



Research article

Neural Schrödinger bridge choice modeling through entropy-regularized stochastic transport of latent behavioral preferences

Irsa Sajjad^{1,*} and Mashail M. Al Sobhi^{2,*}

¹ Department of Mathematics, Faculty of Engineering and Computing, National University of Modern Languages, Islamabad, Pakistan

² Department of Mathematics, Umm-Al-Qura University, Makkah 24227, Saudi Arabia

* **Correspondence:** Email: irsa.sajjad@numl.edu.pk, mmsobhi@uqu.edu.sa.

Abstract: Entropy-regularized stochastic transport in latent utility space is explored as a tool for behavioral decision modeling under uncertainty in this study, proposing a neural Schrödinger bridge choice model (NSBCM). The NSBCM differs from traditional discrete choice models, which assume that preferences are either latent utilities or random coefficients that are constant across choice contexts. The NSBCM differs from traditional discrete choice models, which assume preferences are either latent utilities or random coefficients fixed across choice contexts. The terminal latent state is then decoded to alternative-specific choice probabilities using an energy-based decoder. The resulting architecture learns the neural drift field, stochastic transport process, and probabilistic choice decoder simultaneously. We test the NSBCM on the publicly available student travel mode choice dataset and benchmark it against the multinomial logit model, the mixed logit model, the neural choice network model, the variational diffusion choice model, and the energy-based preference network model. The proposed model achieves 95.84% accuracy, an F1 score of 0.947, a negative log-likelihood of 0.214, and a Brier score of 0.081. The ablation, transport stability, and contextual perturbation analyses also suggest that Schrödinger transport, entropy regularization, and energy-based decoding contribute to stable and calibrated predictions. The results corroborate the potential of continuous-time latent transport to represent dynamic and heterogeneous choice behavior under uncertainty.

Keywords: neural Schrödinger bridge choice model; stochastic optimal transport; entropy-regularized diffusion; latent preference evolution; behavioral decision modeling; deep generative behavioral modeling

1. Introduction

Understanding and modeling individual choice behaviors are fundamental to econometrics, machine learning, and the behavioral sciences [1]. Models based on random utility theory, such as the multinomial logit (MNL) and mixed logit (ML) models, provide intuitive utilities and choice probabilities from both observed and unobserved variables [2]. Machine learning methods have gained traction for extracting intricate [3], nonlinear patterns within large choice sets, improving predictive accuracy but at the expense of economic interpretability [4]. Recently, continuous-time stochastic processes have become crucial tools for generative modeling and inference: Schrödinger bridge formulations provide a principled approach for interpolating between prescribed marginal distributions in a minimum entropy [5] sense relative to a reference process [6]. It connects optimal transport, stochastic control, and deep generative models [7], enabling more general latent processes than those in score-based diffusion techniques [8]. For decision-making, it can approximate latent behavioral dynamics and decision-related uncertainty that are not specified by a static utility. Although there have been theoretical developments, the coupling of Schrödinger bridge methods with choice modeling, particularly for modeling latent processes in discrete decisions, remains underexploited. We theoretically ground Schrödinger bridge methodologies in choice analysis, showing that they possess fundamental properties with respect to existence, identifiability, and estimator consistency, and situate them within a rich literature on choice modeling and learning-based generative processes.

Recent advances in statistical machine learning and behavioral decision analysis have made it clear that learning models based on neural architectures are an important approach to modeling complex and heterogeneous choice environments by incorporating latent probabilistic structures. Recent work has successfully applied hierarchical Bayes mixture formulations and nested latent behavioral representations in discrete–continuous decision systems, as well as uncertainty-aware preference modeling [9]. In turn, probabilistic behavioral segmentation has proven effective for adaptive decision analysis under stochastic contextual variation using Gaussian negative binomial latent class frameworks [10]. Multimodal distribution-matching networks and embedded neural behavioral architectures have also demonstrated that deep statistical learning can significantly enhance latent representation capabilities and the consistency of predictive decisions in complex behavioral environments [11]. In addition, hybrid neural choice modeling frameworks [12] reinforce the value of combining statistical interpretation with nonlinear learning of latent utilities for more sophisticated predictive behavioral systems. The proposed neural Schrödinger bridge choice model (NSBCM) framework builds on these developments by combining two novel extensions to existing behavioral learning paradigms: Entropy-regularized Schrödinger transport and uncertainty-aware probabilistic decoding of decisions based on energy. Given these advances, the NSBCM proposes two related extensions of existing behavioral learning models: Entropy regularization of the Schrödinger transport of latent preferences and uncertainty-aware, energy-based decoding of discrete choices. The specific contributions are as follows.

- (i) For the purpose of formulating this discrete choice behavior, we consider a controlled stochastic transport problem between a behavioral prior and a distribution, given the context at the terminal in latent utility space.

-
- (ii) We model preferences' evolution using a stochastic continuous-time differential equation, whose neural drift field models nonlinear, path-dependent behavior adaptation.
 - (iii) We construct an entropy-regularized learning objective and an energy-based terminal decoder, then derive the choice probabilities under latent uncertainty jointly.
 - (iv) We set out conditions for the existence and uniqueness of the controlled diffusion, the existence of the bridge optimizer, identification of the energy decoder, identification of the transport mechanism, and the consistency of the estimator.
 - (v) We analyze the NSBCM on public travel mode choice data and perform predictive, calibration, ablation, transport stability, and perturbation robustness analyses, compared with econometric and neural baselines.

2. Literature review

Random utility models were the starting point of the econometric tradition of discrete choice, and MNL is a basic model because of its closed-form probabilities and interpretation [1]. Extensions such as ML and latent class models address unobserved heterogeneity [2] and enable flexible substitution patterns [9]. Reviews have described a wide range of econometric choice models, including the evolution of nested structures, mixture models, and perturbed utility models, which extend classic random utility models [13]. Integrated choice and latent variable models combine structural equations with discrete choice models to provide a richer behavioral representation, capturing psychological and other factors not observed in the choice process [14]. Preference estimation with unobserved choice-set heterogeneity highlights methodological approaches for handling unobserved choice environments [15].

Alongside econometric models, machine learning classifiers have also gained popularity in recent years for modeling mode choice and other discrete choices, achieving better predictive performance than the widely used econometric logit models in some empirical contexts [16]. A variety of machine learning approaches (e.g., decision trees, neural networks) have been documented in passenger mode choice models, and their interpretability and behavioral consistency challenges have been cataloged [16].

Schrödinger bridges can be interpreted as optimal entropy-regularized stochastic interpolants between distributions and are useful for generative modeling and time-series synthesis [17]. Methodologically, score-based training and Sinkhorn-based estimators estimate the time-varying drifts that specify the bridge processes [18]. Recent theoretical developments of Schrödinger bridges for generative modeling place them in a wider context of diffusion and score models, bridging the gap between the stochastic control and optimal transport interpretations.

Soft terminal constraints relax exact endpoint matching [18], while jump-diffusion bridges can represent discontinuities in time-series data [19]. In travel research, Krueger et al. [20] fused data sources for ride-sourcing demand estimation, Singh Chauhan et al. [21] examined causal structure in travel mode choice, and Grzenda et al. [22] combined survey, global positioning system, environmental, and spatial data for modeling urban mode-choice. Martín-Baos et al. [23] compared machine learning and econometric choice methods on predictive and behavioral criteria. In a separate application, Lai et al. [24] investigated controllable nonlinear dynamics in a memristive network for telemedicine image encryption. Huang and Xu [25] used hierarchical drift-diffusion modeling and electroencephalography to examine choice overload and evidence accumulation. The latter studies are

contextual examples of nonlinear state evolution and decision dynamics, rather than direct baselines for the NSBCM. Neural ordinary differential equations (ODEs) use a neural network to model the derivative of a hidden state, offering a deterministic continuous-depth model [26]. Neural stochastic differential equation (SDEs) build on this notion by adding stochastic "noise" to the latent trajectory, enabling distributions over paths rather than a single deterministic path. Instead, score-based diffusion models specify a forward noising process and train a score field in reverse time, pushing the noise toward a desired data distribution [27]. These formulations offer flexible mechanisms for continuously transforming and representing uncertainty. However, their typical goals are not always to achieve minimum relative entropy with respect to a behavioral reference process or to set a terminal discrete-choice likelihood. This distinction matters because the NSBCM is a controlled neural SDE. This methodological difference is not just due to the inclusion of Brownian noise or continuous time. Instead, the control field is learnt with a path space KL divergence penalty to a reference diffusion, while the terminal potential is computed using an energy-based likelihood of the observed discrete choice. Neural ODE and generic neural SDE models are thus useful nested or ablation models to assess whether bridge regularization adds value beyond continuous-depth or stochastic dynamics, respectively.

The Schrödinger bridge problem chooses the most probable stochastic evolution between given marginal distributions by minimizing the relative entropy with respect to a reference path law [6,8]. Modern computational methods estimate bridge drifts, including maximum likelihood, score matching, iterative proportional fitting, static entropic optimal transport potentials, and simulation-free score and flow matching [5,17]. These advances relate stochastic control, diffusion modeling and entropy-regularized optimal transport.

Most neural Schrödinger bridge studies focus on generating continuous data, synthesizing images, interpolating biological states, or simulating time series. Classic formulations also assume that samples from both endpoint distributions are available. In discrete-choice data, the distribution of the unobservables, namely the terminal preference distribution, cannot be directly observed and is only observed in conjunction with the selected alternative and its covariates. This is because directly importing a conventional two-marginal bridge does not address how a terminal latent distribution should be tied to a categorical behavioral outcome. Another potential drawback is that bridge estimation can be numerically intensive, as it may require multiple drift network calculations or forward-backward simulations.

Efficiency is now an explicit design consideration in advanced transportation models. Fast-TrafficNet combines spatiotemporal neural controlled differential equations with block-term decomposition to compress sparse, high-dimensional traffic information and reduce computation [28]. Ye et al. modeled higher-order intercity relations by prototype-guided spatiotemporal hypergraph fusion and selectively fused auxiliary information [29]. These are not direct NSBCM baselines but methods for forecasting traffic flow or the revitalization index. However, they show that these methods complement predictive comparisons in transportation applications, such as structured compression, sparse relational learning, and explicit runtime evaluation.

