Received: 28 October 2020 | Accepted: 15 December 2021

DOI: 10.1111/rssa.12804

ORIGINAL ARTICLE

ROYAL STATISTICAL SOCIETY | Series A Statistics in Society

A dynamic choice model to estimate the user cost of crowding with large-scale transit data

Prateek Bansal[1] | **Daniel Hörcher**[2] | **Daniel J. Graham**[2]

[1]Department of Civil and Environmental Engineering, National University of Singapore, Singapore

[2]Transport Strategy Centre, Department of Civil and Environmental Engineering, Imperial College London, London, UK

Correspondence
Prateek Bansal, Department of Civil and Environmental Engineering, National University of Singapore, Singapore.
Email: prateekb@nus.edu.sg

Abstract

Efficient mass transit provision should be responsive to the behaviour of passengers. Operators often conduct surveys to elicit passenger perspectives, but these can be expensive to administer and can suffer from hypothetical biases. With the advent of smart card and automated vehicle location data, operators have reliable sources of revealed preference (RP) data that can be utilized to estimate transit riders' valuation of service attributes. To date, effective use of RP data has been limited due to modelling complexities. We propose a dynamic choice model (DCM) for population-level longitudinal RP data to address prominent challenges. In the DCM, riders are assumed to follow different decision rules (compensatory and inertia/habit) and temporal switching between decision rules based on experience-based learning is also formulated. We develop an expectation–maximization algorithm to estimate the DCM and apply our model to estimate passenger valuation of crowding. Using large-scale data of 2 months with over four million daily trips by an Asian metro, our DCM estimates show an increase of 47% in passenger's valuation of travel time under extremely crowded conditions. Furthermore, the average passenger follows the compensatory rule on only 25.5% or fewer trips. These results are valuable for supply-side decisions of transit operators.

KEYWORDS

crowding valuation, dynamic preferences, expectation–maximization, inertia, smart card data

1 | INTRODUCTION

Millions of transit riders make decisions related to the departure time and routes on a daily basis. Individual preferences manifest into passenger flows on urban rail networks. Exceeding travel flows cause well-known challenges in the form of crowding, accident risk and eventually the degradation of the quality of travel. To ensure efficient mass transit provision and circumvent congestion, the granular mechanism (i.e. individual behaviour) behind the observed aggregate passenger flows is required to be understood.

Stated preference (SP) surveys have been a workhorse method for data collection in travel behaviour modelling (Hensher, 1994). With the emergence of automated fare collection (AFC) using smart cards and automated vehicle location (AVL) data, transit operators have access to new sources of large-scale revealed preference (RP) data to estimate passengers' behaviour and their valuation of service attributes. Both SP and RP data have their strengths and weaknesses. Whereas SP surveys capture the behaviour of a sample of decision-makers in hypothetical scenarios, new RP data sources provide information about the actual travel preferences of the entire population. Studies based on SP data suffer from hypothetical bias (i.e. potential differences between actual and stated preferences) and collecting data from a representative sample of the population remain challenging. However, unlike RP data, SP data offer complete information about the sociodemographic characteristics of decision-makers and ensures enough variation in the attributes of the alternatives faced by decision-makers.

Nevertheless, new sources of population-level RP datasets hold promise to bring to a paradigm shift in travel behaviour research, if the following challenges can be tackled—(a) difficulty in identifying inherent choice experiments, creating choice sets of riders, and retrieving attributes of the alternatives faced by riders, (b) complexity of matching multiple large-scale datasets to obtain the rider's route or departure time choices and (c) scarcity of statistical tools to model day-to-day travel preferences of riders. We address these challenges in this study and show how behaviour and preferences of riders in crowded urban rail networks can be modelled using large-scale AFC and AVL datasets. The main statistical contribution of this study hinges at developing a new model (and its estimator) to capture the dynamic semi-compensatory behaviour of urban rail riders.

1.1 | Empirical context

Quantitative measurement of the user valuation of crowding in public transport is important for predicting the future outcome of policies and optimizing supply-side decisions. Crowding cost estimate is the main input for assignment models that predict the distribution of travel flows between various network segments by replicating individual route choice decisions. These estimates are actively used in other areas of service planning and regulation. Infrastructural investment appraisal, network planning, timetabling and pricing are among the most frequent applications.

Intuitively, the user valuation of crowding can be empirically identified based on the information on instances when passengers avoid travelling in crowded conditions by selecting a longer or expensive route. However, how to recognize such inherent variations in empirical data and how to model behaviour dynamics of riders to derive their valuation of crowding are challenging research questions.

Traditionally, most of the research articles and consultancy reports use SP surveys and estimate the rider's perceived value of crowding in terms of a *crowding multiplier*—the ratio of value-of-travel-time under crowded and uncrowded conditions (Bansal et al., 2021; Wardman & Whelan, 2011). SP studies generally elicit preferences of riders in a hypothetical binary route choice experiment. Assuming that riders choose a utility-maximizing route after trading-off crowding density, travel time and cost, these studies estimate discrete choice models (DCMs) to obtain the crowding multiplier (See Bansal et al., 2019; Yap et al., 2020, for the review).

Whereas the hypothetical bias and poor representation are the main limitation of SP data, the required information to estimate DCMs (riders' route preferences and attributes of all available routes) is difficult to obtain using conventional RP datasets (Tirachini et al., 2016). Due to these challenges, early crowding valuation studies relying on the RP data either deviated from DCMs (Kroes et al., 2014) or complemented RP data with SP data (Batarce et al., 2015). However, the emerging use of AFC data provides an alternative way to collect the required RP data. Tirachini et al. (2016) first illustrate how AFC data can be used to estimate the standing penalty of mass rapid transit users in Singapore, that is, the disutility of standing measured in the equivalent travel time loss. In another study, Hörcher et al. (2017) integrate AFC data with AVL data to estimate the crowding multiplier of the Hong Kong Mass Transit Railway (MTR) users. A recent study by Yap et al. (2020) estimates the crowding valuation of urban tram and bus users in the Netherlands using the framework similar to Hörcher et al. (2017).

1.2 | Empirical research gaps

We identify two research gaps in the crowding valuation literature. First, whereas day-to-day variation in route preferences and learning behaviour of riders are hard to capture in SP experiments, previous RP studies also rely on static choice models. This implies that strong assumptions had to be made on how users form expectations about attribute levels on available routes at every choice occasion. For example, Hörcher et al. (2017) assume that the observed travellers are experienced enough to know the *average* train occupancy on available routes, at the time of day when their journeys begin. Second, even if expectations of transit attributes are correctly recovered, a regular subway user, for example, a daily commuter, might not actively make a compensatory (utility-maximizing) route choice. In fact, a rider can adhere to the same route until a bad experience occurs. However, none of the previous public transport studies model such non-compensatory (e.g. inertia or habit) decision rules. The riders who choose routes based on inertia or habit should not contribute towards the crowding cost valuation in the same way as those who follow compensatory decision rules. Therefore, the segmentation of such decision rules is crucial to accurately estimate the crowding multiplier.

A dynamic choice model that accounts for expectation formation of riders and identifies the temporal variation in the adopted decision rules (compensatory vs. non-compensatory) can address both research gaps. We succinctly review the relevant literature in the following subsection to highlight the limitations of the existing choice models and the need to develop a new model.

1.3 | Literature review

Most of the previous studies modify the systematic part of the indirect utility in DCMs to incorporate different decision rules (Elrod et al., 2004; Swait, 2001). A few studies have used the latent class specification to simultaneously model multiple decision rules considering one-to-one correspondence between a decision rule and a latent class (Dey et al., 2018; Hess et al., 2012; Swait & Adamowicz, 2001). However, these models are static, that is, they assume that the decision-maker uses the same decision rule across all choice occasions. Moreover, these studies do not account for the learning behaviour of the decision-maker and validate models using rather simplistic SP datasets.

Another strand of the literature in psychology suggests using cumulative prospect theory (CPT) and instance-based learning theory (IBLT) for dynamic behaviour modelling. CPT is particularly used in eliciting day-to-day travel mode (Ghader et al., 2019; Yang et al., 2017) and route choices (Jou & Chen, 2013; Yang & Jiang, 2014), but it is not ideal for experience-based learning (Erev et al., 2010; Jou & Chen, 2013). Unlike CPT, IBLT is appropriate for such situations. IBLT relies on the power law of forgetting (Gonzalez et al., 2003). Tang et al. (2017) first illustrate the application of IBLT in DCMs by integrating it into the mixed logit model. The IBLT-based model incorporates temporal variation in preferences of the decision-maker to some extent but assumes that decision-makers adhere to a compensatory decision rule.

Our review suggests that, ironically, the existing DCMs that account for heterogeneity in decision rules are not dynamic, and the ones that incorporate preference dynamics assume decision-makers to follow a compensatory decision rule. Thus, the off-the-shelf DCMs cannot be used to estimate riders' valuation of transit crowding.

1.4 | Contributions

In this study, we propose a dynamic latent class model (DLCM) which integrates IBLT-based expectation formation framework to account for learning behaviour of riders, specifies compensatory and non-compensatory (i.e. inertia/habit) choice processes of riders as latent classes, and allows riders to dynamically transition between these classes based on the differences between the expected and the experienced level of services and other historical attributes. Apart from heterogeneity in decision rules, our model can also account for the unobserved heterogeneity in preferences of riders. DLCM has a structure similar to the non-homogeneous Hidden Markov Models (HMMs)—the latent model of DLCM includes the transition between decision rules (i.e. states), and the measurement model represents the choice model conditional on the decision rule. Whereas the latent model constitutes a series of logistic regression models with time-invariant and time-varying covariates (Bartolucci et al., 2012, 2014; Vermunt et al., 1999), the measurement model relies on a DCM with choice-specific covariates (Train, 2009). We extend the expectation–maximization (EM) algorithm for HMMs to estimate DLCM. Finally, we adapt the Viterbi algorithm to predict the most likely sequence of decision rules of a rider, conditional on her observed route choices (Arulampalam et al., 2002).

