is developed here, where tour arrival time represents one choice dimension and tour return time represents another. While the BVMNP model has been used in previous studies (see, e.g., Golob and Regan 2002 and Zhang et al. 2008), no previous work has investigated choice contexts with more than three or four alternatives. In addition, the estimation procedure used here varies from traditional methods, in order to accommodate the large number of alternatives. In this bivariate context, one must specify two separate utility functions (one for tour arrival and another for tour return), as follows:

$$U_{aj} = V_{aj} + \varepsilon_{aj} \tag{2}$$
$$U_{rl} = V_{rl} + \varepsilon_{rl} \tag{3}$$

Here, U_{aj} and U_{rl} denote latent utilities for arrival and return time alternatives *j* and l, V_{aj} and V_{rl} are systematic utility components, and ε_{aj} and ε_{rl} are random error components. The set of arrival time alternatives is identical to the set of return time alternatives, with arrival time alternatives indexed by $j = 1, \ldots, \mathcal{J}$ and return time alternatives indexed by $k = 1, \ldots, \mathcal{J}$. While this specification does not allow for a utility component specifically related to tour/activity duration, it does significantly reduce the number of choice alternatives. For instance, if time-of-day is modeled in 30-minute intervals over the 24-hour day period (as it is here), this results in 96 alternatives (and utility values), rather than the 1,176 needed for the joint choice model.

2.2 Error Correlation Structure

Since one cannot reasonably assume independence of alternatives, the correlation structure of the error components deserves some attention. While it is theoretically feasible to estimate the entire covariance matrix without imposing any pre-specified structure, there is a clear ordering of alternatives, which evokes certain expectations for covariance properties. With this in mind, a specific structure is imposed here.

One can imagine a variety of correlation structures. Here, a pseudo-AR specification of covariance components is pursued. Components of the covariance matrix are formulated directly, with the upper left and lower right quadrants taking on forms *similar to* a typical AR1 process (though it is worth noting that the formulation cannot be directly interpreted as an AR1 process). Off-diagonal quadrants are formulated slightly different, due to bivariate interactions between work arrival and return times, though covariance components appear similar to those of an AR1 process. The covariance matrix is specified as follows:

$$\Sigma = \begin{bmatrix} \mathcal{W}_a & \mathcal{C} \\ \mathcal{C}' & \mathcal{W}_r \end{bmatrix} \tag{4}$$
$$\mathcal{W}_a = \begin{bmatrix} w_{a11} & \cdots & w_{a1\mathcal{J}} \\ \vdots & \ddots & \vdots \\ w_{a\mathcal{J}1} & \cdots & w_{a\mathcal{J}\mathcal{J}} \end{bmatrix}, w_{apq} = \lambda_a^{|t_p - t_q|}\sigma_a^2 \tag{5}$$
$$\mathcal{W}_r = \begin{bmatrix} w_{r11} & \cdots & w_{r1\mathcal{J}} \\ \vdots & \ddots & \vdots \\ w_{r\mathcal{J}1} & \cdots & w_{r\mathcal{J}\mathcal{J}} \end{bmatrix}, w_{rpq} = \lambda_r^{|t_p - t_q|}\sigma_r^2 \tag{6}$$
$$\mathcal{C} = \begin{bmatrix} c_{11} & \cdots & c_{1\mathcal{J}} \\ \vdots & \ddots & \vdots \\ c_{\mathcal{J}1} & \cdots & c_{\mathcal{J}\mathcal{J}} \end{bmatrix}, c_{pq} = \begin{cases} 0 & \text{for } p > q \\ \sigma_a\sigma_r\lambda_d^{|(t_q - t_p) - (\mu_1 + \mu_2 t_p)| + 1} & \text{for } p \le q \end{cases} \tag{7}$$

Here, $\mathcal{W}_a$ and $\mathcal{W}_r$ are covariance matrices defining the error structure within the arrival time dimension and within the return time dimension, and $\mathcal{C}$ and $\mathcal{C}'$ are the covariance matrices defining the error structure across arrival and return times. Arrival time- and return time-specific variances are denoted by σ_a^2 and σ_r^2, respectively, λ_a and λ_r are correlation coefficients of arrival and return utility components separated by 1 hour, and $\lambda_d^{|(t_q-t_p)-(\mu_1+\mu_2 t_p)|+1}$ is the correlation coefficient between arrival and return utility components. One last item needing attention here is the role of μ_1 and μ_2 (which define the "baseline" duration on which elements of C are based). Essentially, the model posits that some activity duration may be highly desired (e.g., 8 hours for full-time workers arriving at work around 9 am), and this term allows correlations across arrival and return time utilities to be highest for such durations. The "baseline" duration is equal to some constant, μ_1, plus an additional term that varies over the work arrival time, $\mu_2 t_p$. Of course, it is not reasonable to view μ_1 and μ_2 as two fixed values, since one expects differences across individuals or classes of individuals. This is particularly important since work duration is not reflected in the systematic utility equations. Here, μ_1 and μ_2 are taken to be two separate parameters each, two for full-time workers making no additional tours during the day ($\mu_{1,\text{full}}$ and $\mu_{2,\text{full}}$) and two for part-time workers and/or those making additional tours ($\mu_{1,\text{part}}$ and $\mu_{2,\text{part}}$), adding a layer of observed heterogeneity to the model.[1] Of course, one may expect preferred durations to vary with other traveler attributes, and one could control for those too. The reason μ_1 and μ_2 are differentiated here only between full-time workers (with no additional tours) and part-time workers (and/or those making additional tours) is that this distinction seems most important. If μ_1 and μ_2 differed for each individual, Σ would also differ for each individual, requiring computation of distinct Σ's for each observation, which can be computationally expensive. By allowing μ_1 and μ_2 to vary over only two traveler groups, the estimation process is streamlined, with only two covariance matrices, Σ_{full} and Σ_{part}, thus facilitating demonstration of the model here.

In this paper, 30-minute time intervals serve as the choice alternatives. Since there are very few individuals choosing times very early in the day and very late in the day (for both arrival and return choice dimensions), boundary alternatives were needed, essentially grouping many 30-minute alternatives into a single alternative. Since these boundary intervals may exhibit very different properties than non-boundary alternatives, their variances and correlations adjusted, as described below.

First, each of the four boundary alternatives' error terms is allowed to have its own variance parameter, denoted by σ_{a1}^2, $\sigma_{a\mathcal{J}}^2$, σ_{r1}^2, and $\sigma_{r\mathcal{J}}^2$. Second, correlation coefficients between utility components are assumed to be inversely related to each alternative's interval size. More specifically, these parameters are assumed to vary across each pair of alternatives as follows:

$$\text{Corr}(p_a, q_a) = \lambda_a^{|t_{p_a}-t_{q_a}|}\left(\frac{\text{size}_p}{\text{size}_{\text{def}}}\right)^{-\tau_{ar}}\left(\frac{\text{size}_q}{\text{size}_{\text{def}}}\right)^{-\tau_{ar}} \quad (8)$$

$$\text{Corr}(p_r, q_r) = \lambda_r^{|t_{p_r}-t_{q_r}|}\left(\frac{\text{size}_p}{\text{size}_{\text{def}}}\right)^{-\tau_{ar}}\left(\frac{\text{size}_q}{\text{size}_{\text{def}}}\right)^{-\tau_{ar}} \quad (9)$$

$$\text{Corr}(p_a, q_r) = \lambda_d^{|t_{q_r}-t_{p_a}-(\mu_1+\mu_2 t_{p_a})+1|}\left(\frac{\text{size}_p}{\text{size}_{\text{def}}}\right)^{-\tau_d}\left(\frac{\text{size}_q}{\text{size}_{\text{def}}}\right)^{-\tau_d} \quad (10)$$

[1] The distinction here was chosen because average travel durations for full-time workers in the data sample were found to be about 8.2 hours, while average travel durations for both part-time workers and those making additional tours during the day were found to be about 6 hours.