This creates a key methodological gap between the two main research streams. Econometric and neural choice models offer alternative-specific probabilities and are typically not based on the minimum relative entropy of stochastic transport for forming latent preferences. In contrast, neural ODE/SDE, diffusion, and Schrödinger bridge methods offer continuous latent transformations and typically do not rely on an energy-normalized discrete-choice likelihood when a terminal latent

distribution is not observed. The NSBCM tackles this interplay by using a latent prior distributed over individuals, a controlled stochastic path law, a KL penalty with respect to a reference diffusion, and an energy-based terminal choice decoder. The contribution is its ability to integrate these components and to mathematically characterize the induced choice-conditioned terminal law. The framework is theoretically grounded in finite discrete-choice problems, while the empirical findings in this study focus specifically on travel mode choice. This distinction is needed to avoid the risk of reading a demonstrated cross-domain generality into the word "extensibility" as in the case of "generality".

3. Methodology

Classical econometric choice models capture various types of behavioral dynamics; however, recent machine learning-based models cannot adequately model continuously evolving latent behavioral dynamics under uncertainty. Most conventional approaches assume static utilities or fixed random coefficients, thereby limiting their usefulness for modeling temporally adaptive preferences or the stochastic evolution of decisions. Recent advances in stochastic optimal transport and Schrödinger bridge theory have inspired the present investigation, which proposes a new generative choice-modeling paradigm that captures path-dependent latent utility transitions via entropy-regularized stochastic transport mechanisms. The proposed approach combines controlled diffusion dynamics, neural latent representations, and energy-based probabilistic decoding within a unified mathematical framework for discrete decision analysis. Figure 1 illustrates the proposed NSBCM. Entropy-regularized stochastic transport of latent behavioral preferences is achieved via continuous-time Schrödinger bridge dynamics, stochastic drift learning, and energy-aware probabilistic decision decoding in heterogeneous contextual environments.




Figure 1. Model architecture of the proposed study.

3.1. Problem formulation and conceptual framework

Consider N decision-makers or choice situations indexed by $n = 1, \dots, N$. In a choice situation n , the decision-maker selects one alternative from a finite choice set $\mathcal{C}_n$. Let $x_{ntj} \in \mathbb{R}^P$ represent the observed qualities of the alternative j , and let $s_n \in \mathbb{R}^D$ represent individual qualities. The choice observed is denoted as $y_{nt} \in \mathcal{C}_{nt}$, or the one-hot encoding $y_{ntj} \in \{0, 1\}$ with $\sum_j y_{ntj} = 1$. We introduce a NSBCM, which is fundamentally different from classical discrete choice theory, in which utilities are treated as fixed latent variables or random coefficients. In our model, preferences are generated by a continuous-time stochastic transport process in latent utility space. In particular, we propose a latent process $z_{nt}(\tau) \in \mathbb{R}^K$, which exists in a normalized time domain $\tau \in [0, 1]$, which serves as a prior distribution that transforms into a context-conditioned posterior distribution. This model views decision-making as an active process of updating latent preferences rather than a fixed assessment of utilities, thereby incorporating behavioral heterogeneity into a stochastic control problem. The first latent variable follows a prior distribution conditioned on personal attributes as follows:

$$z_{nt}(0) \sim p_0(z | s_n), \quad (1)$$

where the terminal distribution $p_1(z | X_{nt}, s_n)$ is implicitly defined in terms of the observed choice outcomes and the choice environment $X_{nt} = \{x_{ntj}\}_{j \in \mathcal{C}_{nt}}$. The main modeling goal is to define the transformation between p_0 to p_1 as the most likely stochastic path, given that it is as close as possible to a reference diffusion process.

3.2. Controlled stochastic dynamics in latent utility space

The latent evolution of preferences follows a SDE

$$dz_{nt}(\tau) = u_\theta(z_{nt}(\tau), \tau, X_{nt}, s_n) d\tau + \sigma dW_\tau, \quad (2)$$

where u_θ is a parameterized drift (control) function, W_τ denotes a parameterized standard Brownian motion, and $\sigma > 0$ is a diffusion constant that controls the stochastic variability. The drift function is implemented with a deep neural network, allowing it to model nonlinear relationships between state and contextual variables. The uncontrolled reference process is represented by

$$dz(\tau) = \sigma dW_\tau, \quad (3)$$

which characterizes a Wiener measure on path space. The control u , which allows the model to control latent trajectories, steers them towards distributions that align with the observed choices, and is close to the reference diffusion in a probabilistic sense. This formulation incorporates preference formation within a stochastic optimal control paradigm, in which heterogeneity arises from both the initial conditions and path-dependent evolution.

3.3. Schrödinger bridge formulation and optimal transport perspective

Let P_s^0 denote the reference path law generated by the uncontrolled diffusion in Eq (3), with the initial marginal $\mu_0(\cdot | s)$. The classical two-marginal Schrödinger bridge is defined as

$$Q^{\text{SB}} = \arg \min_{Q \ll P_s^0} \text{KL}(Q \parallel P_s^0) \text{ subject to } Q_0 = \mu_0, Q_1 = \mu_1. \quad (4)$$

This hard-marginal formulation assumes that both the initial distribution μ_0 and the terminal distribution μ_1 are independently specified. In the present supervised discrete-choice setting, however, the terminal latent preference distribution is not directly observed. Instead, the observed alternative supplies likelihood information about the terminal latent state. The NSBCM therefore uses the softly constrained problem

$$Q^* = \arg \min_{\substack{Q \ll P_s^0 \\ Q_0 = \mu_0}} \{ \mathbb{E}_Q [-\log p_\eta(y, l, Z_1 X s)] + \lambda \text{KL}(Q \parallel P_s^0) \}. \quad (5)$$

For the controlled process

$$dZ_\tau = u(Z_\tau, \tau, X, s) d\tau + \sigma dW_\tau,$$

Girsanov's theorem gives the equivalent stochastic control representation

$$\mathbb{E}^u \left[-\log p_\eta(y, l, Z_1 X s) + \frac{\lambda}{2\sigma^2} \int_0^1 \|u(Z_\tau, \tau, X, s)\|^2 d\tau \right]. \quad (6)$$

The first term promotes terminal states that explain the observed choice, whereas the second penalizes excessive deviation from the reference diffusion. The optimal terminal marginal is

$$q_1^*(z, l, X s y) = \frac{\rho_1^0(z | s) p_\eta(y, l, z X s)^{1/\lambda}}{\int \rho_1^0(v | s) p_\eta(y, l, v X s)^{1/\lambda} dv}, \quad (7)$$

where ρ_1^0 is the terminal density induced by the reference process. Equation (7) explicitly connects the constructed terminal distribution with the marginal constraints of the classical Schrödinger bridge. Rather than imposing an independently observed terminal marginal, the NSBCM obtains a likelihood-tilted terminal marginal by reweighting the reference density according to the probability of the observed choice.

A larger value of λ keeps the controlled path law closer to the reference diffusion, whereas a smaller value concentrates more strongly around terminal states that assign a high probability to the observed alternative. If an independently estimated terminal marginal μ_1 were available, the soft likelihood condition could be replaced by $Q_1 = \mu_1$, thereby recovering the classical two-marginal Schrödinger bridge.

The optimization objective and probabilistic interpretation of the NSBCM differ from those of neural ODEs, neural SDEs, and conventional diffusion models. A choice model based on a neural ODE uses deterministic latent dynamics, $dZ_\tau = f\theta(Z_\tau, \tau, X, s) d\tau$, so a consistent initial state and context yield a consistent terminal state. This formulation suits nonlinear continuous uncertainty, but it does not

represent pathwise stochastic uncertainty or impose a Schrödinger path-space constraint. A standard neural SDE adds noise to the input signal by $dZ_\tau = f_\theta(Z_\tau, \tau, X, s)d\tau + \sigma dW_\tau$.

However, stochastic dynamics alone do not form a Schrödinger bridge. This is because a terminal prediction loss-only trained neural SDE does not need to choose the minimum relative entropy path law with respect to a reference diffusion. Conventional diffusion generative models typically learn the forward noising and reverse denoising processes to approximate an empirical data distribution through score-matching or denoising objectives. They mainly focus on learning the generative distribution. In contrast, the NSBCM models discrete-choice prediction as a form of likelihood-regularized, entropy-minimizing stochastic transport. It jointly reduces choice loss at the end of the path and path-space divergence from the reference process. The unique feature of the NSBCM is therefore not continuous time or stochastic noise per se. It is the objective that tilts toward the most likely latent distribution, connected to a choice-informed terminal distribution via minimum control energy transport between the two distributions. This structure regularizes the full latent path but does not dictate the process based solely on endpoint prediction accuracy.

3.4. Energy-based terminal choice model

The utility of each alternative is determined at the terminal time $\tau = 1$ by the latent state $z_{nt}(1)$ using an energy-based formulation. In particular, we give an energy function. We define an energy function

$$E_{ntj}(z_{nt}(1)) = E_\eta(x_{ntj}, s_n, z_{nt}(1)), \quad (8)$$

where E_η is a neural network that is parameterized by η . The distribution of the conditional choice probability given the latent state takes the form of a Gibbs distribution

$$p(y_{nt} = j \mid z_{nt}(1)) = \frac{\exp(-E_{ntj}(z_{nt}(1)))}{\sum_{m=1}^{J_{nt}} \exp(-E_{ntm}(z_{nt}(1)))}. \quad (9)$$

Marginalizing over the latent state results in

$$p(y_{nt} = j \mid X_{nt}, s_n) = \int \frac{\exp(-E_{ntj}(z))}{\sum_m \exp(-E_{ntm}(z))} p_1(z \mid X_{nt}, s_n) dz, \quad (10)$$

which can capture both observed and unobserved heterogeneity by integrating over the transported latent distribution.

3.5. Variational objective and learning framework

The training goal is based on the path-space KL divergence that admits the representation

$$\text{KL}(Q \parallel P) = \mathbb{E}_Q \left[\frac{1}{2\sigma^2} \int_0^1 \|u_\theta(z(\tau), \tau)\|^2 d\tau \right]. \quad (11)$$

Together with the negative log-likelihood of observed choices, this yields the overall loss function

$$\mathcal{L}(\theta, \eta) = \sum_{n,t} \mathbb{E} [-\log p(y_{nt} | z_{nt}(1))] + \lambda \mathbb{E} \left[\int_0^1 \|u_\theta(z_{nt}(\tau), \tau)\|^2 d\tau \right], \quad (12)$$

where λ balances data fidelity and transport regularization. This target combines likelihood-based estimation with an entropy-regularized control cost to achieve stable, interpretable learning dynamics. The latent process marginal density $p(z, \tau)$ satisfies the Fokker–Planck equation

$$\frac{\partial p}{\partial \tau} = -\nabla_z \cdot (u_\theta p) + \frac{\sigma^2}{2} \Delta p. \quad (13)$$

Substituting the optimal drift yields

$$\frac{\partial p}{\partial \tau} = -\sigma^2 \nabla_z \cdot (p \nabla_z \log \psi_\theta) + \frac{\sigma^2}{2} \Delta p, \quad (14)$$

that controls the nonstationarity of the time distribution of preferences. The equation provides a mechanistic description of how heterogeneity evolves under deterministic control and random diffusion.

