Estimating Daily Trip Distributions from Aggregate Trip Data

Teethat Vongvanich*[1] and Francesco Viti[1]

[1]Department of Engineering, University of Luxembourg, Esch-sur-Alzette, Luxembourg.
* *teethat.vongvanich@uni.lu*

Short summary

Travel demand models typically produce single point estimates of trip flows, discarding the uncertainty inherent in behavioral heterogeneity and external factors. This paper extends the utility-based Bayesian MCMC framework of Scheffer et al. (2021) in three ways. First, the full MCMC posterior is propagated through the forward model via posterior predictive sampling, yielding a distribution of hourly trips — characterised by a mean and 95% credible band — rather than a single point estimate. Second, the aggregate trip distribution is decomposed into activity-specific components estimated from aggregate trip counts and activity participation data, without individual travel records. Third, scalability is evaluated on random subsamples (5%, 10%, 25%) and demographic strata, with outputs linearly rescaled to population size. Applied to a dataset synthetically derived from the Luxembourg national travel survey (27,490 respondents, 13,714 qualifying tours, four activity types), the posterior predictive mean closely matches observed demand ($R^2 = 0.963$). A 5% random sample rescaled by a factor of 20 achieves $R^2 = 0.969$. Scalability is bounded by sample representativeness, not data volume.
Keywords: Transport demand modelling, Trip distribution estimation, Trip purpose

1 Introduction

Travel demand modeling estimates trips and their temporal dynamics to support transport planning decisions. Two dominant paradigms exist: trip-based models are computationally scalable but lack behavioral realism (McNally, 2007), while activity-based models capture detailed behavior but require extensive individual-level data and are computationally intensive (Bowman & Ben-Akiva, 2001; Charypar & Nagel, 2005). A fundamental challenge persists: within-day and day-to-day demand estimates remain unreliable, and current models neglect uncertainty quantification (Rasouli & Timmermans, 2012). This limitation is critical because real demand is better represented as a distribution reflecting behavioral heterogeneity and external factors, yet models continue to produce only point estimates.

Existing approaches exhibit specific limitations. Trip-based models ignore trip purpose, activity scheduling, and duration constraints (Ortúzar & Willumsen, 2011; Cantelmo et al., 2018), providing limited insight into temporal demand distribution. Activity-based models, though behaviorally rich, are difficult to apply and validate at large scale where calibration and consistency with macroscopic emergent behavior remain open challenges (Moeckel et al., 2017). Their application requires data resources and computational power attainable in major cities such as Zurich or Toronto (Becker et al., 2019; Gao et al., 2010) but typically unavailable in small and medium-sized cities. Dynamic trip matrix estimation methods focus on consistency between trip matrices and traffic counts (Antoniou et al., 2016; Cascetta et al., 1993) but rarely capture the behavioral component of demand explicitly (Flötteröd, 2009). The critical gap remains: no approach successfully combines behavioral consistency with aggregate data while incorporating uncertainty quantification.

Scheffer et al. (2021) demonstrated that utility maximization principles can be combined with Bayesian MCMC calibration to estimate activity-specific zone-level demand from aggregate trip counts. The model is parameterized by activity-specific utility primitives — marginal utility functions that describe how individuals value performing a given activity at a given time of day. These parameters are zone-specific and are estimated from aggregate hourly trip counts (departures and arrivals by zone), together with activity participation counts per zone. The framework does not require individual travel diaries during calibration; the likelihood is evaluated against aggregate counts only.

This study extends the work of Scheffer et al. (2021) in three ways. First, rather than retaining only the maximum a posteriori point estimate of utility parameters, we propagate the full MCMC posterior through the forward model via posterior predictive sampling. This yields a distribution of hourly trips — characterized by a mean and 95% credible band — for each activity type, capturing parameter uncertainty explicitly. Second, we decompose the aggregate trip distribution into activity-specific components, showing that trip primitives for each activity can be recovered from aggregate data and interpreted behaviorally. Third, we analyze scalability: the full procedure is repeated on subsamples of the population (5%, 10%, 25%, and demographic strata), with outputs linearly rescaled to population size. A 5% random sample reproduces the full-population trip distribution accurately; demographically skewed subsamples do not. This shows that scalability is bounded by the representativeness of the calibration sample, not by data volume. The approach is applied to a synthetic dataset derived from a national travel survey of Luxembourg comprising 27,490 respondents and four activity types (Work, School, Leisure, Other). Results are presented here for a single-zone configuration; multi-zone results with destination choice are presented separately at the conference. The remainder of the paper is organized as follows. Section 2 describes the model. Section 3 describes the trip distribution generation procedure. Section 4 presents the case study. Section 5 reports the results. Section 6 concludes.