Here, size_{def} is 30 minutes (the default interval size), and size_p and size_q are the interval sizes of alternatives *p* and *q* (measured in minutes). λ_a, λ_r, and λ_d are the same as described above, and τ_{ar} and τ_d are two new (non-negative) parameters to be estimated. The assumption of non-negativity presumes that correlations between boundary alternatives and other alternatives are smaller than those across non-boundary alternatives. It should be clear that the size terms only come into play when one or both of the alternatives are boundary alternatives. Moreover, these size terms reduce to the original model specification if τ_{ar} or τ_d is zero. Finally, it is worth noting that τ_{ar} affects both the correlations across arrival alternatives and those across departure alternatives. Thus, it is assumed that the way in which the interval size affects correlation patterns across arrival times is the same as the way it affects correlation patterns across return times.

2.3 Systematic Utility Construction

Finally, the systematic utility specifications for arrival time and return time utilities take the forms shown in equations 3.2 and 3.3. Since each alternative represents a discrete time interval, the *t* in the utility equation is taken to be the midpoint of the time interval. For notational convenience, the systematic utilities for arrival time and return time alternatives are rewritten as follows:

$$V_{i,aj} = \mathcal{X}_{ij}\beta_a + \sum_{p=1}^{\mathcal{P}} \eta_{ap} g_{i,ap}(t_j) \tag{11}$$
$$V_{i,rl} = \mathcal{X}_{il}\beta_r + \sum_{p=1}^{\mathcal{P}} \eta_{rp} g_{i,rp}(t_l) \tag{12}$$

Here, $g_{i,ap}(t_j)$ and $g_{i,rp}(t_l)$ represent network characteristics of type *p* (such as travel time and reliability) for arrival and return time intervals t_j and t_l; and β_a, β_r, η_a, and η_r are parameters to be estimated. Similar to MNL and other choice models, covariates that do not vary over time alternatives (e.g., an individual's gender or age) cannot be introduced in the normal way. One can imagine any number of continuous forms to use in this context. Here, $\mathcal{X}_{ij}$ and $\mathcal{X}_{il}$ represent row vectors of individual-specific attributes interacted with cyclical functions (like the utility formulations of Abou Zeid et al. 2006 and Popuri et al. 2008). That is, $\mathcal{X}_{ij}$ and $\mathcal{X}_{il}$ have the following forms:

$$\mathcal{X}_{ij} = \begin{bmatrix} X_{i1}\sin(2\pi t_j/24) \\ X_{i1}\sin(4\pi t_j/24) \\ \vdots \\ X_{i1}\sin(2Q_1\pi t_j/24) \\ X_{i1}\cos(2\pi t_j/24) \\ \vdots \\ X_{i1}\cos(2Q_1\pi t_j/24) \\ X_{i2}\sin(2\pi t_j/24) \\ \vdots \\ X_{i2}\cos(2Q_2\pi t_j/24) \\ \vdots \\ X_{iK}\sin(2\pi t_j/24) \\ \vdots \\ X_{iK}\cos(2Q_K\pi t_j/24) \end{bmatrix}', \quad \mathcal{X}_{il} = \begin{bmatrix} X_{i1}\sin(2\pi t_l/24) \\ X_{i1}\sin(4\pi t_l/24) \\ \vdots \\ X_{i1}\sin(2Q_1\pi t_l/24) \\ X_{i1}\cos(2\pi t_l/24) \\ \vdots \\ X_{i1}\cos(2Q_1\pi t_l/24) \\ X_{i2}\sin(2\pi t_l/24) \\ \vdots \\ X_{i2}\cos(2Q_2\pi t_l/24) \\ \vdots \\ X_{iK}\sin(2\pi t_l/24) \\ \vdots \\ X_{iK}\cos(2Q_K\pi t_l/24) \end{bmatrix}',$$

$$\mathcal{X}_i = \begin{bmatrix} \mathcal{X}_{i1} & 0 & g_{i,a}(t_1) & 0 \\ \vdots & 0 & \vdots & 0 \\ \mathcal{X}_{i\mathcal{J}} & 0 & g_{i,a}(t_{\mathcal{J}}) & 0 \\ 0 & \mathcal{X}_{i1} & 0 & g_{i,r}(t_1) \\ 0 & \vdots & 0 & \vdots \\ 0 & \mathcal{X}_{i\mathcal{J}} & 0 & g_{i,r}(t_{\mathcal{J}}) \end{bmatrix}$$

$$\begin{bmatrix} V_{i,a} \\ V_{i,r} \end{bmatrix} = \mathcal{X}_i \begin{bmatrix} \beta_a \\ \beta_r \\ \eta_a \\ \eta_r \end{bmatrix} \tag{13}$$

The number of individual-specific attributes is K, with each individual attribute interacted with $2Q_k$ cyclical functions (Q_k for sine functions and Q_k for cosine functions). Some covariates may be interacted with fewer than $2Q_k$ cyclical functions by restricting the applicable elements of β_a and β_r to be zero. There are a couple of reasons for selecting this utility form. First, it allows utility to take on a rich assortment of shapes, including multimodal ones. In addition (and as pointed out be Abou Zeid et al. 2006 and Popuri et al. 2008), 24 hours is a multiple of each cyclical function's period, which offers day-to-day consistency in the utility function (e.g., utilities at 0 and 24 hours are identical). And by construction, the systematic utilities are linear in unknown parameters.

3. BVMNP PARAMETER ESTIMATION

Estimation of the BVMNP model can be performed via MCMC simulation. Bayesian techniques are particularly well-suited for estimation of the BVMNP (or any MNP for that matter) since classical methods generally rely on simulated maximum likelihood estimation (MSLE) to avoid numerical evaluation of multi-dimensional integrals involved in the likelihood (McFadden 1989 and Geweke et al. 1994).

In the standard Bayesian construction of the MNP model (see, e.g., Albert and Chib 1993, McCulloch and Rossi 1994, and Zhang et al. 2008, among others), one need not evaluate choice probabilities at all. For the MNP model, the dependent variable, Y_i, can take on values 1, 2, …, $\mathcal{J}$, where Y_i's value simply indexes the chosen alternative. With the latent random utility specification of the model, the probability of Y_i taking on a value *q* is given by the following:

$$P(Y_i = q) = P(U_{iq} \geq \max_{p \in \mathcal{J}} U_{ip}) \tag{14}$$

In other words, the choice probability of alternative *q* is equivalent to the probability that the latent utility associated with alternative *q* is the maximum utility value. Here, U_i is treated as a random (nuisance) parameter to be estimated and is normally distributed (under the MNP model specification), with mean given by the systematic utility, V_i, and variance given by Σ. For the BVMNP model, Y_i is simply taken to be bivariate, with joint choice probability of arrival time q_1 and return time q_2 given by the following:

$$P(Y_i = [q_1, q_2]) = P(U_{i,aq_1} \geq \max_{p \in \mathcal{J}} U_{i,ap} \cap U_{i,rq_2} \geq \max_{p \in \mathcal{J}} U_{i,rp}) \tag{15}$$

The joint choice probability of arrival time q_1 and return time q_2 is equivalent to the probability that the latent utility associated with arrival time alternative q_1 is the maximum utility across all arrival time alternatives and that the latent utility associated with return time alternative q_2 is the maximum utility across all return time alternatives. Bayesian estimation proceeds via a three-step Gibbs sampler as follows:

Step 1: Draw $U_i \mid V_i, \Sigma_i, X_i, Y_i \; \forall \; i$
Step 2: Draw $\lambda, \mu, \sigma, \tau \mid U_i, V_i, X_i, Y_i \; \forall \; i$
Step 3: Draw $\beta_a, \beta_r, \eta \mid U_i, \Sigma_i, X_i, Y_i \; \forall \; i$

The Gibbs sampler does not generate draws for ${\sigma_a}^2$ or ${\sigma_r}^2$ (the variances of non-boundary alternatives) here. It is well known that the MNP requires one element of Σ to be fixed for identification purposes (see, e.g., McCulloch and Rossi 1994). However, with the BVMNP, one element of Σ must be fixed for each nominal measure (Zhang et al. 2008). Thus, ${\sigma_a}^2$ or ${\sigma_r}^2$ are fixed at 1 for identification purposes, though boundary alternative variances are estimated.