3.6. Gradient-based optimization

Let the terminal log-likelihood be defined as

$$\ell_{nt}(z) = \sum_j y_{ntj} \log p(y_{nt} = j | z). \quad (15)$$

The gradient of the objective with respect to the model parameters is given by

$$\nabla_\theta \mathcal{L} = \mathbb{E} \left[\nabla_\theta \ell_{nt}(z_{nt}(1)) + \lambda \int_0^1 u_\theta^\top \nabla_\theta u_\theta d\tau \right]. \quad (16)$$

Using the chain rule, we have

$$\nabla_\theta \ell_{nt}(z(1)) = \nabla_z \ell_{nt}(z(1)) \cdot \frac{\partial z(1)}{\partial \theta}, \quad (17)$$

which demands a distinction by the SDE. In practice, this is accomplished using either stochastic adjoint methods or pathwise reparameterization techniques, enabling efficient end-to-end training.

3.7. Special case and model reduction

In the limiting case where $u_\theta \equiv 0$, the latent process becomes pure diffusion, and $z_{nt}(1) = z_{nt}(0)$. This causes the model to collapse into a latent utility formulation that is identical to conventional ML models. This makes the proposed framework a rigid generalization of the classical discrete choice models.

3.8. Interpretative implications and conceptual contribution

The suggested framework in discrete choice modeling reconceives preferences as dynamically changing objects driven by stochastic transport processes. In contrast to classical models that treat heterogeneity as a time-invariant random variable, the current model accounts for preference formation in a path-dependent manner, where individuals' choices are the final output of a controlled stochastic mechanism. This view combines econometric choice theory, stochastic control, optimal transport, and deep learning, providing improved predictive accuracy and behavioral interpretability.

4. Theoretical foundations of the NSBCM

The proposed NSBCM framework is a flexible stochastic transport model for latent behavioral evolution, but rigorous theoretical guarantees are needed to ensure its mathematical validity, identifiability, and statistical consistency. The importance of the formal analysis of the existence of the induced stochastic process, the well-posedness of the Schrödinger bridge problem associated with it, and the recoverability of the latent parameters becomes evident, particularly since the model combines controlled SDEs, entropy-regularized optimal transport, and energy-based probabilistic decoding in a single latent preference system. The part of the theory developed in the following section is thus the establishment of the existence results for the controlled diffusion process, the characterization of the optimal bridge dynamics, the proof of the identifiability properties of the energy decoder, and the demonstration of the asymptotic consistency of the proposed estimator under appropriate regularity conditions.

4.1. Mathematical setup

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_\tau\}_{\tau \in [0,1]}, \mathbb{P})$ be a filtered probability space on which a K -dimensional standard Brownian motion W_τ . The latent utility state $z_{nt}(\tau) \in \mathbb{R}^K$ evolves over time $\tau \in [0,1]$ by the controlled diffusion, where n is the observation n at time t , k is the number of dimensions (the dimension of the latent space), and z_{nt} refers to the latent utility at that time.

$$dz_{nt}(\tau) = u_\theta(z_{nt}(\tau), \tau, \mathcal{X}_{nt}, s_n) d\tau + \sigma dW_\tau, \sigma > 0, \quad (18)$$

with the initial law

$$z_{nt}(0) \sim p_0(\cdot | s_n). \quad (19)$$

The process that is uncontrolled in the reference is expressed as

$$dz(\tau) = \sigma dW_\tau. \quad (20)$$

Let $P_{s_n}^0$ be the law of the uncontrolled diffusion initiated at $p_0(\cdot | s_n)$, and let $Q_{\theta, \eta}$ be the law of the controlled process with the parameter θ . The terminal-time latent state of $z_{nt}(1)$ determines the alternative-specific energies

$$E_{ntj}(z) = \mathcal{E}_\eta(x_{ntj}, s_n, z), \quad (21)$$

and the conditional choice probabilities

$$p_{\eta}(y_{nt} = j \mid z, \mathcal{X}_{nt}, s_n) = \frac{\exp\{-E_{ntj}(z)\}}{\sum_{m=1}^{J_{nt}} \exp\{-E_{ntm}(z)\}}. \quad (22)$$

The marginal observed data probability is

$$p_{\theta, \eta}(y_{nt} = j \mid \mathcal{X}_{nt}, s_n) = \int_{\mathbb{R}^K} p_{\eta}(y_{nt} = j \mid z, \mathcal{X}_{nt}, s_n) q_{\theta}(z \mid \mathcal{X}_{nt}, s_n) dz, \quad (23)$$

where $q_{\theta}(\cdot \mid \mathcal{X}_{nt}, s_n)$ is the induced terminal law of the controlled diffusion. The Schrödinger bridge interpretation has the path-space goal of

$$\mathcal{J}(Q) = \text{KL}(Q \parallel P_{s_n}^0) - \lambda \mathbb{E}_Q [\log p_{\eta}(y_{nt} \mid z_{nt}(1), \mathcal{X}_{nt}, s_n)]. \quad (24)$$

Equivalently, using Girsanov's theorem, the optimization over the path laws can be written as an optimization over admissible drifts u as follows:

$$\mathcal{J}(u) = \mathbb{E}^u \left[\frac{1}{2\sigma^2} \int_0^1 \|u(z(\tau), \tau, \mathcal{X}, s)\|^2 d\tau - \lambda \log p_{\eta}(y \mid z(1), \mathcal{X}, s) \right]. \quad (25)$$

For compactness of notation below, write $X = (\mathcal{X}, s, y)$.

4.2. Assumptions

The theoretical outcomes are based on regularity, normalization, identification, and convergence assumptions. Each mathematical condition has an implication for empirical discrete-choice modeling. **Assumption A1:** For each fixed X , the drift $u_{\theta}(z, \tau, X)$ is Borel-measurable in (z, τ) , continuous in θ , and globally Lipschitz in z : there is an $L_u < \infty$ less than infinity such that, given $z, z' \in \mathbb{R}^K$, satisfies $\|u_{\theta}(z, \tau, X) - u_{\theta}(z', \tau, X)\| \leq L_u \|z - z'\|$. Moreover, it has linear growth: $C_u < \infty$ exists such that $\|u_{\theta}(z, \tau, X)\| \leq C_u(1 + \|z\|)$. If two similar latent preference states exist, they should be treated similarly, and the size of the behavioral drift may not increase beyond control. Thus, if a small change in latent utility leads to an implausibly sharp change in the implied trajectory of preferences, it is hard to avoid the conclusion that it is a false change in latent utility. From the practitioner's perspective, this means the stable simulation of preference evolution and helps practitioners avoid interpreting a numerical explosion as behavior adaptation.

Assumption A2: For each s , the initial distribution $p_0(\cdot \mid s)$ has a finite second moment. The initial population can be very heterogeneous, skewed, or multimodal but the dispersion of preferences in the underlying population must be finite. We exclude pathological populations with infinite latent variance, and average preference magnitude, uncertainty, and Monte Carlo (MC) summaries are meaningful.

Assumption A3: For each j , $\mathcal{E}_{\eta}(x_j, s, z)$ is continuous in (η, z) , twice continuously differentiable in

z , and the constants $C_E, L_E < \infty$ exist such that $|\mathcal{E}_{\eta}(x_j, s, z)| \leq C_E(1 + \|z\|^2)$, $\|\nabla_z \mathcal{E}_{\eta}(x_j, s, z)\| \leq$

$C_E(1 + \|z\|)$, and $\|\nabla_z \mathcal{E}_{\eta}(x_j, s, z) - \nabla_z \mathcal{E}_{\eta}(x_j, s, z')\| \leq L_E \|z - z'\|$. This is because all alternative

specific energies and the decoder parameters vary smoothly with the latent preference state. The inability to create a discontinuous probability reversal prevents small behavior changes. It also provides stable gradient-based estimation, local sensitivity analysis, and finite integration of the decoder over the latent distribution.

Assumption A4: The parameter spaces Θ and H for θ and η are compact. The drift and decoder parameters are kept within finite admissible regions and are not allowed to become too large. These restrictions can be applied in practice in various ways, such as bounded parameterizations, weight constraints, or regularization. They allow reasonable behavioral responses and make the optimization problem well-posed within the parameter space.

Assumption A5: The normalized energy of one alternative for every $(\mathcal{X}, s, z)$ is defined; equivalently, the energies are distinguished only by a common additive shift. The convenient normalization is $\sum_{j=1}^J E_j(z, \mathcal{X}, s) = 0$. The softmax probabilities depend only on differences in the energy values, and thus, the energy functions are identifiable only up to a common additive shift. We impose a normalization. If we add a constant to all alternatives on the same energy scale, the choice probability does not change. The zero-sum constraint thus constrains the arbitrary location of the energy scale, analogous to choosing a reference utility in a traditional discrete choice model. It does not change the relative preferences or substitution patterns.

Assumption A6: Suppose that if there are two parameter values η, η' , and for almost every $(\mathcal{X}, s, z)$, $E_{\eta,j}(z, \mathcal{X}, s) - E_{\eta,m}(z, \mathcal{X}, s) = E_{\eta',j}(z, \mathcal{X}, s) - E_{\eta',m}(z, \mathcal{X}, s) \forall j, m$, then $\eta = \eta'$. After eliminating the common energy shift using A5, two decoder parameter vectors cannot be identical with respect to the relative attractiveness of the two alternatives for each pair of alternatives. Thus, the effects of relative choice preferences can be correlated with a unique decoder parameterization, not with observationally redundant representations.

Assumption A7: If we assume that two parameters θ, θ' are related $q_{\theta}(\cdot | \mathcal{X}, s) = q_{\theta'}(\cdot | \mathcal{X}, s)$ for almost every $(\mathcal{X}, s)$, then $\theta = \theta'$. The two different transport parameterizations cannot generate the same terminal distribution of latent preferences, given the context. This condition detects the preference evolution mechanism from the terminal choice decoder and provides an opportunity to interpret the estimated drift parameters as a separate representation of behavioral adaptation.