2 Model

The model estimates activity-specific, time-dependent trips from aggregate trip data. The underlying idea is that individuals choose when and where to travel in order to maximize the utility they accumulate over a full day. Individuals choose a destination zone, a departure time, and an activity end time to maximize the utility of performing that activity as opposed to any other alternative activity. Zone-specific utility parameters are calibrated from observed zonal trip counts via Bayesian MCMC.

Inputs, Outputs, and Notation

Let $\mathcal{A}$ denote the set of activity types and $\mathcal{Z}$ the set of zones. Model inputs, outputs, and key notation are summarized in Table 1.

Marginal Utility Functions

Marginal utility functions represent the instantaneous utility gained from performing an activity at a given time (Ettema & Timmermans, 2003). The functional form adopted here is unimodal and flexible enough to represent any activity type:

$$u(t,\, t_s;\, \Theta) \;=\; \frac{U^{\max}\,\beta\,\gamma\,\exp\big[\beta\big(t-(\alpha+\tau t_s)\big)\big]}{\Big(1+\exp\big[-\beta\big(t-(\alpha+\tau t_s)\big)\big]\Big)^{\gamma+1}} \tag{1}$$

The parameter vector is $\Theta = (U^{\max},\, \alpha,\, \beta,\, \gamma,\, \tau)$. A behavioral interpretation is as follows. $U^{\max}$ represents the maximum marginal utility at the saturation point, controlling the gain from performing one activity relative to any other alternative including the travel required to reach it. α is the location of the saturation point: for clock-based activities ($\tau = 0$) it is the time of day at which marginal utility peaks, measured in hours from midnight; for duration-based activities ($\tau = 1$) it is the elapsed time from activity start at which marginal utility peaks. β determines the concentration around the saturation point — a larger β gives a sharper, more peaked utility profile, and is the primary determinant of estimated activity duration. γ governs asymmetry around the saturation point. The parameter τ is fixed per activity type and is not estimated. Clock-based activities such as Work have start times concentrated in the morning and are well described by a fixed time of day. Duration-based activities such as Leisure are less dependent on clock time and are better described by their typical duration (Ettema & Timmermans, 2003).

Three marginal utility functions are parameterized independently per activity a and zone: $u^{-}_{a,y}$ for the pre-departure period (indexed by origin zone y), $u_{a,z}$ for the out-of-home activity, and $u^{+}_{a,z}$ for the post-return period (both indexed by destination zone z). This zone-specific parameterization allows spatial heterogeneity in preferences across the network.
Figure 1 illustrates the three-phase decomposition for a single activity tour.

Table 1: Model inputs, outputs, and notation

Quantity	Description	Dimensions	Notation
Inputs			
Trip flows	Time-dependent trip counts by zone and activity	zones × time	$T^a_{y\to}(h)$, $T^a_{\to z}(h)$
Participation	Activity participants per zone per day	zones × activities	N^y_a
Travel times	Zone-to-zone travel times	zones × zones	tt^a_{yz}
Priors	Parameter prior distributions	per parameter	$\pi(\Theta)$
Outputs			
OD flows	Hourly activity-specific OD matrices	zones × zones × time	$\mathrm{OD}^a_{yz}(h)$
Distributions	OD mean, std, credible bands	zones × zones × time	$\bar{\mathrm{OD}}_{yz}(h)$, $\sigma_{yz}(h)$
Parameters	Posterior parameter distributions	per (a, zone)	$\Theta^-_{a,y}$, $\Theta_{a,z}$, $\Theta^+_{a,z}$
Other notation			
t, Δt	Time of day (min); time-step (15 min)		
t_d, t_s, t_e	Departure, activity start, activity end times		
u_t	Travel disutility per minute ($u_t < 0$, fixed)		
θ, λ	Logit scale; likelihood weight in MCMC score		