In step 1, a normal random walk Metropolis-Hastings (MH) step (see, e.g., Gamerman and Lopes 2006) is used to draw an individual's utility values simultaneously. The proposal density for the MH step is a multivariate normal, with mean equal to the current utility values, and covariance given by $\mathscr{k}\Sigma_i$. Here, Σ_i is the utility covariance matrix for individual *i*, computed from the current values of covariance parameters, and $\mathscr{k}$ is a deflation factor to increase the probability of proposal acceptance. The deflation factor was set to $\mathscr{k} = 0.05$, after calibrating the parameter to achieve approximately 25% proposal acceptance. It was found that this MH algorithm is much more computationally stable than the typical Gibbs sampling algorithm, and it reduces computation time per iteration by nearly one half. Unfortunately, since utility values are more restricted in their movements from one iteration to the next, the algorithm is slow to converge.

In the second step of the Gibbs sampler, a draw of the covariance matrix parameters is generated. Priors on λ_a, λ_r, and λ_d are specified to be independent uniform distributions over the interval from 0 to 1, reflecting a belief that there should be positive correlation across alternative utilities. Priors on $\mu_{1,\text{full}}$ and $\mu_{1,\text{part}}$ are specified to be independent normal distributions with means of 9 and 6 (hours), respectively, and variances of 2 each, while priors on $\mu_{2,\text{full}}$ and $\mu_{2,\text{part}}$ are specified to be normal distributions, each with means and variances of 0 and 1, respectively. Independent gamma priors are employed for σ_{a1}^2, $\sigma_{a\mathcal{J}}^2$, σ_{r1}^2, and $\sigma_{r\mathcal{J}}^2$, each with shape and scale parameters of 2 and 1, respectively, while τ_{ar} and τ_d are assumed to follow independent gammas with shape and scale parameters of 1 each in the prior. The gamma prior restricts these parameters to be positive. Thus, the full conditional posterior distribution of the variance parameters can be written as follows[2]:

[2] When conditioned on U_i and V_i, λ, μ, σ_{a1}^2, $\sigma_{a\mathcal{J}}^2$, σ_{r1}^2, $\sigma_{r\mathcal{J}}^2$, and τ are independent of X_i and Y_i.

$$p(\lambda, \mu, \sigma, \tau \mid U_i, V_i \, \forall \, i) \propto H_1 H_2 \qquad (16)$$

where,

$$H_1 = |\Sigma_{\text{full}}|^{-n_{\text{full}}/2} \left|\Sigma_{\text{part}}\right|^{-n_{\text{part}}/2} \exp\left(-\frac{1}{2}\left(\sum_i (U_i - V_i)' \Sigma_i^{-1} (U_i - V_i)\right)\right)$$

$$H_2 = \left(\sigma_{a1}^2 \sigma_{a\mathcal{J}}^2 \sigma_{r1}^2 \sigma_{r\mathcal{J}}^2\right) \exp\left(-\frac{1}{2}\left(\sum_{q=1}^{4} \frac{\left(\mu_q - \bar{\mu}_q\right)^2}{\sigma_q^2}\right) - \left(\sigma_{a1}^2 + \sigma_{a\mathcal{J}}^2 + \sigma_{r1}^2 + \sigma_{r\mathcal{J}}^2 + \tau_{ar} + \tau_d\right)\right)$$

$$\sum_{q=1}^{4} \frac{\left(\mu_q - \bar{\mu}_q\right)^2}{\sigma_q^2} = \frac{\left(\mu_{1,\text{full}} - 9\right)^2}{4} + \frac{\left(\mu_{1,\text{part}} - 6\right)^2}{4} + \left(\mu_{2,\text{full}}\right)^2 + \left(\mu_{2,\text{part}}\right)^2$$

Since the density here is not in any standard form (with respect to the parameters), a MH step is used to draw these parameters. The proposal density is assumed to be normal, with mean given by the current draw of the parameters (i.e., a normal random walk) and variance initially taken to be very small and updated during the estimation process to aid in generating good proposals (see, e.g., Holden et al. 2009).

In the last step of the Gibbs sampler, a draw of β_a, β_r, and η is generated from the full conditional posterior distribution. For notational convenience, write $\beta = [\beta_a{}', \beta_r{}', \eta']'$. Here, the prior for these parameters is chosen to be multivariate normal with mean $\bar{\beta}$ and covariance matrix Σ_β. Thus, the full conditional posterior is proportional to the following[3]:

$$p(\beta \mid U_i, \Sigma_i, X_i \, \forall \, i) \propto$$
$$\exp\left(-\frac{1}{2}\left(\left(\beta - \bar{\beta}\right)' \Sigma_\beta{}^{-1} \left(\beta - \bar{\beta}\right) + \sum_i (U_i - \mathcal{X}_i \beta)' \Sigma_i{}^{-1} (U_i - \mathcal{X}_i \beta)\right)\right) \qquad (17)$$

Suppose Ω and Λ are given by the following:

$$\Lambda = \Sigma_\beta{}^{-1} \bar{\beta} + \sum_i \mathcal{X}_i{}' \Sigma_i{}^{-1} U_i \qquad (18)$$
$$\Omega = \left(\Sigma_\beta{}^{-1} + \sum_i \mathcal{X}_i{}' \Sigma_i{}^{-1} \mathcal{X}_i\right)^{-1} \qquad (19)$$

Expression 17 suggests that β is proportional to a multivariate normal distribution with mean given by $\Omega\Lambda$ and covariance matrix Ω. Thus, β is drawn from a multivariate normal distribution. Here, vague prior parameters are specified, with $\bar{\beta}$ taken to be a vector of zeroes, off-diagonal elements of Σ_β taken to be zeroes, and diagonal elements of Σ_β set to be very large. Since each network variable should affect utilities negatively, the normal prior distributions for η were truncated above at zero. That is all that is needed to generate the MCMC draws for this BVMNP model.

4. DATA DESCRIPTION

The data used here come from the 2000 San Francisco Bay Area Travel Survey (BATS). The survey collected travel information for roughly 17,000 households over a 2-day period. The observational unit is the travel tour (over 100,000 recorded tours), with network attributes provided for each of five TOD periods and seven modes. In addition, each record was coded with a variety of demographic and travel information. Since the analysis is focused on work-related travel, the sampling frame was restricted to the first home-based work tour made on a weekday for an individual (over the 48-hour survey period), limiting the sample to about 18,000 tours.

---

[3] Note that conditional on Σ_{full}, Σ_{part}, U_i, and X_i, β is independent of Y_i.

Model estimation is very computationally burdensome due to the large number of utility values that must be drawn for each individual (one for each alternative). So, an n = 997 random sample of tours was used in model estimation.

Since time-varying network variables (as shown in equations 11 and 12) are not contained in the data (though the data does contain network information across 5 broad time-of-day periods), regression equations were developed to impute average travel time and its variance. A similar methodology to that of Popuri et al. (2008) was used to this end. For brevity, the results of these models are not presented here, though it should be noted that the models appropriately predict travel times and variances to be highest during typical AM and PM peak periods. Lemp (2009) provides more details on these results.

5. EMPIRICAL RESULTS

Given the nature of the full conditional distribution for the random utilities, the Gibbs sampler was very slow to converge. It is well-known that in high-dimensional utility choice models convergence could be an issue due to poor mixing; however, researchers supplement this by ensuring that the resulting estimates are contextually reasonable, which is the approach we adopted. (For details, see Rossi et al. 2005, and the many references therein.) It is important to note that maximum likelihood estimates are also difficult to obtain in these models. There is simply no guarantee of obtaining global maxima when the likelihood surface is as complicated (i.e., highly multimodal) as the one in this paper.