Assumption A8: The class $\{\log p_{\theta,\eta}(y | \mathcal{X}, s) : (\theta, \eta) \in \Theta \times H\}$ is Glivenko–Cantelli. The sample's log-likelihood always converges uniformly to its population counterpart for all parameter values that are possible. For this reason, it is not the preferred model simply because of random sample fluctuations at one parameter setting, if the sample size is large enough. For practitioners, this assumption relates to the stability of out-of-sample population behavior, which is the behavior measured in empirical model fitting.

Assumption A9: There is a unique $(\theta_0, \eta_0) \in \Theta \times H$ that maximizes the population criterion $M(\theta, \eta) = \mathbb{E}_0[\log p_{\theta,\eta}(Y | \mathcal{X}, S)]$. These assumptions are typical of the existence, identification, and consistency theorems in nonlinear latent variable models. At the population level, one combination of the transport and decoder parameters best approximates the actual decision process. This means that there is at most one local optimum that is observationally equivalent to the global optimum, and there

is only one target the estimator can converge to.

4.3. Existence of the controlled diffusion and path measure

Theorem 4.1. Under Assumptions A1 and A2, for each fixed $(\theta, \mathcal{X}, s)$, the SDE $dz(\tau) = u_\theta(z(\tau), \tau, \mathcal{X}, s) d\tau + \sigma dW_\tau$, $z(0) \sim p_0(\cdot | s)$ admits a unique, strong solution $z(\tau)$ on $[0, 1]$. Moreover, $\mathbb{E} \left[\sup_{0 \leq \tau \leq 1} \|z(\tau)\|^2 \right] < \infty$.

Proof. Since u_θ is globally Lipschitz in z and, by Assumption A1, has linear growth, the classical existence and uniqueness theorem for Itô SDEs applies. Therefore, for every square-integrable initial condition $z(0)$, a unique strong adapted solution on $[0, 1]$ exists. It remains to show the second-moment bound. By Itô's formula applied to $\|z(\tau)\|^2$, we have

$$d\|z(\tau)\|^2 = 2z(\tau)^\top u_\theta(z(\tau), \tau, \mathcal{X}, s) d\tau + K\sigma^2 d\tau + 2\sigma z(\tau)^\top dW_\tau.$$

Taking expectations and using the linear growth condition

$$\mathbb{E} \|z(\tau)\|^2 = \mathbb{E} \|z(0)\|^2 + 2\mathbb{E} \int_0^\tau z(r)^\top u_\theta(z(r), r, \mathcal{X}, s) dr + K\sigma^2 \tau.$$

Now

$$|z^\top u_\theta(z, \tau, \mathcal{X}, s)| \leq \|z\| C_u(1 + \|z\|) \leq \frac{C_u}{2} + \left(C_u + \frac{C_u}{2}\right) \|z\|^2,$$

so $C_1, C_2 < \infty$ exists such that

$$\mathbb{E} \|z(\tau)\|^2 \leq \mathbb{E} \|z(0)\|^2 + C_1 \tau + C_2 \int_0^\tau \mathbb{E} \|z(r)\|^2 dr.$$

By Grönwall's inequality

$$\sup_{0 \leq \tau \leq 1} \mathbb{E} \|z(\tau)\|^2 < \infty.$$

The stronger bound with $\sup_\tau \|z(\tau)\|^2$ follows from the Burkholder–Davis–Gundy inequality applied to the martingale term $2\sigma \int_0^\tau z(r)^\top dW_r$, followed by Grönwall again. Hence, we have

$$\mathbb{E} \left[\sup_{0 \leq \tau \leq 1} \|z(\tau)\|^2 \right] < \infty. \quad \blacksquare$$

Theorem 4.2. Under Assumptions A1 and A2, let Q_θ be the law of the controlled process and P^0 is the law of the uncontrolled process $dz = \sigma dW$ with the same initial distribution. If Novikov's condition holds

$$\mathbb{E}_{P^0} \left[\exp \left(\frac{1}{2\sigma^2} \int_0^1 \|u_\theta(z(\tau), \tau, \mathcal{X}, s)\|^2 d\tau \right) \right] < \infty$$

then $Q_\theta \ll P^0$, and

$$\frac{dQ_\theta}{dP^0} = \exp \left(\frac{1}{\sigma} \int_0^1 u_\theta(z(\tau), \tau, \mathcal{X}, s)^\top dW_\tau - \frac{1}{2\sigma^2} \int_0^1 \|u_\theta(z(\tau), \tau, \mathcal{X}, s)\|^2 d\tau \right).$$

Proof. This is the standard Girsanov theorem. Since the controlled drift is progressively measurable and Novikov's condition holds, the exponential local martingale

$$\varepsilon_1 = \exp \left(\frac{1}{\sigma} \int_0^1 u_\theta^\top dW_\tau - \frac{1}{2\sigma^2} \int_0^1 \|u_\theta\|^2 d\tau \right)$$

is a true martingale with expectation 1. Define a new probability measure Q_θ as

$$\frac{dQ_\theta}{dP^0} = \varepsilon_1.$$

Under Q_θ , the process

$$\tilde{W}_\tau = W_\tau - \frac{1}{\sigma} \int_0^\tau u_\theta(z(r), r, \mathcal{X}, s) dr$$

is a Brownian motion. Therefore under Q_θ , we have

$$dz(\tau) = \sigma d\tilde{W}_\tau + u_\theta(z(\tau), \tau, \mathcal{X}, s) d\tau,$$

which is exactly the controlled diffusion. Thus Q_θ is the path law of the controlled process and is absolutely continuous with respect to P^0 . ■

4.4. Existence of the Schrödinger bridge optimizer

Theorem 4.3. Consider the path-space objective $J(Q) = \text{KL}(Q \parallel P^0) - \lambda \mathbb{E}_Q[\log p_\eta(y \mid z(1), \mathcal{X}, s)]$. Under Assumptions A1–A4 and the integrability condition $\sup_{Q \ll P^0} \mathbb{E}_Q[|\log p_\eta(y \mid z(1), \mathcal{X}, s)|] < \infty$, there is at least one minimizer Q^* of $J(Q)$ over all path laws $Q \ll P^0$.

Proof. Let us define the path space as

$$J(Q) = \text{KL}(Q \parallel P^0) + \Phi(Q),$$

where $\Phi(Q) = -\lambda \mathbb{E}_Q[\log p_\eta(y \mid z(1), \mathcal{X}, s)]$. The KL divergence is convex and lower semicontinuous in Q under weak convergence together with absolute continuity. The terminal log-likelihood term is continuous in the terminal marginal when $\log p_\eta(y \mid z(1), \mathcal{X}, s)$ is continuous in $z(1)$ and uniformly integrable; this follows from Assumption A3 and the bounded-growth condition. Take a minimizing sequence $\{Q_m\}$ such that

$$J(Q_m) \downarrow \inf_Q J(Q).$$

Since $\text{KL}(Q_m \parallel P^0)$ is bounded above along the minimizing sequence up to an additive constant, the sequence $\{Q_m\}$ is tight on path space by standard entropy-tightness arguments. By Prokhorov's theorem, there is a weakly convergent subsequence, still denoted Q_m , such that $Q_m \Rightarrow Q^*$.

By the lower semicontinuity of KL, $\text{KL}(Q^* \parallel P^0) \leq \liminf_{m \rightarrow \infty} \text{KL}(Q_m \parallel P^0)$. By continuity of the expected terminal log-likelihood under the assumed integrability, we have

$$\mathbb{E}_{Q_m} [\log p_\eta(y \mid z(1), \mathcal{X}, s)] \rightarrow \mathbb{E}_{Q^*} [\log p_\eta(y \mid z(1), \mathcal{X}, s)].$$

Therefore

$$\mathcal{J}(Q^*) \leq \liminf_{m \rightarrow \infty} \mathcal{J}(Q_m) = \inf_Q \mathcal{J}(Q).$$

Hence, Q^* is a minimizer. ■

4.5. Characterization of the optimal drift

Theorem 4.4. Let P_s^0 denote the path law of the reference diffusion $dZ_\tau = \sigma dW_\tau$, $Z_0 \sim \mu_0(\cdot \mid s)$ and consider the softly constrained Schrödinger bridge objective $\mathcal{J}(Q) = \mathbb{E}_Q[-\log p_\eta(y, \cdot, Z_1 X s)] + \lambda \text{KL}(Q \parallel P_s^0)$ over path laws satisfying $Q \ll P_s^0$ and $Q_0 = \mu_0$. Assume that $p_\eta(y, \cdot, z X s) > 0$ and that $0 < \mathcal{Z}_\eta(X, s, y) = \mathbb{E}_{P_s^0}[p_\eta(y, \cdot, Z_1 X s)^{1/\lambda}] < \infty$. Then $\mathcal{J}(Q)$ has the unique minimizer Q^* , whose Radon–Nikodym derivative with respect to the reference path law is $\frac{dQ^*}{dP_s^0} = \frac{p_\eta(y, \cdot, Z_1 X s)^{1/\lambda}}{\mathcal{Z}_\eta(X, s, y)}$.

The corresponding terminal marginal is $q_1^*(z, \cdot, X s y) = \frac{\rho_1^0(z \mid s) p_\eta(y, \cdot, z X s)^{1/\lambda}}{\int \rho_1^0(v \mid s) p_\eta(y, \cdot, v X s)^{1/\lambda} dv}$, where $\rho_1^0(\cdot \mid s)$ is the terminal density generated by the reference process. Furthermore, define the backward potential $\varphi_\tau(z, X, s, y) = \mathbb{E}_{P_s^0}[p_\eta(y, \cdot, Z_1 X s)^{1/\lambda} \mid Z_\tau = z]$. If φ_τ is positive and differentiable with respect to z , the optimal controlled process satisfies $dZ_\tau = u^*(Z_\tau, \tau, X, s) d\tau + \sigma dW_\tau$, where the optimal drift is $u^*(z, \tau, X, s) = \sigma^2 \nabla_z \log \varphi_\tau(z, X, s, y)$.