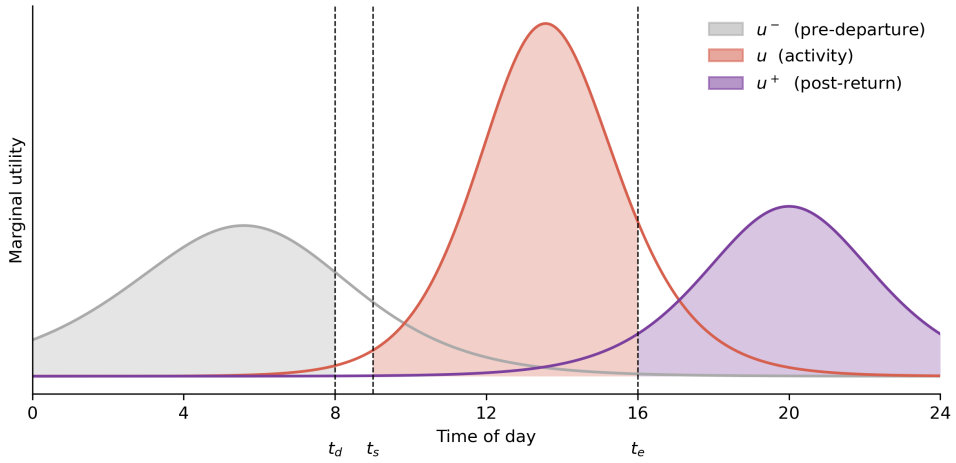


Figure 1: Accumulated utility for a sequence of activities

Accumulated Utility and Tour Choice

To maximize total accumulated utility over a full day to perform an activity, individuals seek an optimal destination, departure time, and activity end time. The total utility consists of four components: the utility of any activity performed before departure (u^-), the disutility of travel, the utility of the out-of-home activity (u), and the utility of any activity performed after return (u^+). This decomposition allows the estimation problem to focus on one activity at a time. For a person based in zone y choosing activity a in zone z, departing at t_d and ending at t_e:

$$\hat{U}_a(t_d, t_e \mid y, z) \;=\; \sum_{t=0}^{t_d} u^-_{a,y}(t)\,\Delta t \;+\; u_t\, tt^a_{yz} \;+\; \sum_{t=t_s}^{t_e} u_{a,z}(t,\, t_s)\,\Delta t \;+\; \sum_{t=t_e}^{T} u^+_{a,z}(t,\, t_e)\,\Delta t \qquad (2)$$

where $t_s = t_d + tt^a_{yz}$ and $T = 1440$ min. The choice set covers all feasible triples (z, t_d, t_e) subject to per-activity constraints on departure window and activity duration. A multinomial logit model gives the choice probability for departure time and end time:

$$P_a^{yz}(t_d, t_e) \;=\; \frac{\exp\big(\hat{U}_a(t_d, t_e \mid y, z)\,/\,\theta\big)}{\displaystyle\sum_{z'}\sum_{t'_d}\sum_{t'_e} \exp\big(\hat{U}_a(t'_d, t'_e \mid y, z')\,/\,\theta\big)} \tag{3}$$

For each origin y, activity a, and feasible alternative (z, t_d, t_e), a fractional flow $N_a^y \cdot P_a^{yz}(t_d, t_e)$ is generated. Because the model operates on complete tours, each flow contributes simultaneously to four hourly OD matrices: the outgoing departure, the outgoing arrival, the return departure, and the return arrival. The return leg follows deterministically from t_e and tt^a_{yz} and is not assumed or fitted separately.

Bayesian MCMC Calibration

Parameter values are inferred via MCMC, yielding a posterior distribution over the parameter vector that reflects consistency with observed aggregate trip counts and the prior.

Starting from initial values drawn from activity-specific priors, the sampler proposes new candidate parameters at each iteration. Proposals are accepted or rejected based on how well the resulting trips match observed hourly departures and arrivals, and on the plausibility of the parameters given the prior. The acceptance score is $S(\Theta) = \mathcal{L}(\Theta)/\lambda + \log \pi(\Theta)$, where $\mathcal{L}$ is the log-likelihood of the observed counts given the model, π is the prior, and $\lambda = 0.5$ is a weight controlling the relative influence of the two terms. Independent chains are run in parallel; the first 33% of each chain is discarded as burn-in and the post-burn-in samples are pooled into a single posterior. The complete estimation process is summarized as:

$$\pi(\Theta) \xrightarrow{\text{MCMC}} p(\Theta \mid \text{data}) \xrightarrow{\text{Eq. 1}} u_a(t, t_s) \xrightarrow{\text{Eq. 2–3}} P_a^{yz}(t_d, t_e) \xrightarrow{N_a^y} T^a_{y\rightarrow}(h) \xrightarrow{\mathcal{L}} S(\Theta) \tag{4}$$

3 Daily Trips as Distributions

Standard MCMC calibration yields a full posterior over the parameter vector Θ. The conventional approach discards this posterior and uses only the maximum a posteriori estimate to produce a trip estimate. Instead, we propagate the full posterior through the forward model to obtain a distribution of trips, characterized by a mean, standard deviation, and credible bands for each hour.

Posterior Predictive Sampling

Let $\{\Theta^{(s)}\}_{s=1}^{S}$ be S parameter vectors drawn uniformly without replacement from the post-burn-in MCMC chain. For each draw s, the forward model (Equations 2–3) is evaluated to produce one complete set of hourly trips. The S realizations are stored in an accumulator array of shape $(S \times |\mathcal{Z}| \times |\mathcal{Z}| \times 24)$. In this study $S = 200$, drawn with a fixed random seed for reproducibility.

Summary Statistics

The posterior predictive mean and standard deviation are computed across the S realizations:

$$\bar{\text{OD}}_{yz}(h) = \frac{1}{S}\sum_{s=1}^{S} \text{OD}_{yz}^{(s)}(h), \qquad \sigma_{yz}(h) = \sqrt{\frac{1}{S}\sum_{s=1}^{S}\big(\text{OD}_{yz}^{(s)}(h) - \bar{\text{OD}}_{yz}(h)\big)^2} \tag{5}$$

Empirical 95% credible bands are the 2.5th and 97.5th percentiles across the S realizations. Both the symmetric band $(\bar{\text{OD}} \pm \sigma)$ and the empirical percentile band are computed; the latter captures non-Gaussian tails in the posterior.