The utility function utilizes eight individual-specific variables plus a constant. These variables include an indicator for males, age of the individual, an indicator for part-time workers, an indicator for high income households (over $75,000 per year), household size, the number of tours undertaken by the individual over the entire day (excluding the modeled tour), travel distance to the destination, and a variable indicating whether the destination zone is coded as central business district (CBD). Table 1 presents the estimation results.

Table 1: Model Estimation Results

Variable		Arrival-Specific Utility Mean Estimate	Arrival-Specific Utility 95% Interval	Return-Specific Utility Mean Estimate	Return-Specific Utility 95% Interval
LOS Variables	Travel Time (min)	-0.0038	(-0.0092, -0.0037)	-0.0058	(-0.0157, -0.0053)
	Travel Time Variance (min^2)	-0.0011	(-0.0026, -0.0011)	-0.0005	(-0.0016, -0.0005)
	Cost ($)	-0.0465	(-0.1369, -0.0342)	-0.0554	(-0.1577, -0.0467)
Constant Interactions	Sin(2*pi*t/24)	0.3822	(-0.1965, 0.4000)	-2.947	(-3.982, -2.923)
	Sin(4*pi*t/24)	-0.9215	(-1.378, -0.9134)	-0.5503	(-1.074, -0.5524)
	Cos(2*pi*t/24)	-1.447	(-2.502, -1.427)	-0.2693	(-1.035, -0.2668)
	Cos(4*pi*t/24)	-0.2851	(-0.8649, -0.2742)	0.1037	(-0.3165, 0.1049)
Male Indicator Interactions	Sin(2*pi*t/24)	0.3037	(0.0426, 0.3086)	0.2977	(-0.1330, 0.3036)
	Sin(4*pi*t/24)	0.3416	(0.1454, 0.3418)	0.1131	(-0.1049, 0.1127)
	Cos(2*pi*t/24)	0.1152	(-0.3719, 0.1032)	0.3879	(0.0652, 0.3855)
	Cos(4*pi*t/24)	0.0955	(-0.1537, 0.0889)	-0.0759	(-0.2826, -0.0763)
Age Interactions	Sin(2*pi*t/24)	0.0026	(-0.0088, 0.0025)	0.0108	(-0.0118, 0.0116)
	Sin(4*pi*t/24)	0.0024	(-0.0064, 0.0020)	0.0144	(0.0030, 0.0146)
	Sin(6*pi*t/24)	-0.0063	(-0.0093, -0.0062)	0.0127	(0.0095, 0.0128)
	Cos(2*pi*t/24)	-0.0034	(-0.0307, -0.0032)	-0.0062	(-0.0188, -0.0062)
	Cos(4*pi*t/24)	0.0013	(-0.0123, 0.0012)	0.0093	(-0.0007, 0.0093)
	Cos(6*pi*t/24)	0.0084	(0.0041, 0.0086)	0.0011	(-0.0021, 0.0012)
Part-Time Indicator Interactions	Sin(2*pi*t/24)	-1.232	(-1.948, -1.217)	-0.2844	(-1.391, -0.2618)
	Sin(4*pi*t/24)	-0.2684	(-0.8433, -0.2492)	0.5054	(-0.1222, 0.5034)
	Cos(2*pi*t/24)	0.1918	(-1.323, 0.2485)	1.307	(0.4414, 1.2943)
	Cos(4*pi*t/24)	0.0105	(-0.7837, 0.0346)	0.3018	(-0.3227, 0.3025)
High Inc. HH Indicator Interactions	Sin(2*pi*t/24)	-0.0796	(-0.3313, -0.0799)	-0.5753	(-1.002, -0.5713)
	Sin(4*pi*t/24)	-0.1613	(-0.3490, -0.1654)	-0.2843	(-0.4918, -0.2812)
	Cos(2*pi*t/24)	-0.0900	(-0.5749, -0.0860)	-0.3883	(-0.6743, -0.3867)
	Cos(4*pi*t/24)	0.0042	(-0.2509, 0.0086)	0.1498	(-0.0512, 0.1497)
HH Size Interactions	Sin(2*pi*t/24)	-0.0429	(-0.1264, -0.0414)	0.0692	(-0.0649, 0.0674)
	Sin(4*pi*t/24)	0.0102	(-0.0561, 0.0107)	0.0487	(-0.0195, 0.0476)
	Cos(2*pi*t/24)	0.1341	(-0.0046, 0.1338)	0.0721	(-0.0194, 0.0715)
	Cos(4*pi*t/24)	-0.0076	(-0.0854, -0.0075)	0.0400	(-0.0346, 0.0417)
No. Other Tours Interactions	Sin(2*pi*t/24)	-0.3042	(-0.4990, -0.3026)	0.7399	(0.4726, 0.7449)
	Sin(4*pi*t/24)	-0.0425	(-0.1789, -0.0423)	0.1729	(0.0112, 0.1751)
	Cos(2*pi*t/24)	-0.0825	(-0.3842, -0.0778)	-0.1693	(-0.3740, -0.1698)
	Cos(4*pi*t/24)	-0.0384	(-0.1981, -0.0420)	-0.0595	(-0.1980, -0.0554)
Travel Distance Interactions	Sin(2*pi*t/24)	0.0175	(0.0059, 0.0172)	-0.0089	(-0.0307, -0.0091)
	Sin(4*pi*t/24)	0.0032	(-0.0053, 0.0033)	0.0023	(-0.0064, 0.0023)
	Cos(2*pi*t/24)	0.0381	(0.0208, 0.0386)	-0.0028	(-0.0141, -0.0026)
	Cos(4*pi*t/24)	0.0225	(0.0125, 0.0226)	0.0050	(-0.0047, 0.0051)
CBD Dest. Indicator Interactions	Sin(2*pi*t/24)	-0.2002	(-0.6283, -0.1903)	-1.290	(-1.998, -1.299)
	Sin(4*pi*t/24)	-0.1589	(-0.4452, -0.1633)	-0.4269	(-0.8165, -0.4136)
	Cos(2*pi*t/24)	0.6770	(-0.0066, 0.6754)	-0.3724	(-1.009, -0.3411)
	Cos(4*pi*t/24)	0.4269	(0.0048, 0.4291)	0.2123	(-0.1545, 0.2189)

Table 1 (Cont'd): Model Estimation Results

Variable	Mean Estimate	95% Interval
σ_{a1}^2	1.278	(0.8876, 1.225)
$\sigma_{a\mathcal{J}}^2$	1.839	(0.2227, 1.581)
σ_{r1}^2	4.459	(3.153, 4.359)
$\sigma_{r\mathcal{J}}^2$	1.784	(0.2363, 1.588)
λ_a	0.7442	(0.7257, 0.7446)
λ_r	0.7335	(0.7127, 0.7343)
λ_d	0.8039	(0.7971, 0.8046)
τ_{ar}	0.0512	(0.0014, 0.0483)
τ_d	0.0295	(0.0018, 0.0271)
$\mu_{1,\text{full}}$	9.892	(9.750, 9.878)
$\mu_{1,\text{part}}$	10.17	(9.798, 10.13)
$\mu_{2,\text{full}}$	-0.0814	(-0.1064, -0.0785)
$\mu_{2,\text{part}}$	-0.0881	(-0.1363, -0.0817)

Implied median values of travel time (VOTTs) for the model are \$5.97/hour and \$6.76/hour for arrival and return journeys (with much higher mean VOTTs, at \$19.85/hour and \$16.07/hour, respectively). Implied median values of reliability (VORs) are \$9.95 and \$5.81 per hour of travel time's standard deviation on the home-to-work journey and work-to-home journey, respectively. (Mean VOR estimates are \$13.47 and \$7.17 per hour.)