We illustrate the applicability of the proposed DLCM in addressing the empirical research gaps (as discussed in Section 1.2) by estimating the crowding cost of an Asian metro's riders. In doing so, we calibrate DLCM using a 2-month-long longitudinal dataset on riders' revealed route preferences. We adopt a passenger-to-train-assignment algorithm to recover route choices of riders by integrating AFC and AVL datasets. Whereas parameters of the compensatory class enable us to accurately estimate the crowding multiplier, sequence of decision rules evaluates the validity of the assumption made by previous studies regarding the compensatory decision rule.

The remainder of this paper is organized as follows: Section 2 discusses RP experiment design and the passenger-to-train-assignment algorithm; Section 3 formulates DLCM; Section 4 presents the EM algorithm to estimate DLCM and provides inference procedure; Data processing is detailed in Sections 5 and 6 illustrates the importance of DLCM by investigating the results of the empirical study. Conclusions and future work are discussed in Section 7.

2 | DATA AND EXPERIMENT DESIGN

We use the data from an RP-based route choice experiment. The experiment exploits a unique feature of the network of an Asian metro. Four urban metro lines form an inner-city loop, and thus, an excellent laboratory for RP-based route choice data (see Figure 1 for a visual illustration). We select 32 origin–destination (OD) pairs between Line 1 and Line 2. Passengers on the selected OD pairs can reach their destinations using just two competitive paths, either by transferring to line 3 or line 4. These paths have enough relative variation in travel time and crowding, circumventing the concern of the dominant alternative.

2.1 | Recovering travel experience from automated data

The routes chosen by passengers and the route-specific attributes are obtained by passing the day-to-day data on AFC and AVL through a passenger-to-train assignment algorithm. To recover the crowding experience of observed travellers throughout their journeys, we assign each passenger to a train, and in-vehicle crowding densities are derived by aggregating the number of individual passengers travelling on each train. The assignment algorithm uses a methodology similar to the one introduced by Hörcher et al. (2017). Since Hörcher et al. (2017) use only 1-day AFC–AVL data, they did not focus on the computational efficiency of the assignment algorithm, which takes 2 days in assigning over four million daily trips to trains. However, for this study, we realize a series of efficiency improvements in the assignment algorithm implementation to make the assignment feasible for the 2-month data.

The passenger-to-train assignment algorithm performs the following steps:

- Smart card trips are classified and treated separately according to (a) whether route choice is ambiguous, that is, one or more feasible paths are available between the trip origin and its destination, (b) the number of transfers on the available paths, and (c) whether one or more than one feasible train *itineraries* are available between the tap-in and tap-out times. Note that a train itinerary is feasible if (a) train departure time on the first leg is greater than the tap-in time, (b) train arrival time on the last leg is not greater than the tap-out time and (c) transfer times between journey legs are all positive. Each trip in the AFC data is expected to have at

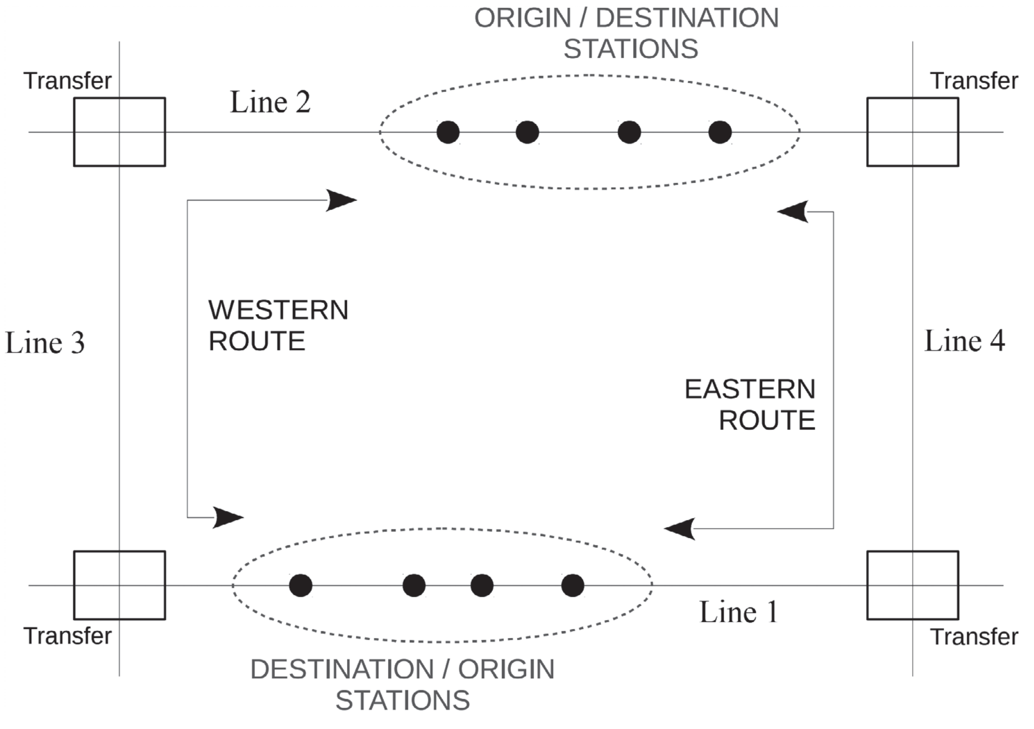

FIGURE 1 Schematic network layout of the metro network

least one feasible itinerary in the AVL dataset. For transfer trips, a train itinerary consists of one specific train in the AVL data on each leg of the journey.

- For trips with only one feasible itinerary, the passenger-to-train assignment is unambiguous: the trip is assigned to the only available train. This leads to a sample of known *egress times* at the destination stations, which help in estimating a non-parametric probability density function (PDF) of the station-specific egress time. Note that egress time is defined as the duration between the train's arrival at the destination and the tap-out time.
- In case of trips with multiple feasible trains but no transfers, the assignment is based on the PDF of the egress times of feasible trains in the data. Hörcher et al. (2017) perform a stochastic assignment in which assignment probabilities are proportional to the egress time densities. In the present application, we assign trips to the most probable itinerary to save computation time. After this stage, we have a representative sample of access times for each origin station, from which we estimate non-parametric PDF of access time.
- In the next step, we focus on one-transfer trips with multiple itineraries. Each itinerary receives an assignment probability based on the PDF of access and egress time, and we select the itinerary with the highest probability product. This process generates a sample of transfer times at each transfer station, which we use to estimate PDF of station-level transfer time.
- Trips with multiple transfers and feasible train itineraries are treated in the same way as the earlier-discussed trip types.
- Trips with multiple feasible routes can be assigned the same way as regular trips, unless the number of transfers differs along the route alternatives. In the latter case, the algorithm selects the most likely route and itinerary based on access and egress time probabilities only.

The computation time of the assignment process depends heavily on the speed of recovering the list of feasible train itineraries from AVL data for individual trips. Substantial computational efficiency gains have been achieved relative to Hörcher et al. (2017) by applying the *Fast Overlap Joins* function available from version 1.12.8 (released in December 2019) of the *data.table* package of R. This data joining technique is ideal for the purpose of the assignment because feasible train itineraries have to satisfy the condition that the train movement interval should be entirely *within* the check-in and check-out intervals. With these amendments, the computation time of a 1-day assignment decreases from the original 2 days to just around 25 min.

2.2 | Key variables

Key trip attributes include travel time, the density of standing passengers, and the probability of standing. Our algorithm infers the probability of standing on the level of OD pairs. The algorithm is detailed in section 3.4.3 of Hörcher et al. (2017). Both the density of crowding and the probability of standing are recovered for each inter-station section of the trip and aggregated to compute the trip- or route-level estimates.

We consider crowding density (measured in passengers per m^2) and standing probability as dummy (or indicator) covariates in all models. We create three indicator variables for crowding density—below one rider/m^2, between one and two riders/m^2 and above two riders/m^2. These categories constitute 39%, 37% and 24% of observations respectively. Similarly, we create three indicator variables for standing probability—below 0.4, between 0.4 and 0.7, and above 0.7. The proportions of observations in these categories are 24%, 45% and 31% respectively.

The observation unit in our analysis is a rider who travels between a specific OD pair. Our dataset covers 2 months, thus allowing for numerous repeated route choice observations of each smart card holder. The smart card number helps us in keeping track of the rider's route choices at several instances. Similar to other RP datasets, we only observe trip attributes (e.g. travel time and crowding level) of the chosen route. We consider that a rider uses IBLT to develop the expectation of attributes on a route based on her previous experiences on that route. A modeller can only know the rider's expectation of attributes on both routes, if she has chosen them at least once in the past. In the absence of knowledge about (expectation of) attributes of both routes, estimation of a compensatory choice model is infeasible. Therefore, we define the initialization period T_I for each rider based on the time until both alternatives are chosen by a rider. For example, if we observe a rider at 10 occasions with the following route choice vector $\{1, 1, 2, 1, 1, 2, 1, 1, 2, 1\}$, we use four occasions to model initialization probabilities in the latent model ($T_I = 4$) and the remaining six occasions for measurement model and transition probabilities in latent model ($T = 6$). Based on the sequence of route choices of smart card holders, we also create new covariates for the DLCM model. Table 1 summarizes all key attributes.