Proof. Define the terminal likelihood functional

$$G(Z_1) = p_\eta(y, \cdot, Z_1 X s)^{1/\lambda}.$$

The proposed optimal path law is therefore

$$dQ^* = \frac{G(Z_1)}{\mathbb{E}_{P_s^0}[G(Z_1)]} dP_s^0.$$

For any admissible path law $Q \ll P_s^0$, we have

$$\begin{aligned}\mathrm{KL}(Q \parallel Q^*) &= \mathbb{E}_Q \left[\log \frac{dQ}{dQ^*} \right] \\ &= \mathrm{KL}(Q \parallel P_s^0) - \frac{1}{\lambda} \mathbb{E}_Q [\log p_\eta(y, l, Z_1 X s)] + \log Z_\eta(X, s, y).\end{aligned}$$

Multiplying by λ and rearranging gives

$$J(Q) = \lambda \mathrm{KL}(Q \parallel Q^*) - \lambda \log Z_\eta(X, s, y).$$

Because the Kullback–Leibler divergence is non-negative, we have

$$J(Q) \geq -\lambda \log Z_\eta(X, s, y),$$

with equality if and only if $Q = Q^*$. Hence, Q^* is the unique minimizer.

The terminal marginal follows by integrating the likelihood-tilted path law over all trajectories terminating at $Z_1 = z$. Finally, applying the Doob h -transform to the reference diffusion with the backward potential φ_τ yields the controlled drift

$$u^*(z, \tau, X, s) = \sigma^2 \nabla_z \log \varphi_\tau(z, X, s, y).$$

Therefore, the optimal NSBCM dynamics constitute a likelihood-tilted, entropy-regularized stochastic transport process.

Theorem 4.5. Under Assumptions A3, A5, and A6, if, for two parameter values η, η' , $p_\eta(y = j \mid z, \mathcal{X}, s) = p_{\eta'}(y = j \mid z, \mathcal{X}, s) \forall j$ for almost every $(z, \mathcal{X}, s)$, then $E_{\eta,j}(z, \mathcal{X}, s) - E_{\eta,m}(z, \mathcal{X}, s) = E_{\eta',j}(z, \mathcal{X}, s) - E_{\eta',m}(z, \mathcal{X}, s)$ for all j, m almost everywhere. Under the normalization in Assumption A5, this implies $\eta = \eta'$.

Proof. If the conditional softmax probabilities are equal, then for any pair j, m , we have

$$\frac{p_\eta(y = j \mid z, \mathcal{X}, s)}{p_\eta(y = m \mid z, \mathcal{X}, s)} = \frac{p_{\eta'}(y = j \mid z, \mathcal{X}, s)}{p_{\eta'}(y = m \mid z, \mathcal{X}, s)}.$$

Using the softmax form, we have

$$\exp\{-E_{\eta,j}(z, \mathcal{X}, s) + E_{\eta,m}(z, \mathcal{X}, s)\} = \exp\{-E_{\eta',j}(z, \mathcal{X}, s) + E_{\eta',m}(z, \mathcal{X}, s)\}.$$

Taking logarithms

$$E_{\eta,j}(z, \mathcal{X}, s) - E_{\eta,m}(z, \mathcal{X}, s) = E_{\eta',j}(z, \mathcal{X}, s) - E_{\eta',m}(z, \mathcal{X}, s).$$

Thus, the energies are identified up to a common additive term independent of j . Under the normalization $\sum_j E_j = 0$, the additive ambiguity vanishes, so $E_\eta = E_{\eta'}$ almost everywhere. By Assumption A6, $\eta = \eta'$. ■

Theorem 4.6. Under Assumptions A5–A7, if $p_{\theta,\eta}(y | \mathcal{X}, s) = p_{\theta',\eta'}(y | \mathcal{X}, s)$ for almost every $(y, \mathcal{X}, s)$, then $\theta = \theta'$ and $\eta = \eta'$.

Proof. Write the marginal likelihood as $p_{\theta,\eta}(y | \mathcal{X}, s) = \int p_{\eta}(y | z, \mathcal{X}, s) q_{\theta}(z | \mathcal{X}, s) dz$. The equality of marginal choice probabilities for all y implies the equality of the induced mixed softmax model. By Theorem 4.5, conditional choice kernels identify energy differences up to normalization. Hence, if the marginal laws coincide, the terminal mixing laws and the energy decoder must induce the same family of conditional logits almost everywhere. Under Assumption A7, the equality of the terminal laws implies $\theta = \theta'$. Then, with identical terminal laws fixed, the equality of the mixed conditional choice model implies the equality of the conditional decoder, and by Theorem 4.5 and normalization, $\eta = \eta'$. ■

Theorem 4.7. Suppose that Assumptions A1–A9 hold. Let the sample log-likelihood function be defined as $\hat{M}_N(\theta, \eta) = \frac{1}{N} \sum_{n=1}^N \sum_{t=1}^{T_n} \log p_{\theta,\eta}(y_{nt} | \mathcal{X}_{nt}, s_n)$, and let the maximum likelihood estimator be $(\hat{\theta}_N, \hat{\eta}_N) \in \arg \max_{(\theta, \eta) \in \Theta \times H} \hat{M}_N(\theta, \eta)$. Then prove that $(\hat{\theta}_N, \hat{\eta}_N) \xrightarrow{p} (\theta_0, \eta_0)$.

Proof. Define

$$m_{\theta,\eta}(X) = \log p_{\theta,\eta}(y | \mathcal{X}, s).$$

By the compactness of $\Theta \times H$ and the continuity of the drift, terminal law, and energy decoder, $m_{\theta,\eta}(X)$ is continuous in (θ, η) almost surely. By Assumptions A1–A3, the energies have at most quadratic growth in z , while Theorem 4.1 gives finite second moments for the latent diffusion. Hence, an integrable envelope $M(X)$ exists such that

$$\sup_{(\theta, \eta) \in \Theta \times H} |m_{\theta,\eta}(X)| \leq M(X), \mathbb{E}[M(X)] < \infty.$$

By Assumption A8, the function class is Glivenko–Cantelli. Therefore

$$\sup_{(\theta, \eta) \in \Theta \times H} |\hat{M}_N(\theta, \eta) - M(\theta, \eta)| \xrightarrow{p} 0,$$

where

$$M(\theta, \eta) = \mathbb{E}_0[\log p_{\theta,\eta}(Y | \mathcal{X}, S)].$$

By Assumption A9, $M(\theta, \eta)$ has a unique maximizer (θ_0, η_0) . Therefore, by the argmax theorem, we have

$$(\hat{\theta}_N, \hat{\eta}_N) \xrightarrow{p} (\theta_0, \eta_0).$$

Theorem 4.8. Consider the penalized criterion $\hat{M}_N^\lambda(\theta, \eta) = \hat{M}_N(\theta, \eta) - \lambda_N R(\theta, \eta)$, where R is continuous and bounded on $\Theta \times H$, and $\lambda_N \rightarrow 0$. Any maximizer $(\hat{\theta}_N^\lambda, \hat{\eta}_N^\lambda)$ satisfies

$$(\hat{\theta}_N^\lambda, \hat{\eta}_N^\lambda) \xrightarrow{p} (\theta_0, \eta_0).$$

Proof. Since R is bounded and $\lambda_N \rightarrow 0$

$$\sup_{(\theta, \eta)} |\lambda_N R(\theta, \eta)| \rightarrow 0.$$

Hence

$$\sup_{(\theta, \eta)} |\hat{M}_N^\lambda(\theta, \eta) - M(\theta, \eta)| \leq \sup_{(\theta, \eta)} |\hat{M}_N(\theta, \eta) - M(\theta, \eta)| + \sup_{(\theta, \eta)} |\lambda_N R(\theta, \eta)| \xrightarrow{p} 0.$$

The rest follows exactly as in Theorem 4.7 by the argmax theorem. ■

Proposition 4.1. If the drift is identically zero, the initial law is degenerate at z_0 , and the energy is linear in the observed attributes, $E_j(z, \mathcal{X}, s) = -x_j^\top \beta(s)$, then the model reduces to a standard multinomial logit.

Proof. If $u_\theta \equiv 0$ and $z(0)$ is degenerate, then $z(1)$ is either deterministic or irrelevant if it does not enter the energy beyond $\beta(s)$. Then

$$p(y = j | \mathcal{X}, s) = \frac{\exp(x_j^\top \beta(s))}{\sum_m \exp(x_m^\top \beta(s))},$$

which is exactly MNL. ■

Proposition 4.2. If $z(1)$ is random but static across repeated choices and the energy takes the form $E_j(z, \mathcal{X}, s) = -x_j^\top \beta(z, s)$, then the model reduces to ML with the mixing law induced by the terminal latent state.

Proof. Conditional on z , the choice probability is MNL

$$p(y = j | z, \mathcal{X}, s) = \frac{\exp(x_j^\top \beta(z, s))}{\sum_m \exp(x_m^\top \beta(z, s))}.$$

Marginalizing over z yields a standard ML integral. ■

We show that there is a unique latent process and a path-space optimizer that guarantees a valid solution to the bridging problem. The optimal bridge drift is defined and provides insight into the nature of the hidden trajectories. We also prove the identifiability of the energy decoder and the full model, thereby enabling the interpretability and recoverability of the underlying parameters. Furthermore, the consistency and reliability of the estimator are demonstrated in large samples. Finally, the framework is sufficiently general to include existing models, reducing to the MNL and ML models as special cases, thereby linking the method to established benchmarks in the field.

Algorithm 1. Training procedure of the proposed NSBCM

Require: Behavioral dataset $\mathcal{D} = \{(X_i, y_i)\}_{i=1}^N$, choice set $\mathcal{C}_i$, latent dimension K , diffusion coefficient σ , entropy regularization weight λ , learning rate η , batch size B , and number of epochs E .

Ensure: Trained drift network b_θ , energy decoder E_ϕ , and behavioral choice estimator $P_\Theta(y_i | X_i)$.

Initialize neural drift parameters θ , decoder parameters ϕ , and latent behavioral prior distribution $p_0(z | X_i)$.

for $e = 1$ to E do

$$dZ_t \leftarrow b_\theta(Z_t, t, X_i) dt + \sigma dW_t$$

$$Z_{t+\Delta t} \leftarrow Z_t + b_\theta(Z_t, t, X_i)\Delta t + \sigma\sqrt{\Delta t}\epsilon_t$$

for each alternative $k = 1, \dots, K$ do

$$E_\phi(Z_T, k, X_i)$$

end for

$$P(y_i = k | Z_T, X_i) \leftarrow \frac{\exp[-E_\phi(Z_T, k, X_i)]}{\sum_{j=1}^K \exp[-E_\phi(Z_T, j, X_i)]}$$

$$\mathcal{L}_{\text{NLL}} \leftarrow - \sum_{i=1}^B \log P(y_i | Z_T, X_i)$$

$$\mathcal{L}_{\text{SB}} \leftarrow \frac{1}{2\sigma^2} \int_0^T \|b_\theta(Z_t, t, X_i)\|^2 dt$$

$$\mathcal{L} \leftarrow \mathcal{L}_{\text{NLL}} + \lambda \mathcal{L}_{\text{SB}}$$

if $\mathcal{L}$ decreases

$$(\theta, \phi) \leftarrow (\theta, \phi) - \eta \nabla_{\theta, \phi} \mathcal{L}$$

else

Retain previous parameter values.

end if

if convergence criterion is satisfied

Break;

else

Continue

end if

end for

Return: Trained drift network b_θ , energy decoder E_ϕ , and behavioral choice estimator $P_\Theta(y_i | X_i)$.