Of course, the effects of these variables also depends on the variability in them. For arrival time choice, it turns out that reliability is more practically significant for the home-to-work journey than is average travel time. This makes sense, since many workers are somewhat constrained in their working hours. There also is an incentive for leaving a buffer period, to ensure arrival at or before work is scheduled to begin. Since it may be more acceptable to arrive 10 minutes early than to arrive 10 minutes late, many may depart from home earlier, rather than later.[4] However, on the return journey, average travel times are more practically significant. One should note, however, that these VOTT and VOR estimates are context-specific (for activity scheduling), and may not be valid for other choice contexts (e.g., mode or route choice) or under different network variable imputation assumptions.

5.1 Effects of Individuals' Characteristics

To understand the effects of time-invariant covariates, the "average" sample individual was assumed (i.e., the sample-average value of each covariate was used); and covariate values were varied one at a time. Figure 1 shows density profiles to illustrate the effect each covariate has on predictive densities of the arrival time at work. Note that time-varying variable effects were omitted here, to highlight the effects of each attribute.

[4] Small et al. (1999) estimated the marginal costs of early arrival to rise with time, from about \$0.028/min at 5 minutes early, up to about \$0.128/min at 15 minutes early. The marginal cost of late arrival, however, was estimated to be 2.5 to 11 times greater, at \$0.31/min.

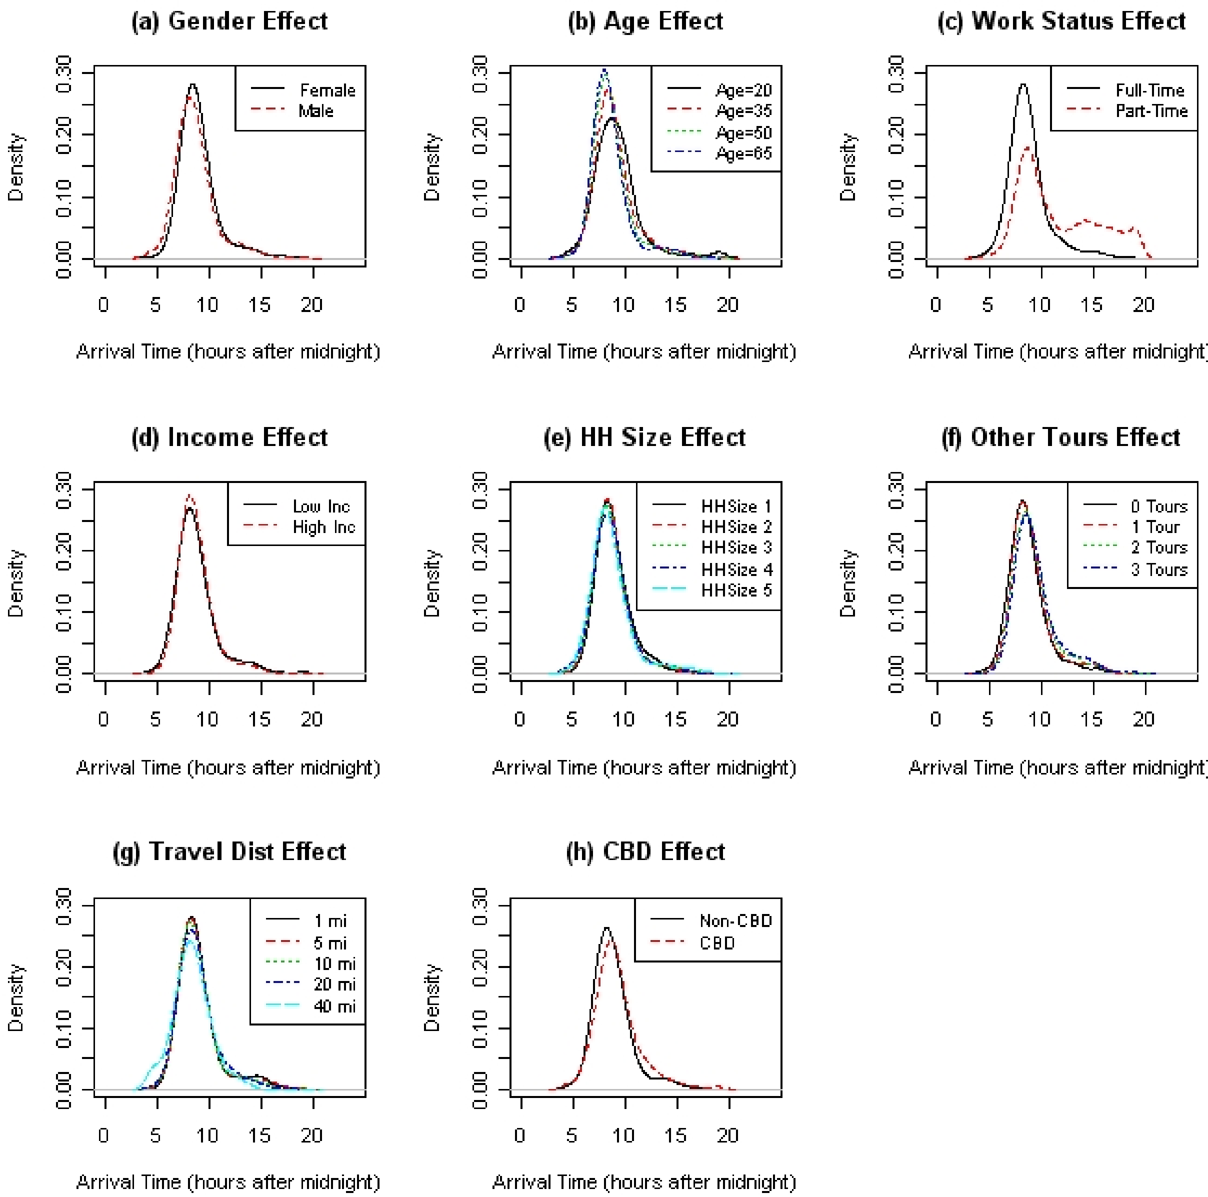

Figure 1: Predicted Arrival Time Density Profiles for Individuals with Different Attributes

Predictive densities (of systematic utility for work arrival time choice) peak near the AM peak period, as expected. Effects of gender, household income, household size, travel distance, and CBD appear rather small (Figures 1a, 1d, 1e, 1g, and 1h), though in line with expectations. For instance, males, those from larger households, and those with longer travel distances tend to arrive earlier, all else equal. The household size effect can, in some sense, be viewed as a proxy for the number of children, and those with children may have obligations such as dropping off children in the morning. Those with longer travel distances may be arriving earlier on average because they need to leave extra buffer time to be sure to arrive at work on time (since longer distances are associated with larger travel time variances) and/or get started earlier in order to arrive home at a reasonable hour, at the end of a long work day. Individuals from high-income households may have more flexibility in work start times, thus explaining why such individuals tend to arrive a bit later, ceteris paribus. Not surprisingly, older individuals (Figure

1b) tend to arrive earlier for work, while part-time worker status (Figure 1c) and added tours (Figure 1f) both tend to make later departure times more desirable. These results seem reasonable, since part-time workers often do not work full 8-hour days and those with additional engagements are less likely to be working typical hours.

Figure 2 shows return time predictive densities for variations in each variable. Males are predicted to return slightly later than females (Figure 2a), and arrive slightly earlier (Figure 1a), on average, so their work durations tend to be slightly longer (by about 20 to 30 minutes). Maybe men are more likely to work overtime than women, and/or females are more likely to have other responsibilities, such as dropping children off at and picking them up from school and child care facilities.