Scalability and Subsample Rescaling

To test whether the model can be calibrated on a fraction of the available data, the full MCMC procedure is repeated on a subsample of N_{sub} persons. All posterior predictive outputs are then

linearly rescaled to population scale by $f = N_{\text{pop}}/N_{\text{sub}}$. For a linear scaling $X \to fX$, $\text{SD}(fX) = f \cdot \text{SD}(X)$, so the standard deviation and percentile bounds are multiplied by f directly. The rescaled distribution is then compared against the observed ground-truth flows.

4 Case Study: Luxembourg

The dataset is synthetically derived from a national travel diary survey of Luxembourg comprising 27,490 respondents. Each record contains the full daily trip chain of one person: departure and arrival times, origin and destination communes, transport mode, and activity type. Activity types are Work, School, Leisure, Other, and Home.

Tours are extracted from the trip chains by the following filters applied in sequence: (1) the person made at least one trip; (2) the final destination is Home; (3) exactly one non-Home activity appears in the chain; (4) activity duration is strictly positive and less than 1440 minutes. This yields 13,714 qualifying tours. External origins and destinations (Belgium, France, Germany) are excluded.

Results are presented here for a single-zone configuration covering all of Luxembourg; multi-zone results with destination choice are presented at the conference. The model supports multi-zone operation, but the 1-zone setup is used here to facilitate interpretation and direct comparison across subpopulations. The mean travel time is 30.1 minutes, the arithmetic mean of all valid one-way trip durations in the survey. In the multi-zone case, travel time is OD-pair specific and is computed as the mean one-way travel time across all observed trips between each origin-destination pair.

Work and School use clock-based utility functions ($\tau = 0$); Leisure and Other use duration-based functions ($\tau = 1$). The travel disutility is fixed at $u_t = -0.5$ utility units per minute, the logit scale parameter is $\theta = 1.0$, and the likelihood weight is $\lambda = 0.5$. These values were selected empirically because they produced stable parameter recovery and realistic demand profiles across all activity types tested.

To evaluate scalability and behavioral validity, the MCMC calibration is repeated on 14 subpopulations. Nine random samples are drawn without replacement at three fractions — 25%, 10%, and 5% — with three random seeds each, yielding approximately 3,430, 1,372, and 686 persons respectively. Five demographic strata are defined deterministically: children (age < 18), adults (age ≥ 18), elderly (age ≥ 65), male, and female. All subpopulations use the same zone boundaries and travel time matrix as the full-population run.

5 Results and Discussion

Utility Parameter Estimates

Table 2 reports the posterior mean and 95% credible interval for $U^{\max}$ and α of the in-activity utility function u for each activity type. α is in hours from midnight for clock-based activities (Work, School) and in hours from activity start for duration-based activities (Leisure, Other). $U^{\max}$ is the maximum utility attained at time α.

Table 2: Posterior mean and 95% credible interval for utility parameters per activity type

	Clock-based ($\tau = 0$)		**Duration-based** ($\tau = 1$)	
	Work	**School**	**Leisure**	**Other**
$U^{\max}$	16.39 [14.66, 17.62]	11.70 [11.25, 12.19]	10.68 [9.63, 11.53]	7.93 [7.25, 8.59]
α	10.19 [8.24, 12.94]	11.35 [8.57, 12.91]	1.42 [0.35, 1.85]	3.63 [2.86, 4.35]

$U^{\max}$ controls the gain from performing the activity relative to any other alternative including the travel required to reach it (Ettema & Timmermans, 2003). Work has the highest $U^{\max}$ by a substantial margin, reflecting that individuals are willing to travel further or accept longer commutes for work than for any other activity. Other has the lowest $U^{\max}$, consistent with its heterogeneous and discretionary nature.

The parameter τ is fixed per activity type and determines how α is interpreted. For clock-based activities ($\tau = 0$), Work and School, α is the time of day at which marginal utility peaks, measured

in hours from midnight. The marginal utility of Work peaks at 10.19h and School at 11.35h, consistent with structured daytime activity periods. For duration-based activities ($\tau = 1$), Leisure and Other, α is the elapsed time from activity start at which marginal utility peaks. The marginal utility of Leisure peaks 1.42 hours after activity start, and that of Other peaks 3.63 hours after activity start, consistent with short discretionary trips for Leisure and a more varied category for Other.

These parameters are recovered entirely from aggregate hourly trip counts and activity participation counts, without access to individual travel records. The fact that they align with behaviorally expected patterns validates that the aggregate signal contains sufficient information to identify activity-specific trip primitives.