5. Experimental evaluation and computational analysis

We chose the student mode choice dataset as a proof-of-concept dataset because of its availability and its socioeconomic, environmental, and travel-related factors, along with five distinct travel

alternatives. Its individual features and alternative attributes offer an interesting context for studying nonlinear preference heterogeneity, stochastic contextual variation, and probabilistic calibration. However, it is small and has few students, so we cannot draw many conclusions from it. To test its external validity, we used the London passenger mode choice (LPMC) data and the Australian intercity mode choice data. The London dataset provides a relatively large sample of revealed preferences for four urban travel options, whereas the Australian dataset provides an intercity context with air, train, bus, and car alternatives. These two datasets challenge transferability because of differences in the sample size, geographic context, populations, and other structural variations. They do not, however, provide empirical validity in other areas, such as transportation. Thus, the empirical findings are presented in terms of travel mode choice, and cross-domain application is suggested as a theoretically informed research line that needs validation.

Algorithm 2. Entropy-regularized Schrödinger bridge latent transport learning

Require: Prior latent behavioral distribution $p_0(z)$, posterior behavioral target distribution $p_1(z | X, y)$, reference diffusion process Q , entropy regularization weight λ , diffusion coefficient σ , and time grid $\{t_m\}_{m=0}^M$.

Ensure: Learned stochastic transport trajectory $\{Z_t\}_{t=0}^T$ and optimal bridge drift b_θ^* .

Initialize the latent prior samples: $Z_0 \sim p_0(z | X)$

$dZ_t \leftarrow \sigma dW_t$

$dZ_t \leftarrow b_\theta(Z_t, t, X) dt + \sigma dW_t$

for each time step t_m do

$b_\theta(Z_{t_m}, t_m, X)$

$Z_{t_{m+1}} \leftarrow Z_{t_m} + b_\theta(Z_{t_m}, t_m, X)\Delta t + \sigma\sqrt{\Delta t} \epsilon_m$

If stochastic instability occurs then

Apply entropy regularization correction.

Else

Continue

end if

end for

$\hat{p}_1(z)$ \Estimate the empirical terminal distribution

$$D_{\text{KL}}(P_\theta \parallel Q) = \frac{1}{2\sigma^2} \mathbb{E} \left[\int_0^T \|b_\theta(Z_t, t, X)\|^2 dt \right]$$

$$J(\theta) \leftarrow -\mathbb{E}[\log P(y | Z_T, X)] + \lambda D_{\text{KL}}(P_\theta \parallel Q)$$

If transport discrepancy decreases then

Update neural drift parameters b_θ .

Else

Adjust entropy regularization strength λ .

End if

$$b_\theta^*(Z_t, t, X) \leftarrow \sigma^2 \nabla_z \log \phi_t(Z_t, X)$$

Return: The learned latent transport trajectory and optimal stochastic drift field.


The proposed NSBCM framework is fully optimized, and Algorithms 1 and 2 present the latent transport learning procedure. Algorithm 1 is the training mechanism comprising stochastic latent drift learning, entropy-regularized optimization, and energy-based probabilistic decoding, while Algorithm 2 describes the Schrödinger bridge's latent transport process that yields a continuous-time evolution of preferences under uncertainty-aware stochastic dynamics. We evaluate the method on a benchmark (the student travel mode choice dataset) that includes heterogeneous transportation choices across modes such as private vehicles, buses, railways, bikes, and walking-based mobility services. The dataset includes sociodemographic, contextual, temporal, and travel-related variables to model latent behavioral uncertainty and changes in preference evolution, including travel cost, travel time, congestion intensity, weather conditions, income level, and trip distance. We replaced all missing observations, one-hot encoded categorical variables, and normalized numerical variables using min–max scaling before training. We split the dataset into training, testing, and validation sets at a 70:15:15 ratio. We built the proposed NSBCM structure in PyTorch and applied it to latent trajectory propagation using GPU-based optimizers, using SDE solvers and Euler–Maruyama discretization. To evaluate the proposed framework fairly and consistently, we compared it with other models, including the MNL, ML, neural choice networks (NCN), variational diffusion choice models (VDCM), and energy-based preference networks (EDPN), under the same experimental conditions.

We chose the hyperparameter settings in Table 1 after numerous validation experiments that optimized latent-transport smoothness, stochastic exploration, and predictive stability. Adding latent dimensionality boosted the model's ability to represent behavior, enabling it to capture more complex preference transitions, but the required latent space was small enough to avoid unstable stochastic trajectories while maintaining interpretability. Similarly, medium diffusion coefficients led to stable preference exploration and low entropy dissipation during transport optimization [33032]. The behavior distribution prior remained smooth and interpolated between the behavior's prior and posterior distribution, without forcing the stochastic transport process to interpret into preference trajectories. We calculated several statistical performance metrics, including classification accuracy, F1-score, negative log-likelihood (NLL), and Brier calibration score, to assess the predictive performance. These metrics capture predictive consistency, probabilistic reliability, and behavioral calibration quality.

Figure 2 illustrates how the proposed NSBCM framework learns latent preferences. The first behavioral prior is continuously and gradually transformed into the context-conditioned terminal preference distribution via continuously timed, entropy-regularized Schrödinger bridge dynamics, yielding smooth stochastic trajectories rather than abrupt changes in latent utility space. The learned neural control field provides a mechanism for the directional control of behavioral states, while the temporal deformation of probability mass shows how behavioral preferences are gradually redistributed over the normalized bridge time. Finally, the geometric structure of the terminal energy-margin representations shows that the transported latent representation produces structured regions of decision confidence, paving the way for a connection between stochastic preference evolution and energy-based probabilistic choice decoding. Together, these findings offer a geometric insight into the controlled diffusion, entropy-regularized transport, latent behavioral heterogeneity, and uncertainty-aware decision formation underpinned by a unified, continuous-time choice modeling framework, the NSBCM.

Table 1. Hyperparameter configuration of the proposed NSBCM framework.

Parameter	Value
Latent utility dimension	64
Drift network depth	5 Layers
Hidden units per layer	256
Diffusion coefficient ((σ))	0.15
Time discretization steps	100
Batch size	128
Initial learning rate	0.0003
Optimizer	AdamW
Training epochs	250
Entropy regularization ((λ))	0.45
Activation function	GELU

**Figure 2.** Entropy-regularized latent preference transport, neural drift dynamics, and terminal decision geometry of the proposed NSBCM framework.

As indicated in Table 2, the NSBCM had the best predictive and probabilistic performance for all three datasets. It obtained accuracy values of $95.84 \pm 0.19\%$, $83.47 \pm 0.21\%$, and $86.93 \pm 0.23\%$ on the student, LPMC, and Australian datasets, respectively, outperforming the strongest baseline, EDPN, by 4.58, 4.16, and 4.25 percentage points. Similar results were seen for the macro-F1, NLL, Brier scores, and Expected Calibration Error (ECE), with the NSBCM achieving the highest score on all datasets, indicating better discrimination, probabilistic accuracy, and calibration. It outperformed the neural SDE by 5.82, 6.05, and 6.02 percentage points, due not only to stochastic continuous-time dynamics but also to likelihood-conditioned Schrödinger transport and path-space entropy regularization. This small standard deviation also shows consistent performance across multiple trials. These results support the NSBCM's generalizability across travel choice contexts and dataset scales, with the caveat that applications beyond transportation still need validation.

As shown in Figure 3, the behavioral contribution analysis reveals that the most influential contextual variables (i.e., demographic and travel-related) in the proposed NSBCM framework are related to the evolution of latent preferences. The learned latent representations consistently show large mean probability contributions for variables related to transportation alternatives, as well as age, income level, transport cost, and transport duration, suggesting that they effectively capture different behavioral sensitivities across contexts. In addition, the proposed entropy-regularized transport mechanism does not appear to produce unstable or randomly distributed patterns of preference importance across variables (as evidenced by a gradual decrease in feature influence for lower-ranked variables). This observation verified that the framework can learn interpretable latent utility structures that represent realistic stochastic behavioral decision processes. The distribution of control cost over predicted behavioral alternatives also illustrates the proposed Schrödinger transport formulation's ability to adapt to changes in latent preference states while maintaining relatively stable stochastic transport energy across behavioral alternatives. The energy density distributions of the violin (Figure 3b) are smooth, suggesting that the learned latent transport process is not energy-chaotic, as would be the case if it resulted in sudden changes in energy during the evolution of the behavior, which is consistent with the theoretical formulation of stochastic transport in the presence of entropy minimization introduced in the methodology section.

Table 2. Comparative predictive performance analysis.