Older individuals are predicted to return earlier than younger ones (Figure 2b), which is not so surprising given that they are predicted to arrive earlier, on average (Figure 1b). Interestingly, part-time workers' return time profiles (Figure 2c) mimic their arrival time profiles (Figure 1c), with return times shifted to later hours, of course. This is very reasonable considering that such workers may have very different work scheduling constraints, as compared to full-time workers. While household size has little effect on arrival times (Figure 1e) and does not appear to shift return times (Figure 2e), it appears to add uncertainty in return time choice. The presence of additional tours has very important effects on a worker's return time (Figure 2f). In particular, such individuals are predicted to return from work much earlier (about 45 to 90 minutes per additional tour, on average), as the number of such additional tours increases. Since these workers obviously have other scheduling considerations for the day, this seems very reasonable. As with travel distance's limited effect on arrival times, its effect on return times is not substantial (Figure 2g). Finally, the effect of traveling to a workplace in the central business district (CBD) is to push return times later in the day (Figure 2h), consistent with these workers' later arrival times (Figure 1h). Perhaps CBD workers enjoy occupations with later start times. Overall, covariates' effects appear reasonable, though somewhat limited in magnitude.

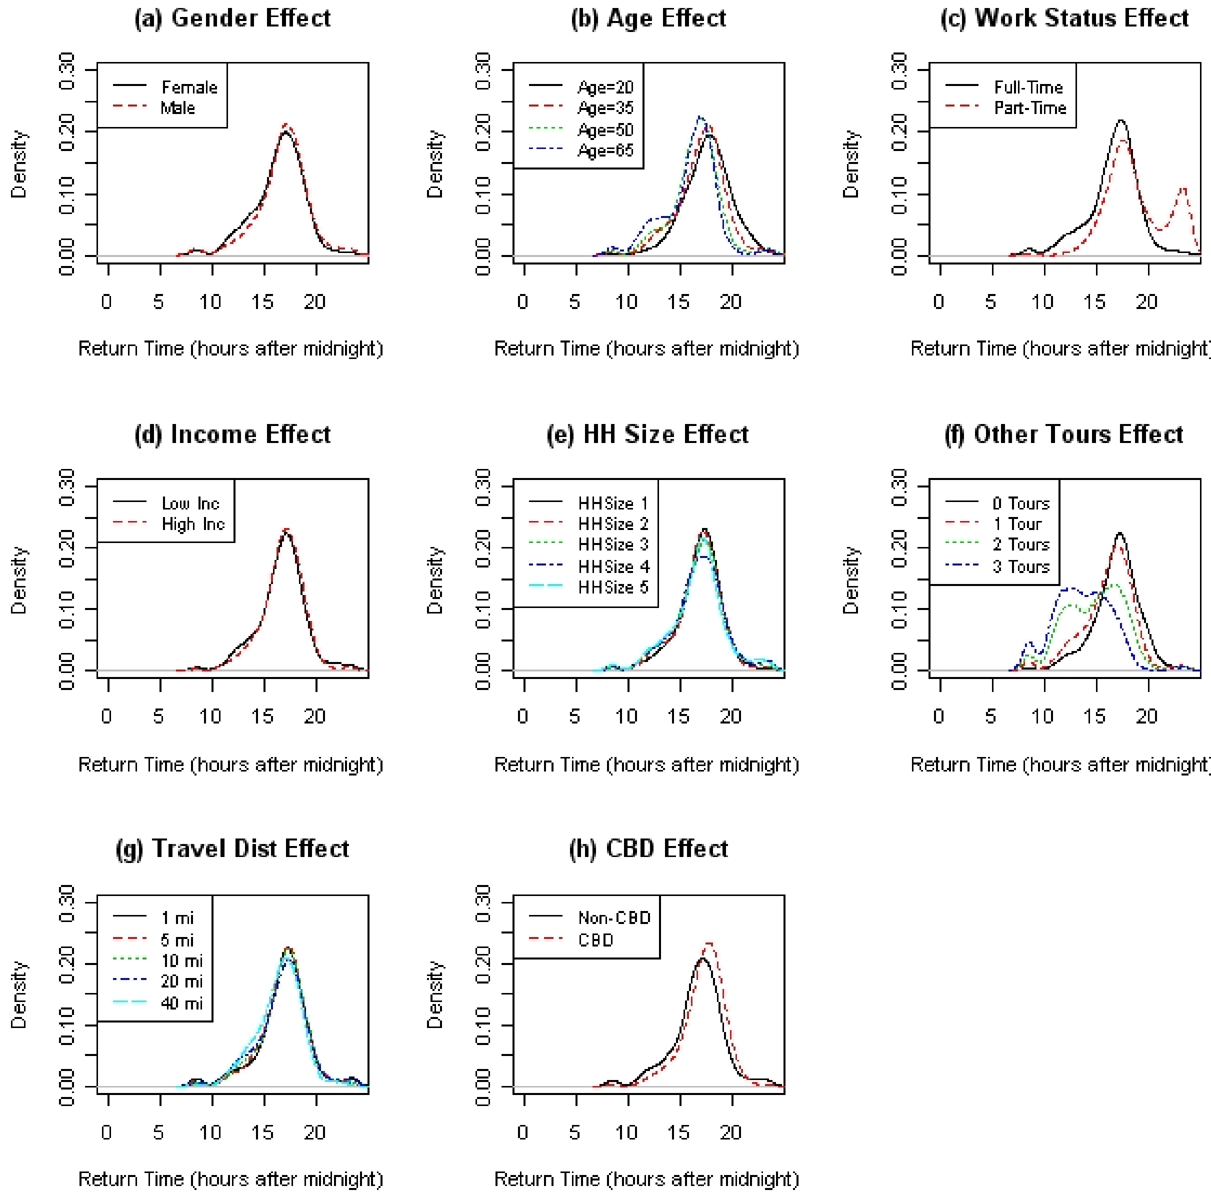

Figure 2: Predicted Return Time Density Profiles for Individuals with Different Attributes

5.2 Out-of-Sample Predictions

To better appreciate the predictive ability of the BVMNP model, relative to a simple joint-choice multinomial logit (MNL) model, out-of-sample prediction was performed, using a 20% random sample. In the case of the MNL, each choice alternative represents the arrival and return time alternative pair – in contrast to the BVMNP model, which represents arrival times and return times as distinct choices. The MNL model was estimated using BIOGEME software and employed 50 randomly chosen alternatives from the set of all 621 joint choice alternatives along with the chosen alternative[5]. Since the MNL model

[5] McFadden (1978) showed that one can use a simple random sample of alternatives for MNL estimation and still obtain consistent parameter estimates.

represents arrival and return time choice jointly, additional utility components related to alternatives' duration (equal to the return time minus arrival time) and its squared value were included in the model, similar to Popuri et al.'s (2008) specification. For consistency with the BVMNP model's duration components, the sample was segmented by full-time workers with no additional tours for the day and all other workers (i.e., part-time workers and those with additional tours). Since the MNL model is estimated using classical techniques, the predictive likelihood is simply a fixed value.

For the BVMNP model (as with any MNP), predictive likelihoods are difficult to compute, due to open-form likelihood expressions. Instead, using random parameter draws from the posterior, utilities were drawn from their corresponding distribution for each individual, taking the maximum utilities to signify the chosen alternative. The probabilities of accurate arrival- and departure-time slot prediction were then averaged over all individuals.

It turns out that predictive accuracies of the BVMNP model are clearly superior to the MNL. In fact, the BVMNP specification beats the MNL nearly 99% of the time. This suggests the BVMNP specification is superior in terms of model fit to the MNL. In the following section, several toll policy simulations are examined using both specifications.

5.3 Policy Simulations

This section examines the consumer surplus (CS) and predicted departure time distribution changes under each of three toll-policy simulations. In the first, it is assumed that $0.15/mile tolls are assessed on all roads during the peak periods and that these result in peak-period travel time *delay* reductions of 50%. In the second, the same tolls are assessed, but peak-period delay reductions are assumed to be just 10%. In the final simulation, $0.30/mile peak-period tolls are assumed to reduce peak period delays by 50%. Note that a 50% delay reduction is not the same as a 50% travel time reduction. Delay is measured in relation to free-flow travel time. Thus, routes with no delay during peak periods do not benefit from the toll policies. The simulations were performed using the estimation sample data for those traveling by auto mode only (since non-auto mode users would be largely unaffected by toll policies), and 1,000 random draws from the collection of posterior draws were employed (with the same 1,000 random draws used in each policy simulation). Since the MNL model was not estimated using Bayesian methods, its CS changes represent point estimates, rather than the distributional estimates that emerge from the BVMNP model.