Activity-Specific Trip Distributions

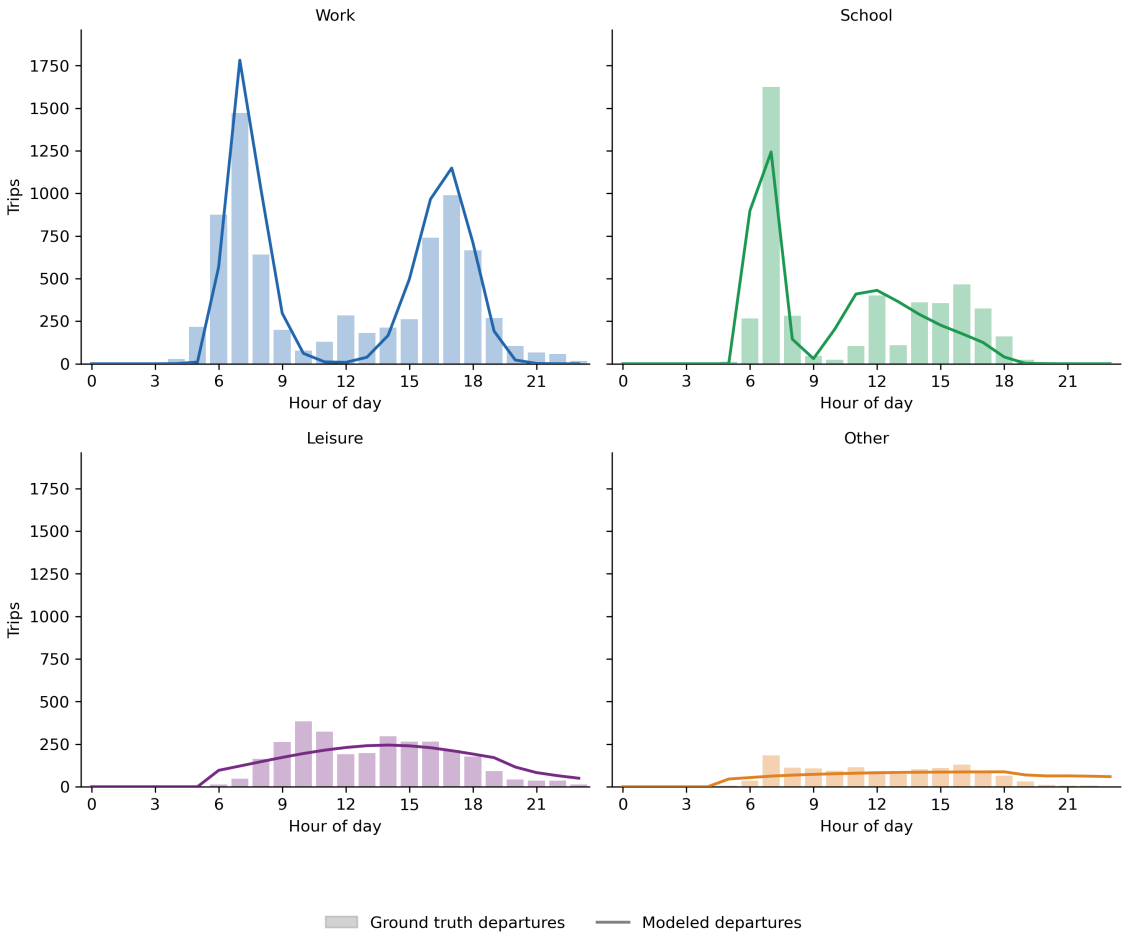


Figure 2: Posterior predictive mean departures by activity type versus ground truth

Figure 2 shows the estimated departure profiles for each activity type separately. It is a single-zone configuration and each subplot shows one activity. Work departures reproduce the sharp morning peak concentrated between 06h and 09h. School departures show a similar but earlier and narrower peak, consistent with fixed school start times. Leisure departures are distributed across midday with no dominant peak, reflecting the discretionary and flexible nature of leisure activities. Other departures show a broad midday distribution intermediate between Leisure and the mandatory activities.

These activity-specific profiles emerge from a single calibration against aggregate trip counts. The model does not receive any information about which trips belong to which activity during calibration; activity-specific structure is recovered through the utility primitives estimated per activity. The alignment between estimated profiles and the ground-truth activity-specific counts confirms that the aggregate signal is sufficient to decompose demand into its activity components.