Dataset	Model	Accuracy (%)	Macro-F1	NLL	Brier	ECE
Student	MNL	81.42 ± 0.48	0.802 ± 0.006	0.582 ± 0.012	0.191 ± 0.005	0.072 ± 0.006
	ML	84.37 ± 0.41	0.831 ± 0.005	0.517 ± 0.010	0.176 ± 0.004	0.064 ± 0.005
	NCN	88.54 ± 0.35	0.876 ± 0.004	0.411 ± 0.009	0.142 ± 0.004	0.051 ± 0.004
	Neural ODE	89.31 ± 0.32	0.884 ± 0.004	0.396 ± 0.008	0.137 ± 0.003	0.047 ± 0.004
	Neural SDE	90.02 ± 0.30	0.891 ± 0.003	0.378 ± 0.008	0.132 ± 0.003	0.043 ± 0.003
	VDCM	90.13 ± 0.31	0.892 ± 0.004	0.372 ± 0.008	0.131 ± 0.003	0.041 ± 0.003
	EDPN	91.26 ± 0.28	0.904 ± 0.003	0.351 ± 0.007	0.122 ± 0.003	0.037 ± 0.003
	NSBCM	95.84 ± 0.19	0.947 ± 0.002	0.214 ± 0.005	0.081 ± 0.002	0.019 ± 0.002
LPMC	MNL	68.74 ± 0.36	0.584 ± 0.006	0.821 ± 0.014	0.284 ± 0.006	0.091 ± 0.006
	ML	70.91 ± 0.34	0.611 ± 0.005	0.756 ± 0.012	0.267 ± 0.005	0.082 ± 0.005

Dataset	Model	Accuracy (%)	Macro-F1	NLL	Brier	ECE
	NCN	74.86 ± 0.29	0.672 ± 0.004	0.654 ± 0.010	0.231 ± 0.004	0.068 ± 0.004
	Neural ODE	76.18 ± 0.28	0.691 ± 0.004	0.621 ± 0.009	0.219 ± 0.004	0.061 ± 0.004
	Neural SDE	77.42 ± 0.27	0.709 ± 0.004	0.589 ± 0.009	0.208 ± 0.004	0.055 ± 0.003
	VDCM	78.06 ± 0.26	0.718 ± 0.003	0.571 ± 0.008	0.201 ± 0.003	0.051 ± 0.003
	EDPN	79.31 ± 0.25	0.735 ± 0.003	0.548 ± 0.008	0.193 ± 0.003	0.046 ± 0.003
	NSBCM	83.47 ± 0.21	0.789 ± 0.003	0.421 ± 0.006	0.151 ± 0.003	0.028 ± 0.002
Australian	MNL	71.86 ± 0.42	0.682 ± 0.007	0.704 ± 0.013	0.249 ± 0.005	0.078 ± 0.006
	ML	74.23 ± 0.39	0.711 ± 0.006	0.648 ± 0.011	0.231 ± 0.005	0.069 ± 0.005
	NCN	78.46 ± 0.35	0.758 ± 0.005	0.561 ± 0.010	0.203 ± 0.004	0.057 ± 0.004
	Neural ODE	79.72 ± 0.33	0.773 ± 0.005	0.533 ± 0.009	0.194 ± 0.004	0.051 ± 0.004
	Neural SDE	80.91 ± 0.31	0.786 ± 0.004	0.507 ± 0.009	0.185 ± 0.004	0.046 ± 0.003
	VDCM	81.47 ± 0.30	0.793 ± 0.004	0.494 ± 0.008	0.180 ± 0.003	0.043 ± 0.003
	EDPN	82.68 ± 0.28	0.806 ± 0.004	0.471 ± 0.008	0.173 ± 0.003	0.038 ± 0.003
	NSBCM	86.93 ± 0.23	0.854 ± 0.003	0.349 ± 0.006	0.129 ± 0.003	0.023 ± 0.002




Figure 3. Explainability, stochastic behavioral uncertainty, probabilistic confidence analysis, and latent behavioral contribution evaluation of the proposed NSBCM framework: (a) Behavioral feature importance analysis, (b) distribution of entropy-regularized control, (c) comparative density distributions of prediction confidence, and (d) top latent behavioral drivers influencing stochastic preference evolution in the proposed NSBCM framework.

The analysis of density prediction confidence in Figure 3 also shows that the proposed NSBCM framework yields significantly more concentrated and sharper probability density functions than other baseline models. The high-confidence regions have a higher overall density, indicating greater confidence and stronger consistency in decision calibration during prediction. By comparison, uncertainty calibration and the stability of latent representation learning in traditional machine learning and econometric models are less widespread and less stable, implying less confidence. The latent behavioral driver analysis also corroborates that sociodemographic characteristics, travel preferences, and contextual mobility conditions strongly influence the stochastic evolution of preferences in the learned latent utility space. The strong correlation between feature contribution patterns and confidence calibration behavior suggests that the proposed NSBCM framework can effectively integrate interpretable behavioral representation learning with entropy-controlled probabilistic decision modeling. Taken together, these findings confirm the proposed framework's strong performance in both stable, uncertainty-aware behavioral prediction and interpretable latent dynamics of preferences in complex real-world decision scenarios.

According to the ablation results in Table 3, the Schrödinger bridge transport mechanism is the most important for stabilizing latent preference evolution and increasing predictive consistency. Without the transport formulation, the model adopts the characteristics of the traditional latent utility model, limiting its ability to capture the path dependency of behavioral change. Similarly, without entropy regularization, the stochastic trajectories become unstable, and the diffusion behavior becomes irregular when optimizing latent spaces. To weaken probabilistic calibration, we also remove the energy-based decoder. The results support the idea that the proposed NSBCM framework is not merely based on stochastic transport with a single neural architecture component but on the combined interaction among entropy-controlled stochastic transport, controlled diffusion dynamics, and energy-based probabilistic decoding. Beyond predictive performance analysis, we studied latent transport geometry and stochastic trajectory stability using the Wasserstein transport distance, drift variance, and entropy-evolution statistics. These measures quantify the smoothness and efficiency of probabilistic interpolation between behavioral distributions (see Figure 4).

Table 3. Ablation analysis of the proposed NSBCM components.

Configuration	Accuracy (%)	NLL
Full NSBCM	95.84	0.214
Without Schrödinger transport	91.47	0.356
Without entropy regularization	92.13	0.341
Without the energy decoder	90.82	0.389
Deterministic drift only	88.54	0.427




Figure 4. Comparative probabilistic performance and calibration across the student, LPMC, and Australian datasets.

The results showed that the NSBCM had the lowest ECE, Brier score, and NLL across all three datasets, reflecting improved classification performance, more reliable confidence estimates, and more accurate predictive probability distributions. Compared with the between-run standard deviations, the ECE values for the student and LPMC datasets were 0.019 and 0.028, respectively, and 0.023 for the Australian dataset. Perhaps the NSBCM does not achieve better accuracy with increased predictive confidence, as the Predictive Stability Threshold PST and calibration error remain significantly decreased in each. The proposed multiscale NSBCM framework can learn a latent preference-related transport, as illustrated in Figure 5. Local stability spectra, rotational flow plots, and topological regime maps show that the learned neural drift tends to structure the latent utility space into regular regions of stable transport rather than irregular regions of stochastic dynamics. In addition, the comparison of forward and reverse transport results, and the representation of the time evolution of preferences in the hysteresis diagram, suggest the coherent time evolution of preferences with reduced directional asymmetry, while the local irreversibility profile points to the intermediate transport region where behavioral adaptation is most severe. The co-evolution of Wasserstein discrepancy, stochastic control effort, latent entropy, covariance structure, and KL divergence shows that the bridge gradually evolves toward the terminal preference distribution while keeping stochastic exploration under control. The component contribution analysis also shows the importance of stochastic drift, the energy decoder, Schrödinger transport, and entropy regularization for predictive performance, while the stochastic uncertainty envelope reveals a gradual, bounded change in the variability of latent preferences over the normalized transport time.




Figure 5. Multiscale dynamic analysis of stability, transport irreversibility, stochastic control, component contribution, and uncertainty evolution in the proposed NSBCM framework.

The transport stability results reveal that the proposed NSBCM learns much more “smooth” latent preference trajectories than competing generative NSBCM architectures (see Table 4). Learning probabilistic paths has been shown to be geometrically efficient, reducing the Wasserstein transport distance between prior and posterior latent behavioral distributions. The same holds for the drift variance, which is reduced, showing that the proposed entropy-regularized stochastic control mechanism keeps latent utility evolution from becoming unstable and oscillatory. This behavior is especially relevant in behavioral decision systems, where people’s preferences may change incrementally rather than discontinuously over time. The proposed NSBCM captures this phenomenon by modeling latent preference formation as a controlled stochastic transport process governed by Schrödinger bridge dynamics. To capture uncertainty, we added Gaussian perturbations to travel-related variables, including travel time, congestion intensity, and environmental conditions. These experiments reflect the realistic, noisy behavioral environments commonly found in transportation systems.

Table 4. Latent transport stability and geometric analysis.

Model	Drift variance	Wasserstein distance	Stability index
VDCM	0.314	0.271	0.81
EDPN	0.281	0.233	0.84
NSBCM (proposed)	0.142	0.117	0.96

The robustness results in Table 5 demonstrate that the proposed NSBCM exhibits much greater predictive stability with respect to the level of stochastic perturbations than the other models. Classical econometric methods degrade rapidly because static utility assumptions cannot adapt to dynamic contexts. Deep generative models like VDCM show good robustness but remain sensitive to high noise in the latent diffusion. By contrast, the proposed NSBCM maintains strong predictive consistency, and the entropy-regularized bridge dynamics enable latent preference transport to evolve smoothly even in uncertain environments. This validates the proposed framework's viability for learning uncertainty-aware behavioral representations for complex real-world decision systems.

Table 5. Robustness analysis under stochastic contextual noise.

Noise level	MNL	ML	VDCM	NSBCM
0%	81.42	84.37	90.13	95.84
5%	78.21	82.46	88.31	94.72
10%	74.65	79.82	85.94	93.15
15%	71.38	76.91	82.47	91.08
20%	67.42	72.33	78.26	88.54

The proposed NSBCM framework consists of stochastic latent preference transport coupled with the formation of energy-based decisions, as shown in Figure 6. The predictive entropy manifold and stochastic control field describe how uncertainty is organized and how preference dynamics evolve in the latent utility space, and the terminal energy-margin geometry describes the separation between behavioral alternatives after stochastic transport. Regions near resulting decision boundaries have higher predictive entropy, meaning uncertainty increases where alternative-specific energies become less distinct. The entropy dispersion trajectory also shows the co-evolution of the distributional spread of latent preferences and uncertainty over normalized bridge time, linking continuous-time stochastic adaptation to terminal decision formation. Greater separation of the lowest competitive energy levels always correlates with higher (maximum) choice probability and lower predictive entropy. It also illustrates how the terminal energy-based decoder, the NSBCM maps onto a geometric interpretation of converting context-conditioned stochastic latent dynamics into structured, confidence-aware probabilistic decisions.

In Figure 7, the error distributions for the three tasks (choice probability estimation, latent utility prediction, and log-likelihood approximation) show that the proposed NSBCM consistently yields lower probabilistic estimation error and reduced variance compared with competing deep generative architectures and classical econometric models. High NSBCM error densities in lower-error regions suggest that the entropy-regularized Schrödinger transport formulation effectively stabilizes latent preference evolution and improves uncertainty-aware behavioral inference. By contrast, the baseline models (MNL, ML, NCN, and VDCM) have wider error distributions, driven by less consistent latent transport and less robust probabilistic calibration. Furthermore, the significant decrease in the means and standard deviations indicates that the proposed framework maintains geometric coherence while evolving preferences stochastically, yielding reliable and computationally stable behavioral decision representations in heterogeneous contextual environments.