Figure 3 shows the distribution of CS change (measured as the difference in CS between tolling and status quo simulations) for the three tolling policy simulations under the MNL and BVMNP specifications. Average estimates of CS change per traveler were -$0.90 and -$1.49 for simulation 1 under the BVMNP and MNL specifications, respectively, -$1.51 and -$1.85 for simulation 2, and -$2.41 and -$3.36 for simulation 3. Under both model specifications, Simulation 1's CS changes are estimated to be least negative, not surprisingly since its combination of tolls and delay reductions should offer the greatest value to travelers. In addition, both models predict the CS change under tolling simulation 3 to be more than twice as negative as simulation 1, as expected, since delay reductions are identical for the two simulations, but tolls are twice as large in simulation 3.

Based on the differences between simulations 1 and 2 under both model specifications, it appears that the peak period delay reduction does have a significant effect on CS change. Tolls under these simulations are identical, but peak travel delay is reduced by 40% more in simulation 1 as compared to simulation 2. The most notable difference between the two models here is the magnitude of CS change, which is estimated to be much lower for the MNL than for the BVMNP. This could be because the BVMNP model recognizes the similarities between alternatives near peak periods and peak periods (via error term correlations). In other words, peak shoulder slots may not be viewed as poorly by travelers under the BVMNP specifications, when compared to the MNL, where correlations do not exist. While CS change is important for scenario evaluations and policymaking, it is also important to understand tolls' effects on travelers' scheduling choices.

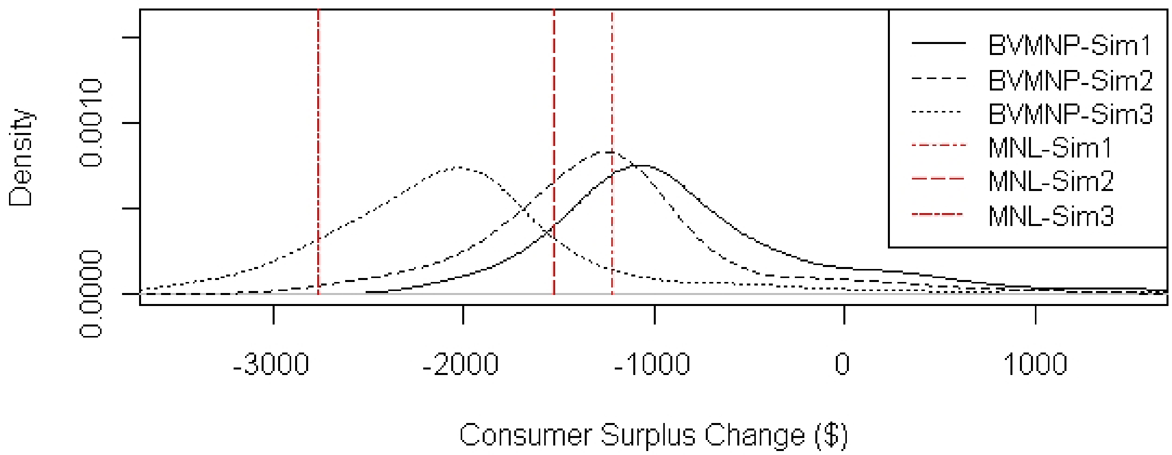

Figure 3: Consumer Surplus Change Distributions Under Toll Policy Simulations

To appreciate the tolls' effects on traveler behavior, Figure 4 illustrates simulated aggregate predictive arrival and return time densities for the two models under each scenario. Examining these distributions, it is immediately clear that the MNL specification implies that workers are rather non-responsive to peak tolling policies, while much larger scheduling shifts emerge under the BVMNP specification.

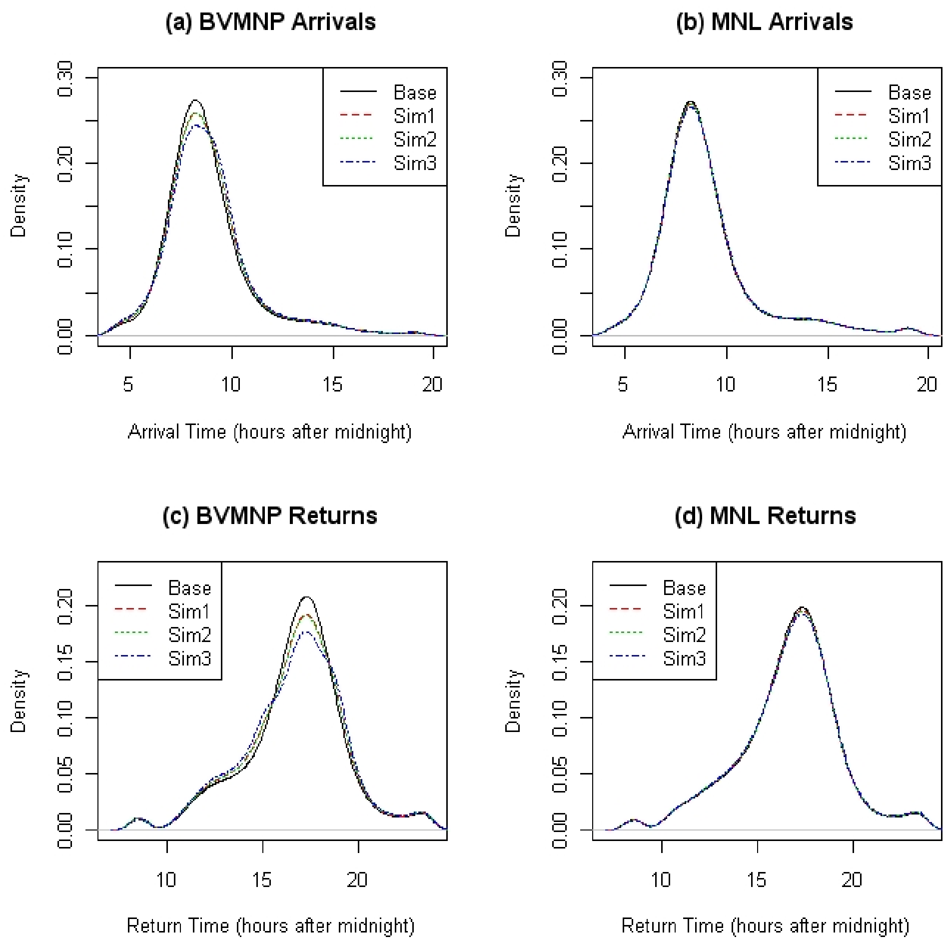

Figure 4: Arrival and Return Time Predictive Densities Under Toll Policy Simulations

To further examine differences in the two models, three time-of-day periods were examined here: an off-peak period (before 5 am or after 10 am for arrivals and before 2:30 pm or after 7:30 pm for returns), a shoulder peak period (5 to 6 am or 9 to 10 am for arrivals and 2:30 to 3:30 pm or 6:30 to 7:30 pm for returns), and the peak period (6 to 9 am for arrivals and 3:30 to 6:30 pm for returns). Given that the simulated tolls are only applied during peaks, one expects shoulder periods to experience a relatively large increase in shares, as compared to off-peak periods (i.e., before 5 am and after 10 am for arrivals and before 2:30 pm and after 7:30 pm for returns). Under the MNL specification, percentage changes for off-peak and shoulder periods are similar in magnitude for arrivals and returns and for each toll policy simulation, highlighting the MNL's independence of irrelevant alternatives assumption. The BVMNP specification, on the other hand, exhibits percentage changes more in line with expectations. The BVMNP predicts peak-period travelers to shift more toward shoulder periods, with the greatest shoulder period shifts under simulation 3, not surprisingly. Of course, this is due to the correlations offered under the BVMNP specification.