Posterior Predictive Trip Distributions

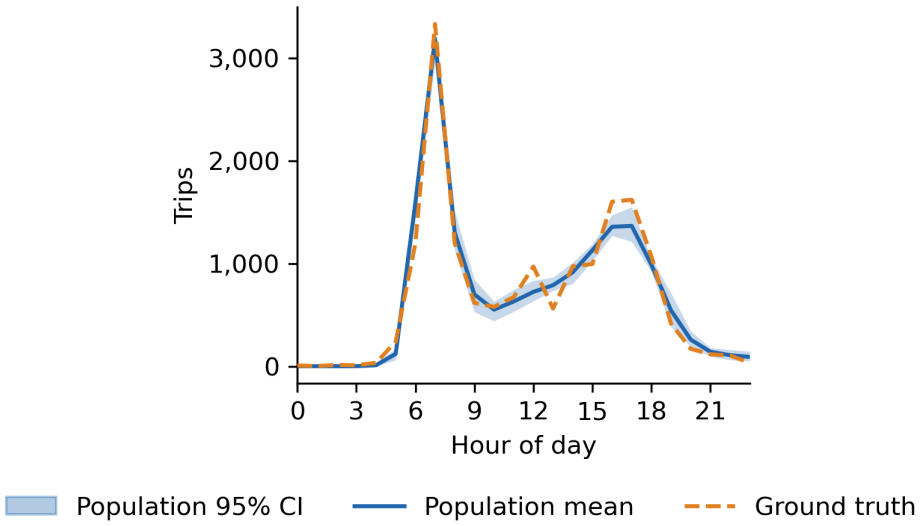


Figure 3: Posterior predictive trip distribution (total, all activities) versus ground truth

Figure 3 shows the posterior predictive distribution of total hourly trips alongside the ground-truth flows. The posterior mean closely tracks the observed profile throughout the day ($R^2 = 0.963$). The 95% credible band is narrow during peak hours, where trip counts are high and the likelihood strongly constrains the parameters, and widens during off-peak hours where counts are low.

Scalability

Figure 4 and Table 3 report scalability results. The R^2 is computed on rescaled hourly trip counts versus ground-truth flows and measures how well each sample reproduces observed demand after linear rescaling by $f = N_{\text{pop}}/N_{\text{sub}}$.

Table 3: Activity composition and scalability of calibration samples

Sample	**N tours**	**Work**	**School**	**Leisure**	**Other**	R^2
Full population	13,714	64%	18%	13%	6%	—
5% sample	686	63%	20%	13%	5%	0.969
Adults	11,580	75%	5%	14%	5%	0.900
Children	2,134	1%	88%	4%	7%	0.634
Elderly	719	4%	0%	77%	19%	−0.061

The random 5% sample (≈686 persons, $f = 20$) achieves $R^2 = 0.969$. Its activity composition closely mirrors the full population, so the estimated utility parameters are transferable under linear rescaling. This result holds across all three random seeds at 5%. At 10% and 25%, rescaled distributions converge further toward the full-population result with narrower credible bands.

The adults sample ($R^2 = 0.900$) reproduces total demand reasonably but underweights the morning peak, because adults account for only 5% School trips whereas the full population includes 18%. The children sample ($R^2 = 0.634$) shows the inverse: School dominates at 88% and Work is near-zero, producing a narrower school-day profile that does not match the full-population shape. The elderly sample produces a negative R^2 (-0.061): with 77% Leisure, their travel is distributed across midday rather than concentrated in morning and evening peaks. Rescaling cannot recover a temporal shape that is not present in the calibration data.

However, failure at the aggregate level does not mean the model fails to learn subsample behavior. Figure 5 shows that for the adults subsample, Work departures are well reproduced while School departures are near-zero — correctly reflecting the activity composition of adults. The aggregate