Figure 6. Energy-based decision geometry and uncertainty evolution in the proposed NSBCM framework.




Figure 7. Error distribution analysis comparing the proposed NSBCM framework against classical econometric and deep generative choice models.

6. Conclusions

This study proposes a novel NSBCM for uncertainty-aware behavioral decision modeling in latent utility space. The framework differs markedly from conventional utility-based methods, modeling preference evolution as a continuous-time stochastic process subject to Schrödinger bridge dynamics, latent neural drift learning, and energy-based probabilistic decoding. By combining stochastic optimal transport, controlled diffusion processes, and entropy-aware probabilistic inference into a single framework, the model offers a mathematically sound representation of uncertainty-aware, nonstationary, and heterogeneous behavioral choices. Based on the theoretical analysis, we proved the existence of the controlled diffusion process, identified the energy decoder, and proved the consistency of the proposed estimator. Experimental results show that the NSBCM outperformed classical econometric and deep generative models in predictive accuracy, F1 score, calibration performance, latent transport stability, and robustness across all scenarios. The latent transport analysis also verified that the framework ensures geometric continuity in the evolution of preferences while retaining a stable, uncertain behavioral representation.

The empirical analysis confirms that the NSBCM can be used on the three assessments of travel mode choice data. It does not provide universal superiority across all domains of behavioral decision making. While the path-space formulation and energy decoder can be applied to different finite choice sets, healthcare decisions, product choice, financial decisions, and policy preferences may have different data-generating mechanisms, dependence structures, constraints, and/or ethical issues. Empirical claims in these domains thus need validation before generalization. Furthermore, the NSBCM imposes a higher computational burden than the MNL and single-pass neural choice models, as it involves repeated stochastic integration during training and inference. This trade-off is highlighted by the reported complexity, runtime, parameter count, and comparison of memory. The study suggests that future research could focus on low-rank drift parameterization, adaptive stochastic solvers, sparse relational conditioning, and testing the new models on nontransport choice data. The results show that entropy-regularized stochastic transport is a useful probabilistic framework for uncertainty-aware travel mode choice.

Author contributions

Irsa Sajjad: Conceptualization, formal analysis, methodology, software, validation, visualization, writing – original draft, writing – review and editing; Mashail M. Al Sobhi: Conceptualization, formal analysis, methodology, validation, writing – original draft, funding acquisition, writing – review and editing. All authors have read and approved the final version of the manuscript for publication.

Use of generative AI tools declaration

The authors declare that they have not used artificial intelligence (AI) tools in the creation of this article.

Data availability

The data that support the findings of this study are openly available at <https://www.kaggle.com/datasets/svenmillerkaggle/student-mode-choice-data>.

Acknowledgments

The authors extend their appreciation to Umm Al-Qura University, Saudia Arabia, for funding this research work through grant number 26UQU4281739GSSR01.

Funding details

This research was funded by Umm Al-Qura University, Saudi Arabia, under grant number 26UQU4281739GSSR01.

Conflict of interest

The authors declare no conflicts of interest.

References

1. K. E. Train, *Discrete choice methods with simulation*, 2 Eds., Cambridge: Cambridge University Press, 2012. <https://doi.org/10.1017/CBO9780511805271>
2. D. McFadden, K. Train, Mixed MNL models for discrete response, *J. Appl. Econ.*, **15** (2000), 447–470. [https://doi.org/10.1002/1099-1255\(200009/10\)15:5%3C447::AID-JAE570%3E3.0.CO;2-1](https://doi.org/10.1002/1099-1255(200009/10)15:5%3C447::AID-JAE570%3E3.0.CO;2-1)
3. T. Hillel, M. Bierlaire, M. Z. E. B. Elshafie, Y. Jin, A systematic review of machine learning classification methodologies for modelling passenger mode choice, *J. Choice Model.*, **38** (2021), 100221. <https://doi.org/10.1016/j.jocm.2020.100221>
4. Y. M. Aboutaleb, M. Danaf, Y. Xie, M. Ben-Akiva, Discrete choice analysis with machine learning capabilities, *arXiv*, 2021, arXiv:2101.10261. <https://doi.org/10.48550/arXiv.2101.10261>
5. L. Winkler, C. Ojeda, M. Oppel, A score-based approach for training Schrödinger bridges for data modelling, *Entropy*, **25** (2023), 316. <https://doi.org/10.3390/e25020316>
6. F. Vargas, P. Thodoroff, A. Lamacraft, N. Lawrence, Solving Schrödinger bridges via maximum likelihood, *Entropy*, **23** (2021), 1134. <https://doi.org/10.3390/e23091134>
7. S. Tang, Foundations of Schrödinger bridges for generative modeling, *arXiv*, 2026, arXiv:2603.18992. <https://doi.org/10.48550/arXiv.2603.18992>
8. Y. Chen, T. T. Georgiou, M. Pavon, Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge, *SIAM Rev.*, **63** (2021), 249–313. <https://doi.org/10.1137/20M1339982>
9. I. Sajjad, M. Malik, A hierarchical Bayesian discrete–continuous mixture model with nested Dirichlet processes, *Statistics*, 2026, 1–34. <https://doi.org/10.1080/02331888.2026.2654459>
10. I. Sajjad, I. A. Nafisah, M. M. A. Almazah, O. A. Alamri, J. G. Dar, A symmetrical analysis of decision making: Introducing the Gaussian negative binomial mixture with a latent class choice model, *Symmetry*, **16** (2024), 908. <https://doi.org/10.3390/sym16070908>
11. I. Sajjad, M. A. Alomair, J. G. Dar, A statistical machine learning approach to multimodal data processing: Hierarchical generalized distribution matching networks (HGDM-Net), *Phys. Scr.*, **101** (2026), 186003. <https://doi.org/10.1088/1402-4896/ae6402>

12. I. A. Nafisah, I. Sajjad, M. A. Alshahrani, O. A. Alamri, M. M. A. Almazah, J. G. Dar, Statistical predictive hybrid choice modeling: Exploring embedded neural architecture, *Mathematics*, **12** (2024), 3115. <https://doi.org/10.3390/math12193115>
13. M. Haghani, M. C. J. Bliemer, D. A. Hensher, The landscape of econometric discrete choice modelling research, *J. Choice Model.*, **40** (2021), 100303. <https://doi.org/10.1016/j.jocm.2021.100303>
14. A. Vij, J. L. Walker, How, when and why integrated choice and latent variable models are latently useful, *Transp. Res. B Methodol.*, **90** (2016), 192–217. <https://doi.org/10.1016/j.trb.2016.04.021>
15. G. S. Crawford, R. Griffith, A. Iaria, A survey of preference estimation with unobserved choice set heterogeneity, *J. Econometrics*, **222** (2021), 4–43. <https://doi.org/10.1016/j.jeconom.2020.07.024>
16. F. Zhang, A Monte Carlo study of integrating discrete choice models and neural networks in transportation decision-making, In: *Proceedings of the 2nd international conference on engineering management, information technology and intelligence (EMITI)*, **1** (2025), 522–526. <https://doi.org/10.5220/0014362200004718>
17. A. A. Pooladian, J. Niles-Weed, Plug-in estimation of Schrödinger bridges, *SIAM J. Math. Data Sci.*, **7** (2025), 1315–1336. <https://doi.org/10.1137/24M1687340>
18. J. Garg, X. Zhang, Q. Zhou, Soft-constrained Schrödinger bridge: A stochastic control approach, *Proc. Mach. Learn. Res.*, **238** (2024), 4429–4437.
19. S. De Marco, H. Pham, D. Zanni, Schrödinger bridges with jumps for time series generation, *arXiv*, 2026, arXiv:2602.20011. <https://doi.org/10.48550/arXiv.2602.20011>
20. R. Krueger, M. Bierlaire, P. Bansal, A data fusion approach for ride-sourcing demand estimation: A discrete choice model with sampling and endogeneity corrections, *Transp. Res. C Emerg. Technol.*, **152** (2023), 104180. <https://doi.org/10.1016/j.trc.2023.104180>
21. R. Singh Chauhan, C. Riis, S. Adhikari, S. Derrible, E. Zheleva, C. F. Choudhury, et al., Determining causality in travel mode choice, *Travel Behav. Soc.*, **36** (2024), 100789. <https://doi.org/10.1016/j.tbs.2024.100789>
22. M. Grzenda, M. Luckner, J. Zawieska, P. Wrona, Combining data from multiple sources for urban travel mode choice modelling, *arXiv*, 2024, arXiv:2407.12137. <https://doi.org/10.48550/arXiv.2407.12137>
23. J. Á. Martín-Baos, J. A. López-Gómez, L. Rodríguez-Benitez, T. Hillel, R. García-Ródenas, A prediction and behavioural analysis of machine learning methods for modelling travel mode choice, *Transp. Res. C Emerg. Technol.*, **156** (2023), 104318. <https://doi.org/10.1016/j.trc.2023.104318>
24. Q. Lai, M. Qin, X. W. Zhao, Reconfigurable multiscroll memristive neural network with application to telemedicine privacy protection, *IEEE Trans. Cybern.*, **56** (2026), 1910–1919. <https://doi.org/10.1109/TCYB.2026.3664460>
25. X. Huang, S. Xu, Mitigating choice overload: The interactive effects of set size and overall preference revealed by hierarchical drift diffusion modeling and electroencephalography, *NeuroImage*, **321** (2025), 121542. <https://doi.org/10.1016/j.neuroimage.2025.121542>
26. R. T. Q. Chen, Y. Rubanova, J. Bettencourt, D. Duvenaud, Neural ordinary differential equations, *Adv. Neural Inf. Process. Syst.*, **31** (2018), 6572–6583.

27. Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, B. Poole, Score-based generative modeling through stochastic differential equations, In: *International conference on learning representations (ICLR)*, 2021.
28. Z. Xu, Z. Lv, J. Li, Fast-TrafficNet: A hybrid model for efficient prediction of nonlinear traffic flow with sparse data, *Chaos Solitons Fractals*, **201** (2025), 117230. <https://doi.org/10.1016/j.chaos.2025.117230>
29. R. Ye, Z. Xu, Z. Lv, A spatio-temporal hypergraph fusion model for transportation revitalization index forecasting, *Comput. J.*, 2026, bxag088. <https://doi.org/10.1093/comjnl/bxag088>



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)