3 | MODEL FORMULATION

At each time step t, we observe the route choice of a smart card holder i, denoted by $y_{it} \in \{1 : \text{route } 1;\ 2 : \text{route } 2\}$. Similar to HMM, we assume an underlying latent process $\boldsymbol{s}$, which

TABLE 1 Description of covariates used in discrete choice models

Attributes	Description
Trip attributes	
Crowding0	Dummy variable for the crowding level below one rider/metre^2 (base)
Crowding1	Dummy variable for the crowding level between one and two riders/metre^2
Crowding2	Dummy variable for the crowding level above two riders/metre^2
SP0	Dummy variable for the standing probability below 0.4 (base)
SP1	Dummy variable for the standing probability between 0.4 and 0.7
SP2	Dummy variable for the standing probability above 0.7
Travel time	Travel time of trip (in hours)
Attributes derived from a sequence of route choices	
Highest choice proportion	(Number of times most chosen route is chosen)/$(T_I + T)$
Proportion of total transitions	(Total number of route transitions)/$(T_I + T)$
Proportion of last transitions	(Number of route transitions until time step t)/t
Choice proportion for a route	(Number of times a route is chosen until time step t)/t
Lagged choice	This dummy takes value 1 at time step t for the route that is chosen at time step $t - 1$

follows first-order Markov chain with state space of two decision rules {1 : compensatory; 2 : inertia}.

In traditional HMMs, conditional on the latent process $\boldsymbol{s}$ and the available exogenous covariates, $\{y_{i1}, \ldots, y_{iT}\}$ are independent (Bartolucci et al., 2014). In the proposed DLCM model, this condition holds for the compensatory class ($s = 1$), but the distribution of y_{it} conditional on the inertia/habit class ($s = 2$) also depends on the lagged choice $y_{i(t-1)}$.

As discussed in the last subsection, we observe a rider for $T_I + T$ periods. The first T_I choices of a rider are used to calibrate the initialization probability, and the remaining T choice occasions are used to calibrate the transition probability and measurement model. However, we do not discard the first T_I choices while computing IBLT-based expectations of route-specific attributes and deriving covariates from route choice sequences of riders.

The measurement model conditional on the compensatory class only includes variables related to the expectations of the route-specific characteristics at the time of choice. On the other hand, the measurement model conditional on the inertia class and latent model include variables associated with the characteristics of the specific route chosen in the last period and aggregate characteristics of the past route choice sequence. In the absence of riders' demographics, we also derive covariates based on the entire (i.e. $T + T_I$) choice sequence of riders and use them in the latent model to capture the heterogeneity across riders.

In this section, we describe the measurement (i.e. discrete choice) model, followed by different components of the latent model. We formulate DLCM for two alternatives and two latent classes, but without loss of generality, it can be extended to any number of alternatives and latent classes.

3.1 | Measurement model

3.1.1 | Conditional on the compensatory class

Conditional on rider i being in the compensatory class at time t (i.e. $s_{it} = 1$), she is assumed to make route choice based on the random utility maximization theory (Train, 2009). She derives the following utility from choosing route j at time t:

$$U_{itj} = \boldsymbol{\phi}_i \boldsymbol{Q}_{itj} + \nu_{itj} = \boldsymbol{\gamma}\mathfrak{S}(\boldsymbol{F}_{itj}) + \chi_i \mathfrak{S}(\boldsymbol{G}_{itj}) + \nu_{itj}, \quad \text{where} \quad \chi_i \sim \text{Normal}(\boldsymbol{\varrho}, \boldsymbol{\Psi}). \tag{1}$$

The vector of observed covariates $\boldsymbol{Q}_{itj}$ constitutes two types of attributes—$\boldsymbol{F}_{itj}$ and $\boldsymbol{G}_{itj}$. The marginal utility associated with attribute vector $\boldsymbol{F}_{itj}$ does not vary across riders, but preference heterogeneity is present for attributes $\boldsymbol{G}_{itj}$. χ_i is assumed to follow normal distribution, but any semi-parametric mixing distribution can be specified depending on the context (see Bansal et al., 2018a, for review). We compute the expected value of $\boldsymbol{F}_{itj}$, that is, $\mathfrak{S}(\boldsymbol{F}_{itj})$, using IBLT (Tang et al., 2017):

$$\mathfrak{S}(\boldsymbol{F}_{itj}) = \sum_{t'} W_{t't}\boldsymbol{F}_{it'j}, \quad \text{where} \quad W_{t't} = \frac{[t-t']^{-\mu}}{\sum_{\tau}[t-\tau]^{-\mu}}. \tag{2}$$

t' and τ belong to a set of past times for which rider i actually observed the variable for route j and μ is a memory decay parameter that captures the rate of forgetting the past experiences. $\mathfrak{S}(\boldsymbol{G}_{itj})$ is also computed using Equation (2). If ν_{itj} is Gumbel-distributed idiosyncratic disturbance term, the probability of choosing route k by rider i at time t, conditional on ride i being in the compensatory class at time t, is:

$$P(y_{it} = k | s_{it} = 1; \boldsymbol{\gamma}, \boldsymbol{\varrho}, \boldsymbol{\Psi}, \mu) = \frac{\exp(V_{itk})}{\exp(V_{it1}) + \exp(V_{it2})}. \tag{3}$$

3.1.2 | Conditional on the inertia class

If a rider is in the inertia class at time t, she is more likely to choose the same route at time t as chosen at $t-1$. Based on this observation, the probability of choosing route 1 by rider i at time t, conditional on ride i being in the inertia class at time t (i.e. $s_{it} = 2$), is:

$$P(y_{it} = 1 | s_{it} = 2, y_{i(t-1)}; \lambda_1, \lambda_2) = \frac{\exp\left(\lambda_1 \mathbb{1}[y_{i(t-1)} = 1] - \lambda_2 \mathbb{1}[y_{i(t-1)} = 2]\right)}{\exp\left(\lambda_1 \mathbb{1}[y_{i(t-1)} = 1] - \lambda_2 \mathbb{1}[y_{i(t-1)} = 2]\right) + 1}, \tag{4}$$

where $\mathbb{1}[.]$ is an indicator function. λ_j captures the likelihood of a passenger to choose route j at time t if $y_{i,(t-1)} = j$ and $s_{it} = 2$. In other words, if λ_1 is higher than λ_2, riders on average encounter more inertia on route 1 compared to route 2. Moreover, if our modelling hypothesis is consistent with the observed data, we would expect λ_1 and λ_2 to be positive. Some route-specific attributes derived from the historical choice sequence of a rider can also be incorporated in the systematic utility. We choose this specific form of choice probability for the inertia class based on the context of the empirical study, but there is a flexibility to capture different non-compensatory decisions rules by increasing the cardinality of the state space of the underlying Markov chain (i.e. with more latent classes).

3.2 | Latent model

3.2.1 | Transition model

A rider's class transition probabilities are likely to depend on the difference between the expected and experienced level of service on the route chosen in the previous period. Moreover, a choice sequence of a rider also provides information about the class of the rider. For example, a consistent route choice across several occasions indicate a rider's inclination towards being in the inertia class.

We model the transition between latent classes using logistic regression. The resulting probability expression for transition from class s to the compensatory class is:

$$P\left(s_{i(t+1)} = 1 | s_{it} = s; \zeta_s, \zeta_s^C, \mu\right) = \frac{\exp\left(m_{its}\right)}{1 + \exp\left(m_{its}\right)},$$
$$\text{where} \quad m_{its} = \zeta_s \left[\boldsymbol{X}_{itj_t} - \mathfrak{S}(\boldsymbol{X}_{itj_t})\right] + \zeta_s^C \boldsymbol{X}_{it}^C, \tag{5}$$

and $\boldsymbol{X}_{itj_t}$ is a vector of attributes (e.g. crowding level) experienced by rider i on chosen route j_t at time t. $\boldsymbol{X}_{it}^C$ is derived from a sequence of choices made by rider i. A proportion of choice transitions made by rider i in the choice sequence observed until time t is one such attribute. The expected value of attributes is obtained using Equation (2).

If $\boldsymbol{X}_{itj_t}$ includes the level-of-service attributes for which "less is better" (e.g. travel time, crowding), we expect ζ_s to be positive. Intuitively, if the experienced level of service is poorer than the prior expectation at time t, a rider is more likely to remain in or switch to the compensatory class at $t + 1$.

3.2.2 | Initialization model

Similar to the transition model, based on the differences between the experienced and expected level of service on the chosen route at $t = T_I$, we obtain the latent class probabilities of a rider after the choice made at $t = T_I$. Thus, the probability of the initial latent class being compensatory is:

$$P\left(s_{i(T_I+1)} = 1; \zeta_0, \zeta_0^C, \mu\right) = \frac{\exp\left(m_{iT_I}\right)}{1 + \exp\left(m_{iT_I}\right)},$$
$$\text{where} \quad m_{iT_I} = \zeta_0 \left[\boldsymbol{X}_{iT_Ij_{T_I}} - \mathfrak{S}\left(\boldsymbol{X}_{iT_Ij_{T_I}}\right)\right] + \zeta_0^C \boldsymbol{X}_{iT_I}^C, \tag{6}$$

and the expectation $\mathfrak{S}(.)$ is computed using Equation (2).