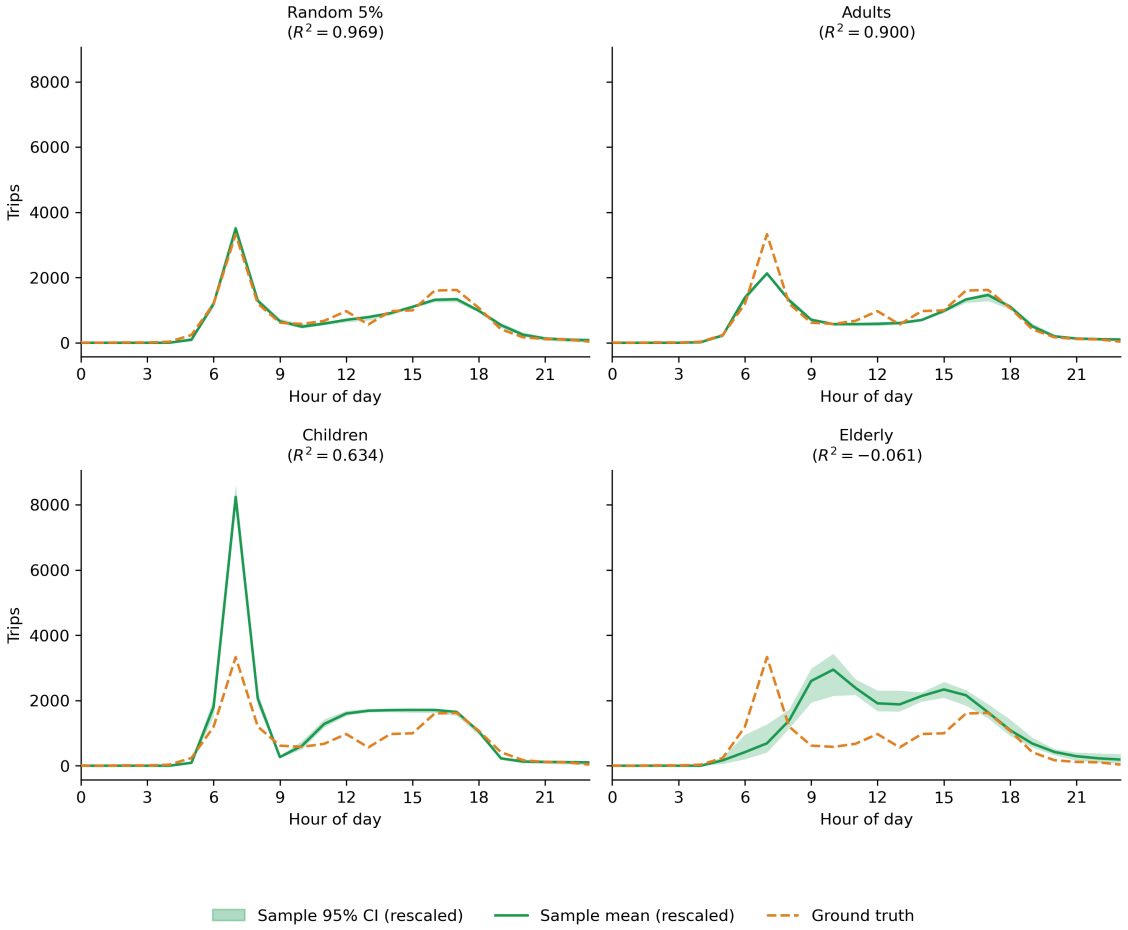


Figure 4: Rescaled trip distributions versus ground truth for four samples.

R^2 is low not because the model is wrong, but because the activity composition of the subsample does not represent the full population. Each subsample recovers the correct trip primitives for its dominant activity.

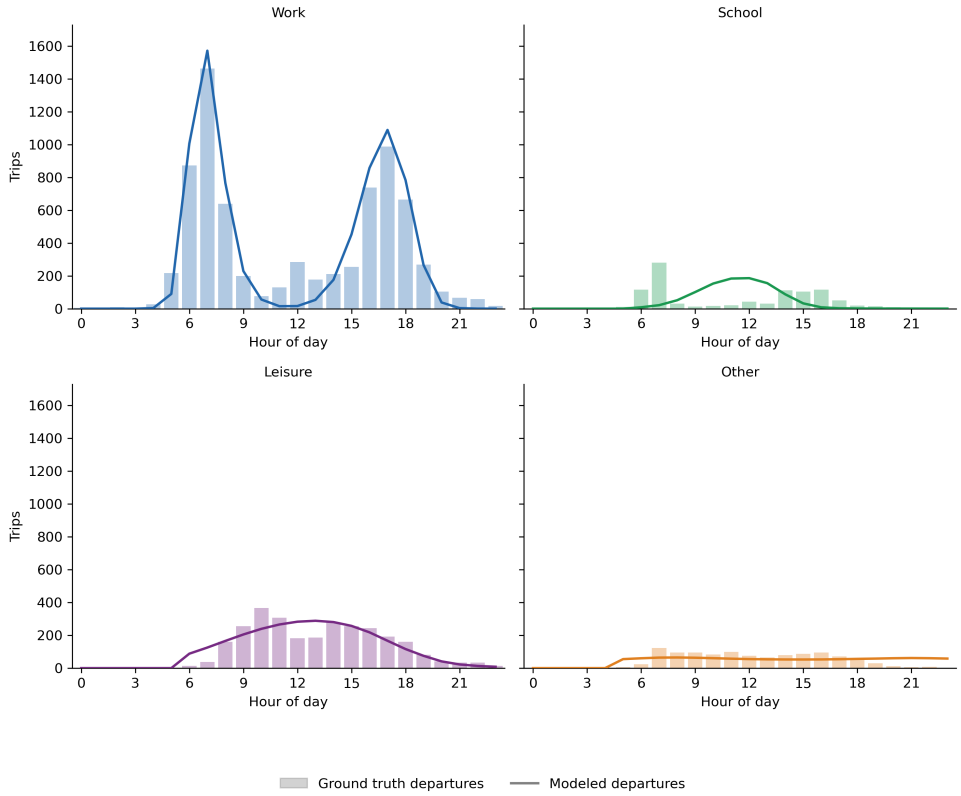


Figure 5: Activity-specific departure estimates for the adults subsample

These results demonstrate that scalability is bounded by the representativeness of the calibration sample, not by data volume. A 5% random sample suffices because it retains the activity composition of the full population. A demographically skewed sample does not; however, each subsample recovers the correct trip primitives for its dominant activity.