6. CONCLUSIONS

Better methods for modeling travelers' tour scheduling choices are clearly needed. This paper develops and applies a bivariate multinomial probit (BVMNP) model as a reasonable option in the context of tour scheduling. Here, work-tour scheduling was found to be most influenced by age (with older individuals arriving and returning from work earlier in the day), worker status (with part-time worker schedules shifted to later times-of-day), and the number of travel tours undertaken by the worker (with those having additional tours arriving slightly later and returning from work much earlier), each of which seems reasonable.

Empirical evidence suggests that the BVMNP performs better than a relatively straightforward, and standard, joint MNL model of all paired timing choices. In addition, the BVMNP model offers more reasonable scheduling predictions under various tolling policy simulations. While the MNL predicted rather small changes in the number of peak period travelers, the BVMNP predicted much more substantial peak-period travel reductions. In addition, the BVMNP consistently predicted relatively large share increases for peak shoulder periods (consistent with expectations), whereas MNL predictions for peak shoulders were generally similar in magnitude to off-peak period share increases.

Bayesian techniques are particularly advantageous in estimating the BVMNP model, where conditional posterior distributions were derived for latent utility variables, covariance components, and utility function parameters separately. Bayesian estimation also provides draws from the multivariate posterior distribution of all parameters. In terms of risk and uncertainty analysis (which is particularly important for toll road analyses), posterior draws offer a natural setting for capturing such uncertainty in demand modeling and other systems.[6] The scenario analyses provided in this paper illustrate how uncertainty passes through the model into welfare estimates.

While great strides in activity-based travel demand modeling have been made in recent years, time-of-day modeling remains a key weakness of most model systems. Due to the very large number of alternatives one must typically consider, a major difficulty in tour scheduling models is how to capture correlation across alternatives, since one does not expect independence across such similar alternatives. The BVMNP model offers researchers and analysts the ability to capture such correlations in a meaningful way, while retaining a random utility framework. Of course, more research and experimentation would be needed to fully appreciate the relative merits and limitations of the model.

ACKNOWLEDGEMENTS

Funding was provided by the Strategic Highway Research Program (SHRP) II, Project CO4, titled “Improving Our Understanding of How Highway Congestion and Pricing Affect Travel Demand”. The authors would like to thank Mark Bradley for assembling and formatting the data used in this work.

REFERENCES

AASHTO (2007) Combating Congestion through Leadership, Innovation, and Resources: A Summary Report on the 2007 National Congestion Summits. American Association of State Highway and Transportation Officials (AASHTO). Accessed July 2010 at http://downloads.transportation.org/CTL-1.pdf.

Abkowitz, M.D. (1981) An Analysis of the Commuter Departure Time Decision. *Transportation*, 10 (3), 283-297.

[6] Capturing uncertainty in model inputs is also important for risk analysis (see, e.g., Lemp and Kockelman 2009).

Abou Zeid, M., T.F. Rossi, and B. Gardner (2006) Modeling Time of Day Choice in the Context of Tour and Activity Based Models. *Transportation Research Record*, 1981, 42-49.

Albert, J.H. and S. Chib (1993) Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association*, 88, 669-679.

Bhat, C.R. (1998) Analysis of Travel Mode and Departure Time Choice for Urban Shopping Trips. *Transportation Research Part B*, 32 (6), 361-371.

Bhat, C.R. and J.L. Steed (2002) A Continuous-Time Model of Departure Time Choice for Urban Shopping Trips. *Transportation Research Part B*, 36 (3), 207-224.

Chin, A.T.H. (1990) Influences on Commuter Trip Departure Time Decisions in Singapore. *Transportation Research Part A*, 24 (5), 321-333.

de Jong, G., A. Daly, M. Pieters, and T. van der Hoorn (2007) The Logsum as an Evaluation Measure: Review of the Literature and New Results. *Transportation Research Part A*, 41, 874-889.

Gamerman, D. and H.F. Lopes (2006) *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, 2nd Edition, Chapman & Hall/CRC, Boca Raton.

Geweke, J., M. Keane, and D. Runkle (1994) Alternative Computational Approaches to Inference in the Multinomial Probit Model. *The Review of Economics and Statistics*, 76 (4), 609-632.

Golob, T.F. and A.C. Regan (2002) Trucking Industry Adoption of Information Technology: A Multivariate Discrete Choice Model. *Transportation Research Part C*, 10, 205-228.

Holden, L., R. Hauge, and M. Holden (2009) Adaptive Independent Metropolis-Hastings. *The Annals of Applied Probability*, 19 (1), 395-413.

Kockelman, K.M. and J.D. Lemp (2009) The Financing of New Highways: Opportunities for Welfare Analysis and Credit-Based Congestion Pricing. Proceedings of the 88th Annual Meeting of the Transportation Research Board, January 2009, Washington, D.C.

Lemp, J.D. (2009) Capturing Random Utility Maximization Behavior in Continuous Choice Data: Application to Work Tour Scheduling. Doctoral Dissertation, Department of Civil, Architectural, and Environmental Engineering, The University of Texas at Austin.

Lemp, J.D., K.M. Kockelman, and P. Damien (2010) The Continuous Cross-Nested Logit Model: Formulation and Application for Departure Time Choice. *Transportation Research Part B* 44 (5): 646-661.

Lemp, J.D. and K.M. Kockelman (2009) Understanding and Accommodating Risk and Uncertainty in Toll Road Projects: A Review of the Literature. *Transportation Research Record*, 2132, 106-112.

McCulloch, R. and P.E. Rossi (1994) An Exact Likelihood Analysis of the Multinomial Probit Model. *Journal of Econometrics*, 64, 207-240.

McFadden, D. (1989) A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration. *Econometrica*, 57 (5), 995-1026.

PB Consult (2005) The MORPC Travel Demand Model: Validation and Final Report. Prepared for the Mid-Ohio Regional Planning Commission as part of the MORPC Model Improvement Project.

Popuri, Y., M. Ben-Akiva, and K. Proussaloglou (2008) Time of Day Modeling in a Tour-Based Context: The Tel-Aviv Experience. Proceedings of the 87th Annual Meeting of the Transportation Research Board, January 2008, Washington, D.C.

Rossi, P.E., G.M. Allenby, and R. McCulloch (2005) *Bayesian Statistics and Marketing*. John Wiley & Sons, Ltd., Hoboken, NJ.

Schofer, J.L. (2005) Summary Statement. Proceedings of the USDOT Expert Forum on Road Pricing and Travel Demand Modeling, Alexandria, VA.

Small, K.A. (1982) The Scheduling of Consumer Activities: Work Trips. *The American Economic Review*, 72 (3), 467-479.

Small, K.A. (1987) A Discrete Choice Model for Ordered Alternatives. *Econometrica*, 55 (2), 409-424.

Steed, J. and C.R. Bhat (2000) On Modeling the Departure Time Choice for Home-Based Social/Recreational and Shopping Trips. *Transportation Research Record*, 1706, 152-159.

Transportation Research Board (2007) Metropolitan Travel Forecasting: Current Practice and Future Direction. TRB Special Report 288, Committee for Determination of the State of the Practice in Metropolitan Area Travel Forecasting, Washington, D.C.

Vovsha, P. and M. Bradley (2004) A Hybrid Discrete Choice Departure Time and Duration Model for Scheduling Travel Tours. *Transportation Research Record*, 1894, 46-56.

Vovsha, P., W. Davidson, and R. Donnelly (2005) Making the State of the Art the State of the Practice: Advanced Modeling Techniques for Road Pricing. Proceedings of the USDOT Expert Forum on Road Pricing and Travel Demand Modeling, Alexandria, VA.

Wang, J.J. (1996) Timing Utility of Daily Activities and Its Impact on Travel. *Transportation Research Part A*, 30 (3), 189-206.

Zhang, X., W.J. Boscardin, and T.R. Belin (2008) Bayesian Analysis of Multivariate Nominal Measures Using Multivariate Multinomial Probit Models. *Computational Statistics and Data Analysis*, 52, 3697-3708.