4 | MODEL ESTIMATION

By combining all three components of the model, we write the conditional likelihood of the model for rider i:

$$P(y_{i1},\ \dots\ ,y_{iT}|\boldsymbol{Q}_i,\boldsymbol{X}_i;\boldsymbol{\Theta}) = \sum_{s_1=1}^{2}\sum_{s_2=1}^{2}\ \dots\ \sum_{s_T=1}^{2}\underbrace{\prod_{t=1}^{T}P\left(y_{it}|q_{its_t}=1,y_{i(t-1)}\right)}_{\text{Choice Model}}\ \underbrace{P(q_{i1s_1}=1)}_{\text{Initialization Model}}\ \dots$$

$$\dots\ \underbrace{\prod_{t=2}^{T}P(q_{its_t}=1|q_{i(t-1)s_{(t-1)}}=1)}_{\text{Transition Model}}, \tag{7}$$

where q_{its} is 1 if the passenger i belongs to class s at time t, else it is zero. The probability distributions on the right-hand side of Equation (7) are conditional on the observed explanatory variables but are omitted here. The parameters are stacked in $\boldsymbol{\Theta} = \{\mu, \zeta_0, \zeta_1, \zeta_2, \zeta_0^1, \zeta_1^1, \zeta_2^1, \zeta_0^2, \zeta_1^2, \zeta_2^2, \gamma, \rho, \boldsymbol{\Psi}, \lambda_1, \lambda_2\}$. Figure 2 shows the schematic diagram of the proposed DLCM. This specification can be viewed as a variant of the traditional heterogeneous hidden Markov models where conditional on the latent class, choice probabilities also depend on the lagged choice. The proposed DLCM can be estimated using the Markov Chain Monte Carlo (MCMC) simulation, but the estimation suffers from the label-switching problem (see Papastamoulis, 2014; Spezia, 2009, for a detailed discussion). Since remedies of the label-switching problem are model specific and not straightforward, we rely on the maximum likelihood estimation.

4.1 | Expectation–maximization (EM) algorithm

Direct maximization of the likelihood is challenging due to a well-known risk of underflow (i.e. the product of probabilities is too small to be represented by the CPU) (Netzer et al., 2017). Moreover, the analytical gradient expression of the likelihood is complex and maximization using numerical gradients can result in prohibitively large computation time. To decompose the likelihood maximization into simplified optimization problems, we estimate DLCM using the EM algorithm (Baum et al., 1970; Dempster et al., 1977; Welch, 2003). Readers can refer to Bansal et al. (2018b) and James (2017) to know more about applications of the EM algorithm in the estimation of choice models. We extend the EM algorithm for the heterogeneous HMMs to account for

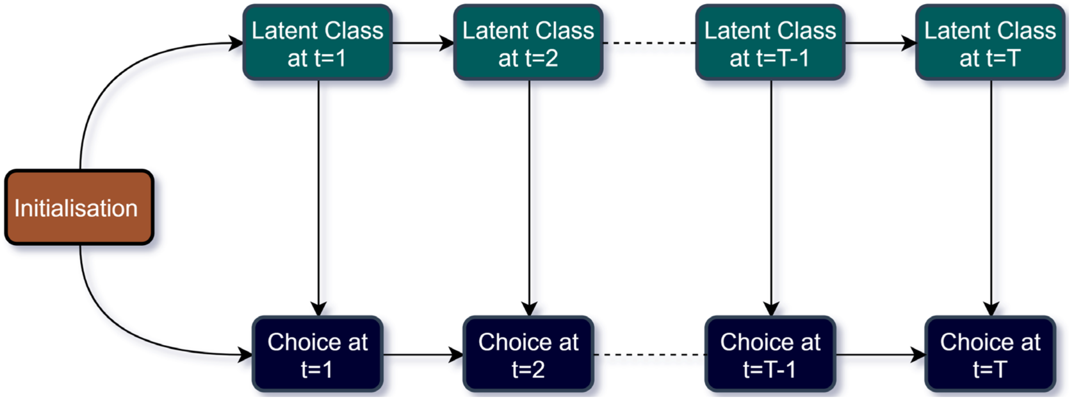

FIGURE 2 The dynamic latent class model [Colour figure can be viewed at wileyonlinelibrary.com]

the auto-correlated choices, preference heterogeneity in the choice model, and riders" learning behaviour.

The EM algorithm was originally developed to deal with the missing data problem. The DLCM likelihood maximization problem also falls under the same category because latent classes can be treated as the missing data. The EM algorithm is a two-step iterative procedure where the conditional expectation of the missing data is obtained in the E-step and then the complete loglikelihood is maximized in the M-step to update the model parameters. The convergence criterion is defined based on the difference in parameter estimates or loglikelihood values of two consecutive iterations.

Assuming the latent class assignment as the missing variable and observations of different riders to be independent, we write the complete loglikelihood $\log L_c$ for a sample of N riders:

$$\begin{aligned}\log L_c =& \log P(\boldsymbol{y}_1, \dots, \boldsymbol{y}_T, \boldsymbol{s}_1, \dots, \boldsymbol{s}_T; \boldsymbol{\Theta})\\ =& \left[\sum_{i=1}^{N}\sum_{t=1}^{T}\sum_{s=1}^{2} q_{its} \log\left[P\left(y_{it}|q_{its}=1, y_{i(t-1)}\right)\right]\right] \dots\\ &\dots + \left[\sum_{i=1}^{N}\sum_{s=1}^{2} q_{i1s} \log\left[P(q_{i1s}=1)\right]\right] \dots\\ &\dots + \left[\sum_{i=1}^{N}\sum_{t=2}^{T}\sum_{s=1}^{2}\sum_{r=1}^{2} [q_{its}q_{i(t-1)r}] \log\left[P(q_{its}=1|q_{i(t-1)r}=1)\right]\right]. \end{aligned} \tag{8}$$

4.1.1 | E-step

Based on the complete loglikelihood $\log L_c$ expression, the E-step in $(k+1)^{\text{th}}$ iteration requires computing the following expectations:

$$\begin{aligned}\pi_{its}^{k+1} &= \mathbb{E}\left[q_{its}|\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right] = P\left[q_{its}=1|\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right],\\ \omega_{itrs}^{k+1} &= \mathbb{E}\left[q_{its}q_{i(t-1)r}|\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right] = P\left[q_{its}q_{i(t-1)r}=1|\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right],\end{aligned} \tag{9}$$

where $\boldsymbol{\Theta}^k$ is a parameter vector in kth iteration of the EM algorithm. To compute expectations in the E-step efficiently, we define forward (α_{its}) and backward (β_{its}) variables as follows:

$$\begin{aligned}\alpha_{its}(\boldsymbol{\Theta}) &= P\big(y_{i1}, \dots, y_{it}, q_{its}=1; \boldsymbol{\Theta}\big)\\ \beta_{its}(\boldsymbol{\Theta}) &= P\left(y_{i(t+1)}, \dots, y_{iT}|y_{it}, q_{its}=1; \boldsymbol{\Theta}\right).\end{aligned} \tag{10}$$

We then compute the π_{its}^{k+1} and ω_{itrs}^{k+1} in terms of forward and backward variables using the Bayes theorem.

$$\begin{aligned}\pi_{its}^{k+1} &= P\left[q_{its}=1|\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right]\\ &= \frac{P\left(\boldsymbol{y}_i|q_{its}=1; \boldsymbol{\Theta}^k\right) P\big(q_{its}=1; \boldsymbol{\Theta}^k\big)}{P\big(\boldsymbol{y}_i; \boldsymbol{\Theta}^k\big)}\\ &= \frac{P\left(y_{i1}, \dots, y_{it}, y_{i(t+1)}, \dots, y_{iT}|q_{its}=1; \boldsymbol{\Theta}^k\right) P\big(q_{its}=1; \boldsymbol{\Theta}^k\big)}{P\big(\boldsymbol{y}_i; \boldsymbol{\Theta}^k\big)}\end{aligned}$$

$$
\begin{aligned}
&= \frac{P\left(y_{i1}, \ldots, y_{it} | q_{its} = 1; \boldsymbol{\Theta}^k\right) P\left(y_{i(t+1)}, \ldots, y_{iT} | y_{it}, q_{its} = 1; \boldsymbol{\Theta}^k\right) P\left(q_{its} = 1; \boldsymbol{\Theta}^k\right)}{P\left(\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right)} \\
&= \frac{P\left(y_{i1}, \ldots, y_{it}, q_{its} = 1; \boldsymbol{\Theta}^k\right) P\left(y_{i(t+1)}, \ldots, y_{iT} | y_{it}, q_{its} = 1; \boldsymbol{\Theta}^k\right)}{P\left(\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right)} \\
&= \frac{\alpha_{its}\left(\boldsymbol{\Theta}^k\right) \beta_{its}\left(\boldsymbol{\Theta}^k\right)}{P\left(\boldsymbol{y}_i; \boldsymbol{\Theta}^k\right)} \\
&= \frac{\alpha_{its}\left(\boldsymbol{\Theta}^k\right) \beta_{its}\left(\boldsymbol{\Theta}^k\right)}{\sum_{s'=1}^{2} \alpha_{its'}\left(\boldsymbol{\Theta}^k\right) \beta_{its'}\left(\boldsymbol{\Theta}^k\right)}.
\end{aligned}
\tag{11}
$$

From now onward, we omit $\boldsymbol{\Theta}_k$ for brevity.

$$
\begin{aligned}
\omega_{itrs}^{k+1} &= P[q_{its} q_{i(t-1)r} = 1 | \boldsymbol{y}_i] \\
&= \frac{P\left(\boldsymbol{y}_i | q_{its} q_{i(t-1)r} = 1\right) P\left(q_{its} q_{i(t-1)r} = 1\right)}{P(\boldsymbol{y}_i)} \\
&= \frac{\begin{bmatrix} P\left(y_{i1}, \ldots, y_{i(t-1)} | q_{i(t-1)r} = 1\right) P\left(y_{it} | y_{i(t-1)}, q_{its} = 1\right) \\ P\left(y_{i(t+1)}, \ldots, y_{iT} | y_{it}, q_{its} = 1\right) P\left(q_{its} = 1 | q_{i(t-1)r} = 1\right) P\left(q_{i(t-1)r} = 1\right) \end{bmatrix}}{P(\boldsymbol{y}_i)} \\
&= \frac{\alpha_{i(t-1)r} P\left(y_{it} | y_{i(t-1)}, q_{its} = 1\right) \beta_{its} P\left(q_{its} = 1 | q_{i(t-1)r} = 1\right)}{\sum_{s=1}^{2} \alpha_{its} \beta_{its}}.
\end{aligned}
\tag{12}
$$

The computation details of forward α_{its} and backward β_{its} variables are provided in the Appendix A.