6 Conclusions

This study extends the framework of Scheffer et al. (2021) in three ways. First, the full MCMC posterior over utility parameters is propagated through the forward model via posterior predictive sampling, yielding a distribution of hourly trips — characterized by a mean and 95% credible band — rather than a single point estimate. Second, the aggregate trip distribution is decomposed into activity-specific components: trip primitives for each activity type are estimated from aggregate data alone and shown to be behaviorally interpretable. Work has the highest utility intensity and peaks at a consistent morning clock-time; School shows a similar but narrower peak; Leisure and Other are characterized by preferred durations rather than clock-times, consistent with their discretionary nature. Third, scalability is analysed by repeating the calibration on subsamples and linearly rescaling the outputs.

Applied to a synthetic dataset derived from a national travel survey of Luxembourg (27,490 respondents, four activity types), the posterior predictive mean closely tracks observed demand throughout the day ($R^2 = 0.963$), and the 95% credible band envelopes the ground-truth flows. A 5% random sample scaled by a factor of 20 reproduces the full-population trip distribution with $R^2 = 0.969$, confirming that a small but representative sample suffices. Demographically skewed subsamples fail at the aggregate level because their activity composition does not match the population — but each subsample correctly reproduces the trip primitives of its dominant activity: adults reproduce work trips, children reproduce school trips. Scalability is therefore bounded by representativeness, not data volume. Results are presented for a single zone covering all of Luxembourg; the model supports multi-zone operation but destination choice is not demonstrated here.

Future work will focus on a multi-zone application with destination choice, the introduction of travel time feedback between the demand model and network conditions, and regression of utility

parameters on zone characteristics to enable transferability across study areas without recalibration.

References

Antoniou, C., Azevedo, C. L., Lu, L., Pereira, F., & Ben-Akiva, M. (2016). W-SPSA in practice: Approximation of weight matrices and calibration of traffic simulation models. *Transportation Research Part C: Emerging Technologies*, *59*, 129–146.

Becker, H., Balac, M., Ciari, F., & Axhausen, K. W. (2019). Assessing the welfare impacts of shared mobility and mobility as a service (MaaS). *Transportation Research Part A: Policy and Practice*, *131*. doi: 10.1016/j.tra.2019.09.027

Bowman, J. L., & Ben-Akiva, M. E. (2001). Activity-based disaggregate travel demand model system with activity schedules. *Transportation Research Part A: Policy and Practice*, *35*(1), 1–28.

Cantelmo, G., Viti, F., Cipriani, E., & Nigro, M. (2018). A utility-based dynamic demand estimation model that explicitly accounts for activity scheduling and duration. *Transportation Research Part A: Policy and Practice*, *114*, 303–320.

Cascetta, E., Inaudi, D., & Marquis, G. (1993). Dynamic estimators of origin-destination matrices using traffic counts. *Transportation Science*, *27*(4), 363–373.

Charypar, D., & Nagel, K. (2005). Generating complete all-day activity plans with genetic algorithms. *Transportation*, *32*(4), 369–397.

Ettema, D., & Timmermans, H. (2003). Modeling departure time choice in the context of activity scheduling behavior. *Transportation Research Record: Journal of the Transportation Research Board*, *1831*, 39–46. doi: 10.3141/1831-05

Flötteröd, G. (2009). *Cadyts — A free calibration tool for dynamic traffic simulations.* (9th Swiss Transport Research Conference)

Gao, W., Balmer, M., & Miller, E. J. (2010). Comparison of MATSim and EMME/2 on Greater Toronto and Hamilton Area network. *Transportation Research Record: Journal of the Transportation Research Board*, *2197*(1), 118–128.

McNally, M. G. (2007). The four-step model. In *Handbook of transport modelling* (pp. 35–53). Emerald Group Publishing Limited.

Moeckel, R., Huntsinger, L., & Donnelly, R. (2017). From macro to microscopic trip generation: Representing heterogeneous travel behavior. *The Open Transportation Journal*, *11*(1), 31–43.

Ortúzar, J. d. D., & Willumsen, L. G. (2011). *Modelling transport* (4th ed.). John Wiley & Sons.

Rasouli, S., & Timmermans, H. (2012). Uncertainty in travel demand forecasting models: Literature review and research agenda. *Transportation Research Record: Journal of the Transportation Research Board*, *2302*(1), 1–10.

Scheffer, A. H. M., Connors, R., & Viti, F. (2021, September). Dynamic modal split incorporating trip chaining: A parsimonious approach to mode-specific demand estimation. In *Transportation research procedia.* Elsevier.