4.1.2 | M-step

After computing π_{its}^k and ω_{itrs}^k in the E-step, the complete loglikelihood is maximized to obtain the parameters for $(k+1)^{\text{th}}$ iteration.

$$
\begin{aligned}
\boldsymbol{\Theta}^{k+1} = \arg\max_{\boldsymbol{\Theta}} \Bigg[& \sum_{i=1}^{N} \sum_{t=1}^{T} \sum_{s=1}^{2} \pi_{its}^k \log\left[P\left(y_{it} | q_{its} = 1, y_{i(t-1)}\right)\right] \\
& + \sum_{i=1}^{N} \sum_{s=1}^{2} \pi_{i1s}^k \log\left[P\left(q_{i1s} = 1\right)\right] \\
& + \sum_{i=1}^{N} \sum_{t=2}^{T} \sum_{s=1}^{2} \sum_{r=1}^{2} \omega_{itrs}^k \log\left[P\left(q_{its} = 1 | q_{i(t-1)r} = 1\right)\right] \Bigg].
\end{aligned}
\tag{13}
$$

We adopt the one-dimensional grid search approach to select the value of the IBLT parameter μ and all other parameters are estimated using the EM algorithm at a given value μ. There are two challenges in estimating the value of μ using the EM algorithm. First, since all three components of the model include the learning parameter μ, the entire objective function is required to be optimized at once. This defies the purpose of the EM algorithm to decompose a complex optimization problem into simpler ones. Second, as the value of μ changes at each iteration of the algorithm, we also need to iteratively update attributes because the expected value of attributes in each component of DLCM depends on μ. Both challenges make the estimation computationally expensive and numerical issues also prevail. However, in our one-dimension grid search strategy, the expected value of attributes is computed only once at the beginning of the algorithm and parameters in the M-step are updated by solving the simpler optimization problems:

$$\{\boldsymbol{\gamma}, \boldsymbol{\varrho}, \boldsymbol{\Psi}\}^{k+1} = \underset{\{\boldsymbol{\gamma}, \boldsymbol{\varrho}, \boldsymbol{\Psi}\}}{\arg\max} \left[\sum_{i=1}^{N} \sum_{t=1}^{T} \pi_{it1}^{k} \log \left[P\left(y_{it} | q_{it1} = 1, y_{i(t-1)} \right) \right] \right] \tag{14}$$

$$\{\lambda_1, \lambda_2\}^{k+1} = \underset{\{\lambda_1, \lambda_2\}}{\arg\max} \left[\sum_{i=1}^{N} \sum_{t=1}^{T} \pi_{it2}^{k} \log \left[P\left(y_{it} | q_{it2} = 1, y_{i(t-1)} \right) \right] \right] \tag{15}$$

$$\left\{\zeta_0, \zeta_0^C\right\}^{k+1} = \underset{\{\zeta_0, \zeta_0^C\}}{\arg\max} \sum_{i=1}^{N} \sum_{s=1}^{2} \pi_{i1s}^{k} \log \left[P(q_{i1s} = 1) \right] \tag{16}$$

$$\left\{\zeta_1, \zeta_1^C\right\}^{k+1} = \underset{\{\zeta_1, \zeta_1^C\}}{\arg\max} \sum_{i=1}^{N} \sum_{t=2}^{T} \sum_{s=1}^{2} \omega_{it1s}^{k} \log \left[P(q_{its} = 1 | q_{i(t-1)1} = 1) \right] \tag{17}$$

$$\left\{\zeta_2, \zeta_2^C\right\}^{k+1} = \underset{\{\zeta_2, \zeta_2^C\}}{\arg\max} \sum_{i=1}^{N} \sum_{t=2}^{T} \sum_{s=1}^{2} \omega_{it2s}^{k} \log \left[P(q_{its} = 1 | q_{i(t-1)2} = 1) \right] \tag{18}$$

Whereas the update for $\{\boldsymbol{\gamma}, \boldsymbol{\varrho}, \boldsymbol{\Psi}\}$ is equivalent to estimating a weighted mixed multinomial logit model, other updates are analogous to the estimation of weighted multinomial logit models. Even after this simplification, the estimation is computationally expensive. For instance, if the EM takes 1000 iterations to converge, the estimation of DLCM involves the estimation of 1000 mixed logit models. For this reason, we rely on the Broyden-Fletcher–Goldfarb–Shanno (BFGS) algorithm with analytical gradients to optimize these functions in Python (see appendix of Krueger et al., 2021, for analytical gradients).

In the EM estimation of DCMs, the information matrix is generally obtained by taking cross-product of M-step score vectors at convergence (James, 2017; Train, 2008). However, Bansal et al. (2018b) illustrate that standard errors estimated using this procedure can be biased. Therefore, we only use the EM algorithm to get the point estimates of parameters and standard errors are obtained by numerically computing the Hessian of true conditional likelihood at the estimated parameters. We use *numdifftools* library in python to numerically compute the Hessian.

4.2 | Most likely sequence of latent classes

Multiple sequences of latent class can result in the observed route choice sequence of a rider, but such sequences of decision rules have different probabilities. Conditional on the sequence of observed route choices and parameter estimates, we estimate the most likely sequence of a rider's latent classes using the Viterbi algorithm (Arulampalam et al., 2002; Forney, 1973; He, 1988). The Viterbi algorithm is a dynamic programming algorithm that recursively computes the probability that the Markov chain for the rider i is in latent class j after t time steps while passing through the most likely sequence of latent classes $s_{i1}, \ldots, s_{i(t-1)}$. This probability can be formally written as $P_{itj} = \max_{s_1, \ldots, s_{(t-1)}} P(s_{i1}, \ldots, s_{i(t-1)}, y_{i1}, \ldots, y_{i(t-1)}, s_{it} = j | \boldsymbol{\Theta})$, where the most probable sequence of latent classes until time $t-1$ is represented by $\max_{s_1, \ldots, s_{(t-1)}}$

operator (Blunsom, 2004). The probability is recursively computed using the following expression: $P_{itj} = \max_k P_{i(t-1)k} a_{ikj} b_j(y_{it})$, where a_{ikj} is the transition probability from latent class k to j and $b_j(y_{it})$ is the probability of observing the route choice y_{it} conditional on the latent class j for rider i at time t. In sum, we compute most likely sequence of latent classes by keeping track of the path of latent classes that led to each latent class and at the end, back trace the best path to the beginning.

Since a rider's choice behaviour (compensatory vs. inertia) is characterized by these classes, the retrieved sequence will provide information about the extent of a rider's inertia. This analysis enables us to evaluate the behavioural assumption of previous studies.

4.3 | Monte Carlo study

To assess the recovery of parameters, the potential convergence issues in the EM estimation and the prediction performance of the Viterbi algorithm, we present a Monte Carlo study as the supplementary material. The results of this study show a good recovery of parameters. Gradient values for all parameters are close to zero at convergence that indicates good convergence properties of the estimator. The prediction accuracy of the Viterbi algorithm in the simulation study is above 80%. We also provide a python code to replicate the Monte Carlo study as supplementary material.

5 | DATA PROCESSING

We follow four sequential screening steps to obtain an appropriate sample for the analysis.

1. We only include riders with at least five observed choices over 2 months. This criterion is satisfied by 20,960 riders in the population.
2. We discard riders who frequently travel between two different OD pairs in the same direction (i.e. on the same route) to avoid the mixed learning and contamination of attribute expectations. Specifically, if a rider travels more than two times on other OD pairs in the same direction of the most travelled OD pair, the rider is excluded from the analysis. We are left with 16,328 riders after this filtering.
3. The riders who choose only one route or transition to the least-chosen route only once across observed choice occasions are discarded from the analysis. Not to our surprise, this criterion eliminates 13,305 riders (81.4%), leaving us with 3023 riders for further analysis. This observation empirically validates our hypothesis that a large proportion of riders follow inertia-based decision rules while choosing subway routes.
4. To ensure that the estimation of transition model utilizes preferences of a rider, the rider should at least have two available occasions ($T \geq 2$) after eliminating initial choice instances T_I. After applying other minor filtering criteria, we are left with 1947 riders and 1,921 riders, respectively, when we consider one (memory = 1) and two past choice occasions (memory = 2) in the computation of attributes using IBLT (see Equation 2). We also conduct analysis assuming contribution of the past three choices (memory = 3) on a route in computing IBLT-based expectations, but we do not present its results here because lower memory specifications result in better model fit.

6 | EMPIRICAL RESULTS

6.1 | Results of static models

As fares are not differentiated based on the route chosen, we derive crowding cost valuations in terms of the equivalent travel time loss—the ratio of the marginal utility of travel time under crowded and uncrowded conditions. Therefore, interactions between travel time and crowding characteristics are included in the model. Readers are referred to (Tirachini et al., 2017) and (Wardman & Whelan, 2011) to see the evaluation of the crowding-dependent value of travel time multipliers.

To highlight the challenges in modelling a dynamic choice process using static models, we estimate them and present their results. We first estimate the multinomial logit (MNL) model using the data created to estimate the choice component of DLCM. The MNL estimates for memories 1 and 2 are presented in Table 2. While computing expectation of route-specific attributes using IBLT, the memory decay parameter μ is considered to be 1. For both samples, the coefficient of travel time is statistical significant (p < 1E-16) and has an intuitive sign. The positive sign on the interaction of travel time with crowding1 × SP2 and crowding2 × SP2 indicate that increase in the expected value of crowding and standing probability decreases the marginal disutility of the expected travel time. However, the marginal disutility of travel time should increase with the increase in crowding characteristics (Bansal et al., 2019; Hörcher et al., 2017; Kroes et al., 2014; Tirachini et al., 2016). Therefore, the direction of the effects of SP2 interaction terms is counter-intuitive. The results remain virtually the same for different values of μ. The sample also includes riders who follow inertia-based decision rules and modelling their choices using a compensatory model might have resulted in this discrepancy.

With the possibility to segment riders following different decision rules, we estimate the static latent class MNL (LC-MNL) model and results are presented in Table 3. The improved likelihood suggests that LC-MNL explains the choice process slightly better, but both sign and magnitude of marginal utilities of interaction terms remain counter-intuitive.

TABLE 2 Results of the multinomial logit model (dependent variable: route choices of Asian metro users)

	Memory = Last choice		**Memory = Last two choices**	
	Estimate	**Std. error**	**Estimate**	**Std. error**
Travel time (in hours)	−23.30	0.427***	−22.69	0.436***
Travel time × crowding1 × SP1	−0.41	0.118***	−0.36	0.118**
Travel time × crowding1 × SP2	0.28	0.181·	0.62	0.191***
Travel time × crowding2 × SP1	−0.36	0.237·	−0.26	0.233·
Travel time × crowding2 × SP2	0.23	0.163·	0.35	0.171*
Constant (route 1)	−0.19	0.029***	−0.17	0.029***
Number of observations	1947		1921	
Total available choice occasions	13,447		13,282	
Loglikelihood at convergence	−6247.5		−6303.8	

p-value: [0, 0.001): "***", [0.001, 0.01): "**", [0.01, 0.05): "*", [0.05, 0.32): "·", [0.32, 1.0): ""

TABLE 3 Results of the latent class multinomial logit model (dependent variable: route choices of Asian metro users)

	Memory = Last choice occasion				Memory = Last two choice occasions			
	Class 1		Class 2		Class 1		Class 2	
	Estimate	Std. error	Estimate	Std. error	Estimate	Std. error	Estimate	Std. error
Travel time (in hours)	−28.99	0.689***	−18.85	0.746***	−27.93	0.722***	−19.09	0.743***
Travel time × crowding1 × SP1	−0.65	0.237**	−0.32	0.216·	−0.78	0.253**	−0.28	0.209·
Travel time × crowding1 × SP2	−0.24	0.341	−0.02	0.380	0.09	0.384	0.08	0.386
Travel time × crowding2 × SP1	−0.13	0.487	−0.75	0.428·	−0.18	0.494	−0.32	0.404
Travel time × crowding2 × SP2	0.32	0.323·	−0.26	0.304	0.85	0.368*	−0.42	0.321·
Constant (route 1)	0.92	0.053***	−1.51	0.067***	1.01	0.056***	−1.47	0.065***
Class-probability parameter	−0.53	0.022***			−0.48	0.021***		
Number of observations	1947				1921			
Total available choice occasions	13447				13282			
Loglikelihood at convergence	−5387.9				−5375.0			

p-value: [0, 0.001): "***", [0.001, 0.01): "**", [0.01, 0.05): "*", [0.05, 0.32): "·", [0.32, 1.0): " "

6.2 | Results of the dynamic latent class model

We estimate the dynamic latent class model (DLCM) for both values of the memory decay parameter μ. In the final specification, we keep covariates with z-value greater than 1 ($p < 0.32$) and set the EM convergence criterion to 10^{-8}. Gradient values of all parameters at convergence are close to zero in all specifications, ensuring convergence to a local optimum. Since parameter estimates are not very sensitive to the value of memory decay parameter μ, we first set μ to 1 and find the model specification. Conditional on this specification, we obtain the optimum value of μ through grid search.

The parameter estimates of the proposed DLCM for $\mu = 1$ are shown in Table 4. We find the value of μ through grid search for DLCM with the memory of two previous choice occasions. To

TABLE 4 Results of the proposed dynamic latent class model (dependent variable: route choices of Asian metro users)

	Memory = Last choice		Memory = Last two choices	
	Estimate	Std. err	Estimate	Std. error
Initialization probability				
Constant	14.55	3.124***	11.66	3.113***
Highest choice proportion	−21.27	3.448***	−18.15	3.359***
Proportion of total transitions	5.53	1.895**	6.65	1.959***
Transition model (compensatory class)				
Constant	−2.61	0.550***	−3.45	0.776***
Standing probability	1.17	0.909·	1.37	1.330·
Proportion of last transitions	−14.80	3.362***	−27.41	6.394***
Proportion of total transitions	24.17	3.998***	38.30	7.707***
Transition model (inertia class)				
Constant	−5.09	0.387***	−5.21	0.373***
Standing probability	1.72	0.917·	0.88	0.784·
Proportion of last transitions	−58.92	6.019***	−57.18	6.275***
Proportion of total transitions	72.69	6.824***	71.77	7.172***
Choice model (compensatory class)				
Travel time (in hours)	−5.23	0.648***	−5.00	0.662***
Travel time × crowding1 × SP1	−0.22	0.194·	−0.29	0.198·
Travel time × crowding1 × SP2	−0.49	0.298·	−0.44	0.318·
Travel time × crowding2 × SP1	−0.65	0.379·	−0.79	0.383*
Travel time × crowding2 × SP2	−0.87	0.266***	−0.97	0.290***
Constant (route 1)	−0.10	0.047*	−0.06	0.047·
Choice model (inertia class)				
Lagged choice	1.49	0.104***	1.50	0.082***
Choice proportion for a route	5.83	0.325***	5.38	0.245***
Constant (route 1)	−0.13	0.092·	−0.11	0.079·
Number of observations	1947		1921	
Total available choice occasions	13,447		13,282	
Loglikelihood at convergence	−3740.6		−3734.9	

p-value: [0, 0.001): "***", [0.001, 0.01): "**", [0.01, 0.05): "*", [0.05, 0.32): "·", [0.32, 1.0): " "

find the optimum value of μ, we create a one-dimension grid between 0.5 and 1.5, at an increment of 0.1. The resulting loglikelihood values at each grid point are presented in Figure 3. The results indicate that the optimum value of μ is 1 and therefore, the specification presented in Table 4 remains the *final* specification.

6.2.1 | Measurement (Choice) model

We first discuss the results of the measurement model conditional on the compensatory class. In our specification, *Travel Time* implies travel time at the base crowding density

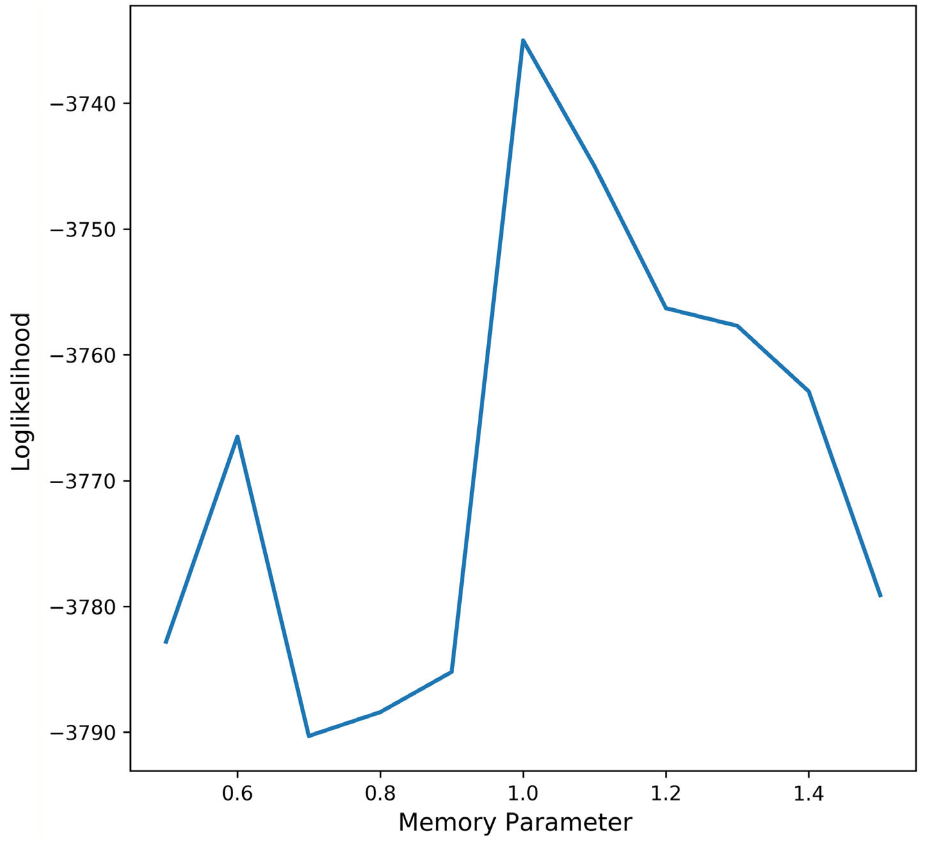

FIGURE 3 One-dimensional grid search to select the memory decay parameter (μ) [Colour figure can be viewed at wileyonlinelibrary.com]

(< 1 riders/metre^2) and the base standing probability (<0.4). Unlike MNL and LC-MNL, sign and magnitude of parameter estimates of the level-of-service attributes are intuitive in the compensatory class of DLCM (see Table 4). In other word, the marginal disutility of travel time is monotonically increasing with the crowding level. Values of travel time under crowding1-SP1, crowding1-SP2, crowding2-SP1 and crowding2-SP2 are 1.06, 1.09, 1.16 and 1.19 times that of the value in the base condition respectively. Since only 3% of riders in our sample have experienced crowding levels above three riders/metre^2, the value of travel time at crowding density greater than two riders/metre^2 (crowding2) can be considered as the value of travel time when crowding density is between two and three riders/metre^2. If we further linearly extrapolate the values of travel time under crowding1-SP2 and crowding2-SP2, metro riders' valuation of travel time appears to increase by around 47% in extremely crowded condition (crowding levels between five and six riders/metre^2 and standing probability above 0.7) relative to the one obtained under uncrowded conditions (crowding density less than one riders/metre^2 and standing probability less than 0.4). We also explore the heterogeneity in crowding cost by specifying a lognormal distribution on the marginal disutility of travel time, but the standard deviation does not turn out be statistically significant ($p > 0.5$).

Since studies based on SP experiments are likely to over-estimate the crowding cost, we compare our crowding cost estimates with those of RP and RP-SP studies (see Tirachini et al., 2017, for an international comparison). The comparison indicates that the crowding cost estimates of DLCM are lower than the crowding multipliers reported by previous studies. In the RP study of Hong Kong MTR, Hörcher et al. (2017) find that increase in the disutility of travel time due to an additional rider per square metre is 0.12 times of the base disutility. Similar results are reported based on SP–RP experiments in Santiago (Batarce et al., 2015) and Paris (Kroes et al., 2014).

There are two potential reasons why our results are not directly comparable to those of previous studies. First, we are using a non-random sample of regular riders to estimate the model—our sample excludes more than 80% of riders who do not switch routes during our study period (see step 3 in Section 5). These riders might have compensated travel time with crowding when they started using metro, but do not update their route preferences during this experiment. Note that a lack of variation in route preferences does not imply that those riders are insensitive to crowding. Their actual crowding valuation can be lower or higher than our estimates but is not empirically identified due to lack of data on their initial preferences. Whereas earlier studies could identify the crowding cost of this subgroup of riders by making a strong assumption that these riders also make compensatory choices at all occasions, we do not make any such assumption and our crowding valuation estimates are therefore identified using only choice occasions when riders make compensatory choices. Second, the specification of the choice model in DLCM is different as compared to earlier experiments because they control for other trip attributes beside travel time and crowding levels.

We now discuss the results of the inertial class. Conditional on being in the inertia class, route choice in the last occasion and the historical frequency of chosen routes could explain the future route choices of a rider. Positive signs on coefficients of both covariates indicate that a rider is more likely to choose the route that she has chosen on the last choice occasion and the one that she has chosen more frequently in the past. The results are consistent with our modelling hypothesis associated with the inertiaclass.

6.2.2 | Initialization and transition model

We first discuss the results of the initialization model. A rider who keeps using the same route and makes fewer transitions over the study period is more likely to be in inertia class at the beginning. These results are aligned with the intuition.

In the transition model for both classes, the same covariates are statistically significant ($p < 0.06$ for all except standing probability). Among the level-of-service attributes, standing probability has a statistically significant association with class transition probabilities ($p < 0.19$). Loosely speaking, positive sign on standing probability indicates that if a rider expects to sit but she has to stand, she is likely to switch to the compensatory class if she is in the inertia class, otherwise is likely to remain in the compensatory class. The proportion of historical transitions at time t and total transitions in the study period also determine the class transition probabilities. Both covariates have different signs and high magnitude. These results imply that keeping the proportion of total transitions constant, if a rider has a higher proportion of historical transitions, she is more likely to choose the same route on the next choice occasion, that is, she is more likely to be in the inertia class. Therefore, a negative sign on the proportion of historical transitions is intuitive and sensible. A positive sign of the proportion of total transitions can be interpreted similarly.

6.2.3 | Recovering latent classes

We also retrieve the latent class of each rider at different choice occasions using the Viterbi algorithm and then we compute the proportion of all choice occasions when riders are in the compensatory class (class 1). On average, riders follow compensatory decision rules only on 25.5% of choice occasions. This proportion is 35.1%, 26.7% and 14.4% for visitors ($T + T_I < 10$), less regular commuters ($10 \leq T + T_I \leq 20$), and regular commuters ($T + T_I > 20$) respectively. In essence, frequent riders are more likely to be in the inertia state.

7 | CONCLUSIONS AND FUTURE WORK

This paper contributes to our understanding of travel behaviour in large urban transport networks, with the aim of improving the precision of economic policy appraisal and supply-side decisions. In the existing literature of demand modelling, transit users' valuation of various travel attributes such as crowding is often estimated by analysing how riders make decisions in various mode or route choice situations. However, previous travel demand studies rely on static choice models, mostly overlooking that riders may form expectations on future travel conditions by learning from past experience. In other words, most of the existing public transport studies assume that riders always follow a fully compensatory decision rule and ignore inertia/habit-based choice behaviour. The present paper's route choice data indicates that more than 80% of riders keep using the same route or shift to another route only once over 2 months. This observation strengthens our belief that incorporating inertia/habits in dynamic choice models is important. No existing choice model can capture such dynamic semi-compensatory behaviour of urban rail riders.

In this study, we propose a DLCM which enables heterogeneity in the decision rules of riders by associating distinct decision rules with latent classes in the observed sample. The model allows riders to transition between latent classes over time, and specifies learning behaviour of travellers using instance-based learning theory. We estimate this model by adapting the EM algorithm and recover the most probable sequence of latent classes for each rider using the Viterbi algorithm. We apply the proposed DLCM to estimate the crowding cost of an Asian metro's riders using a 2-month-long panel data on their revealed route preferences.

From the results of DLCM estimation, we identify several important determinants of route choice behaviour when a rider is in compensatory state. These include the user's expectation on in-vehicle travel time, crowding density and the probability of standing travelling on the available routes. The omission of these factors might lead to biased results in traffic forecasting. The results also enable us to quantify the user valuation of various crowding related nuisance factors in public transit. We find that the marginal disutility of travel time loss under extremely crowded conditions (crowding density between five and six riders/meter^2 and standing probability above 0.7) is around 47% higher than the one obtained under uncrowded conditions (crowding density less than one riders/meter^2 and standing probability less than 0.4). This crowding cost estimate is slightly lower than those reported by previous RP studies. This is perhaps a consequence of an appropriate segmentation of compensatory and inertia-based choice processes. More specifically, the crowding cost estimates in this study are obtained from only choice occasions when riders make compensatory choices.

After aggregating the recovered sequence of latent classes for each rider, we find that an average rider follows a compensatory decision rule only on 25.5% of choice occasions and this proportion further decreases to 14.4% for regular commuters. These results raise a natural question of how the consumer benefit of crowding related interventions should be evaluated for those making non-compensatory choices. Intuition suggests that travellers who are making non-compensatory route choices might also perceive the benefits of crowding reduction. However, the welfare calculation approaches for compensatory choices cannot be directly adopted for other non-compensatory decision rules. The subject of benefit calculation under heterogeneous decision rules thus opens up avenues for future research.

It is worth noting that riders with one or no instance of route switching over the study period are discarded from the analysis due to empirical identification issues. Despite not using the entire population of riders, the implications of our results remain meaningful in two ways. First, crowding valuation estimates remain reliable because they are derived from the measurement model of the compensatory class, but the discarded riders are likely to belong to the inertia class. Second, the estimated proportion of instances when riders follow a compensatory decision rule is an upper bound. We could have analysed the behaviour of the discarded riders using the DLCM model had we observed their preferences since they started using MTR.

We derive the proposed model for two latent classes, but without loss of generality, the model can be extended to multiple classes representing distinct non-compensatory decision rules while accounting for temporal transitions between classes. The potential of DLCM can be explored in understanding dynamic travel behaviour (e.g. preferences for mobility-on-demand services) and various other types of consumer behaviour (e.g. food and healthcare preferences), specially in situations when the longitudinal RP data are easily accessible.

ACKNOWLEDGEMENTS

The authors thank Akshay Vij and Rico Krueger for insightful discussions on dynamic behaviour models. They are also thankful to the joint editor and two anonymous reviewers for their constructive feedback.

ORCID

Prateek Bansal https://orcid.org/0000-0001-6851-8461
Daniel J. Graham https://orcid.org/0000-0002-6971-8737

REFERENCES

Arulampalam, M.S., Maskell, S., Gordon, N. & Clapp, T. (2002) A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2), 174–188.

Bansal, P., Daziano, R.A. & Achtnicht, M. (2018a) Extending the logit-mixed logit model for a combination of random and fixed parameters. *Journal of Choice Modelling*, 27, 88–96.

Bansal, P., Daziano, R.A. & Guerra, E. (2018b) Minorization-maximization (MM) algorithms for semiparametric logit models: bottlenecks, extensions & comparisons. *Transportation Research Part B: Methodological*, 115, 17–40.

Bansal, P., Hurtubia, R., Tirachini, A. & Daziano, R.A. (2019) Flexible estimates of heterogeneity in crowding valuation in the New York City subway. *Journal of Choice Modelling*, 31, 124–140.

Bansal, P., Kessels, R., Krueger, R. & Graham, D.J. (2021) Face masks, vaccination rates and low crowding drive the demand for the London underground during the COVID-19 pandemic. *arXiv preprint arXiv:2107.02394*.

Bartolucci, F., Farcomeni, A. & Pennoni, F. (2012) *Latent Markov models for longitudinal data*. Boca Raton: CRC Press.

Bartolucci, F., Farcomeni, A. & Pennoni, F. (2014) Latent markov models: a review of a general framework for the analysis of longitudinal data with covariates. *Test*, 23(3), 433–465.

Batarce, M., Muñoz, J.C., de Dios Ortúzar, J., Raveau, S., Mojica, C. & Ríos, R.A. (2015) Use of mixed stated and revealed preference data for crowding valuation on public transport in santiago, Chile. *Transportation Research Record*, 2535(1), 73–78.

Baum, L.E., Petrie, T., Soules, G. & Weiss, N. (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, 41(1), 164–171.

Blunsom, P. (2004) Hidden Markov models. *Lecture Notes, August*, 15(18–19), 48.

Dempster, A.P., Laird, N.M. & Rubin, D.B. (1977) Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1), 1–22.

Dey, B.K., Anowar, S., Eluru, N. & Hatzopoulou, M. (2018) Accommodating exogenous variable and decision rule heterogeneity in discrete choice models: application to bicyclist route choice. *PloS One*, 13(11), e0208309.

Elrod, T., Johnson, R.D. & White, J. (2004) A new integrated model of noncompensatory and compensatory decision strategies. *Organizational Behavior and Human Decision Processes*, 95(1), 1–19.

Erev, I., Ert, E., Roth, A.E., Haruvy, E., Herzog, S.M., Hau, R. et al. (2010) A choice prediction competition: choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47.

Forney, G.D. (1973) The viterbi algorithm. *Proceedings of the IEEE*, 61(3), 268–278.

Ghader, S., Darzi, A. & Zhang, L. (2019) Modeling effects of travel time reliability on mode choice using cumulative prospect theory. *Transportation Research Part C: Emerging Technologies*, 108, 245–254.

Gonzalez, C., Lerch, J.F. & Lebiere, C. (2003) Instance-based learning in dynamic decision making. *Cognitive Science*, 27(4), 591–635.

He, Y. (1988) Extended viterbi algorithm for second order hidden Markov process. In: *[1988 Proceedings] 9th international conference on pattern recognition*, IEEE. pp. 718–720.

Hensher, D.A. (1994) Stated preference analysis of travel choices: the state of practice. *Transportation*, 21(2), 107–133.

Hess, S., Stathopoulos, A. & Daly, A. (2012) Allowing for heterogeneous decision rules in discrete choice models: an approach and four case studies. *Transportation*, 39(3), 565–591.

Hörcher, D., Graham, D.J. & Anderson, R.J. (2017) Crowding cost estimation with large scale smart card and vehicle location data. *Transportation Research Part B: Methodological*, 95, 105–125.

James, J. (2017) Mm algorithm for general mixed multinomial logit models. *Journal of Applied Econometrics*, 32(4), 841–857.

Jou, R.-C. & Chen, K.-H. (2013) An application of cumulative prospect theory to freeway drivers' route choice behaviours. *Transportation Research Part A: Policy and Practice*, 49, 123–131.

Kroes, E., Kouwenhoven, M., Debrincat, L. & Pauget, N. (2014) Value of crowding on public transport in île-de-France, France. *Transportation Research Record*, 2417(1), 37–45.

Krueger, R., Bierlaire, M., Daziano, R.A., Rashidi, T.H. & Bansal, P. (2021) Evaluating the predictive abilities of mixed logit models with unobserved inter- and intra-individual heterogeneity. *Journal of Choice Modelling*, 41, 100323.

Netzer, O., Ebbes, P. & Bijmolt, T. H. (2017) Hidden Markov models in marketing. In: *Advanced methods for modeling markets*. Springer, pp. 405–449.

Papastamoulis, P. (2014) Handling the label switching problem in latent class models via the ECR algorithm. *Communications in Statistics-Simulation and Computation*, 43(4), 913–927.

Spezia, L. (2009) Reversible jump and the label switching problem in hidden Markov models. *Journal of Statistical Planning and Inference*, 139(7), 2305–2315.

Swait, J. (2001) A non-compensatory choice model incorporating attribute cutoffs. *Transportation Research Part B: Methodological*, 35(10), 903–928.

Swait, J. & Adamowicz, W. (2001) The influence of task complexity on consumer choice: a latent class model of decision strategy switching. *Journal of Consumer Research*, 28(1), 135–148.

Tang, Y., Gao, S. & Ben-Elia, E. (2017) An exploratory study of instance-based learning for route choice with random travel times. *Journal of Choice Modelling*, 24, 22–35.

Tirachini, A., Sun, L., Erath, A. & Chakirov, A. (2016) Valuation of sitting and standing in metro trains using revealed preferences. *Transport Policy*, 47, 94–104.

Tirachini, A., Hurtubia, R., Dekker, T. & Daziano, R. A. (2017) Estimation of crowding discomfort in public transport: results from Santiago de Chile. *Transportation Research Part A: Policy and Practice*, 103, 311–326.

Train, K.E. (2008) Em algorithms for nonparametric estimation of mixing distributions. *Journal of Choice Modelling*, 1(1), 40–69.

Train, K.E. (2009) *Discrete choice methods with simulation*. Cambridge: Cambridge University Press.

Vermunt, J.K., Langeheine, R. & Bockenholt, U. (1999) Discrete-time discrete-state latent Markov models with time-constant and time-varying covariates. *Journal of Educational and Behavioral Statistics*, 24(2), 179–207.

Wardman, M. & Whelan, G. (2011) Twenty years of rail crowding valuation studies: evidence and lessons from British experience. *Transport Reviews*, 31(3), 379–398.

Welch, L.R. (2003) Hidden Markov models and the Baum-Welch algorithm. *IEEE Information Theory Society Newsletter*, 53(4), 10–13.

Yang, J. & Jiang, G. (2014) Development of an enhanced route choice model based on cumulative prospect theory. *Transportation Research Part C: Emerging Technologies*, 47, 168–178.

Yang, C., Liu, B., Zhao, L. & Xu, X. (2017) An experimental study on cumulative prospect theory learning model of travelers' dynamic mode choice under uncertainty. *International Journal of Transportation Science and Technology*, 6(2), 143–158.

Yap, M., Cats, O. & van Arem, B. (2020) Crowding valuation in urban tram and bus transportation based on smart card data. *Transportmetrica A: Transport Science*, 16(1), 23–42.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Bansal, P., Hörcher, D. & Graham, D.J. (2022) A dynamic choice model to estimate the user cost of crowding with large-scale transit data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 185(2), 615–639. Available from: https://doi.org/10.1111/rssa.12804

APPENDIX A. DETAILS OF THE E-STEP

The required expectations in the E-step depend upon $\alpha_{its}(\mathbf{\Theta})$ and $\beta_{its}(\mathbf{\Theta})$. These forward and backward variables are recursively computed as follows:

$$\begin{aligned}
\alpha_{its} &= P(y_{i1}, \ \dots \ , y_{it}, q_{its} = 1) \\
&= P(y_{i1}, \ \dots \ , y_{it} | q_{its} = 1) P(q_{its} = 1) \\
&= P(y_{i1}, \ \dots \ , y_{i(t-1)} | q_{its} = 1) P(y_{it} | y_{i(t-1)}, q_{its} = 1) P(q_{its} = 1) \\
&= P(y_{it} | y_{i(t-1)}, q_{its} = 1) P(y_{i1}, \ \dots \ , y_{i(t-1)}, q_{its} = 1) \\
&= P(y_{it} | y_{i(t-1)}, q_{its} = 1) \sum_{s'=1}^{2} P(y_{i1}, \ \dots \ , y_{i(t-1)}, q_{i(t-1)s'} = 1, q_{its} = 1) \\
&= P(y_{it} | y_{i(t-1)}, q_{its} = 1) \sum_{s'=1}^{2} P(y_{i1}, \ \dots \ , y_{i(t-1)}, q_{i(t-1)s'} = 1) P(q_{its} = 1 | q_{i(t-1)s'} = 1) \\
&= P(y_{it} | y_{i(t-1)}, q_{its} = 1) \sum_{s'=1}^{2} \alpha_{i(t-1)s'} P(q_{its} = 1 | q_{i(t-1)s'} = 1).
\end{aligned} \tag{A1}$$

We initialize the forward recursion as follows:

$$\begin{aligned}\alpha_{i1s} &= P(y_{i1}, q_{i1s} = 1)\\ &= P(y_{i1}|q_{i1s} = 1)P(q_{i1s} = 1)\end{aligned} \tag{A2}$$

$$\begin{aligned}\beta_{its} &= P(y_{i(t+1)}, \dots, y_{iT}|y_{it}, q_{its} = 1)\\ &= \sum_{s'=1}^{2} P(y_{i(t+1)}, \dots, y_{iT}, q_{i(t+1)s'} = 1|y_{it}, q_{its} = 1)\\ &= \sum_{s'=1}^{2} P(y_{i(t+1)}, \dots, y_{iT}|y_{it}, q_{its} = 1, q_{i(t+1)s'} = 1)P(q_{i(t+1)s'} = 1|q_{its} = 1)\\ &= \sum_{s'=1}^{2} [P(y_{i(t+1)}|y_{it}, q_{i(t+1)s'} = 1)P(y_{i(t+2)}, \dots, y_{iT}|y_{i(t+1)}, q_{i(t+1)s'} = 1) \dots\\ &\quad \dots P(q_{i(t+1)s'} = 1|q_{its} = 1)]\\ &= \sum_{s'=1}^{2} P(y_{i(t+1)}|y_{it}, q_{i(t+1)s'} = 1)\beta_{i(t+1)s'}P(q_{i(t+1)s'} = 1|q_{its} = 1).\end{aligned} \tag{A3}$$

The backward recursion is initiated by considering $\beta_{iTs} = 1 \ \forall i, s$. To validate this recursion, we compute $\beta_{i(T-1)S}$ using Equation (A3) and the considered initial condition.

$$\begin{aligned}\beta_{i(T-1)S} &= \sum_{s'=1}^{2} P(y_{iT}|y_{i(T-1)}, q_{iTs'} = 1)(1)P(q_{iTs'} = 1|q_{i(T-1)s} = 1)\\ &= P(y_{iT}|y_{i(T-1)}, q_{i(T-1)s} = 1).\end{aligned} \tag{A4}$$

The above equation confirms that $\beta_{iTs} = 1 \ \forall i, s$ is appropriate for initialization of the backward recursion.