Perturbed utility Markovian traffic equilibrium: theory and computation

Rui Yao[1] and Kenan Zhang[2]

[1]Technion, Israel
[2]EPFL, Switzerland

Short summary

This paper proposes the perturbed utility Markovian equilibrium (PUME) framework for stochastic traffic assignment. The demand side, the perturbed utility Markovian choice model (PUMCM), characterizes route choice as an undiscounted infinite-horizon Markov decision process in which a convex surplus function at each state replaces the additive random utility shock; its gradient yields link choice probabilities on the full simplex, and corner solutions arise endogenously at dominated actions. Fenchel–Young duality links the state-wise surplus on the policy space to a flow-space perturbation, and strict stage negativity ensures every Bellman-consistent policy proper, so policy evaluation is well posed without the need for discounting. The framework accommodates recursive logit and network GEV as special cases, and extends to α-entmax and sparsemax, which assign zero probability to dominated actions. The equilibrium is formulated as a monotone variational inequality (VI) in cost space, which accommodates asymmetric supply and reduces to the classical Beckmann-type characterization when the supply is integrable. A modified policy iteration solves the inner network loading with a global convergence guarantee; and an accelerated VI algorithm solves the outer equilibrium problem. Numerical experiments illustrate endogenous corner solutions and approximately linear runtime scaling in network size.

Keywords: corner solutions; Markovian traffic equilibrium; modified policy iteration; perturbed utility; variational inequality.

1 Introduction

Traffic equilibrium models predict network flow patterns and support system design and management. Classical stochastic user equilibrium (SUE) couples a cos) mapping with a route choice model derived from additive random utility models (ARUM) (Daganzo & Sheffi, 1977; Sheffi, 1985; Ben-Akiva & Lerman, 1985). To avoid the challenges associated with choice set generation, Markovian traffic equilibrium (MTE) models routing as sequential link choice in an absorbing Markov decision process (MDP), yielding closed-form choice probabilities (Fosgerau, Frejinger, & Karlstrom, 2013; Mai et al., 2015) and a dual characterization on link-cost space that scales with the network (Akamatsu, 1997; Baillon & Cominetti, 2008; Oyama et al., 2022).

Despite this progress, three issues limit the applicability of MTE to general equilibrium modeling and computation. First, existing Markovian route choice models are often ARUM-based and therefore have full support: every feasible outgoing link is chosen with strictly positive probability (McFadden, 1981; Fosgerau, Frejinger, & Karlstrom, 2013; Mai et al., 2015). Consequently, even dominated links carry positive flow, producing implausible detours unless links are excluded ex ante. In many applications, corner solutions (zero-probability actions) are preferable as they generate more plausible behavioral predictions (Fosgerau et al., 2022; Yao et al., 2024; Tan et al., 2024). Second, existing algorithms based on the dual MTE formulation repeatedly solve a fixed-point problem, the Bellman equation, in the network loading step, and Bellman function evaluation relies on dynamic programming. In the undiscounted case, the Bellman operator is nonexpansive but generally not contractive, so value iteration may fail to converge Mai & Frejinger (2022).

Third, classical MTE formulations cast the equilibrium as a convex program, which requires a symmetric structure of the cost functions Baillon & Cominetti (2008); Oyama et al. (2022). Many settings of interest in transportation, including multi-class and asymmetric link interactions, naturally fall outside the optimization framework. The variational inequality (VI) formulation provides a more flexible equilibrium characterization Dafermos (1980), yet it has not been adopted in MTE.

Motivated by these gaps, this paper develops a perturbed utility Markovian equilibrium (PUME) framework. Its demand side is built on a perturbed utility Markovian choice model (PUMCM) that characterizes stochastic route choice as an undiscounted MDP regularized by a convex perturbation function. Accordingly, the Bellman optimality operator is written as a surplus (choice-probability generating) function whose gradient yields link choice probabilities at each state, and we derive conditions under which the induced choice map is well defined on the full simplex and produces corner solutions. PUMCM generalizes the ARUM distributional assumption of classic MTE Hofbauer & Sandholm (2002); Fosgerau & McFadden (2012) and extends the static perturbed utility route choice model of Fosgerau et al. (2022); Yao et al. (2024) to a Markovian setting while retaining scalability. On top of PUMCM, we formulate PUME as a VI in cost space with a monotone supply mapping and establish equilibrium existence and uniqueness under mild conditions. A standard monotone-VI solver suffices for computation, with the PUMCM inner loop handled by modified policy iteration. Our contributions are summarized as follows.

1. *Perturbed utility Markovian choice model.* PUMCM generalizes ARUM/GEV-based Markovian choice models and admits both interior and corner choice probabilities. We characterize how the regularity of the surplus function yields corner solutions and propagates through value functions, policies, and induced flows.
2. *Globally convergent undiscounted network loading.* Using structural properties of PUMCM, we establish existence and uniqueness of the optimal value function and prove global convergence of modified policy iteration (MPI) for the entire model class.
3. *Variational Markovian equilibrium.* We extend the classic optimization formulation of MTE by casting PUME as a monotone VI in cost space that accommodates asymmetric link interactions, prove existence and uniqueness, and show that every PUME induces a variational Nash equilibrium in flow space.

2 Methodology

Perturbed utility Markovian choice model

We model a traveler's route choice toward a destination d as a undiscounted infinite-horizon MDP with a single absorbing terminal state d. Nonterminal states $s \in \mathcal{S}$ represent network nodes and actions $\mathcal{A}_s$ represent feasible outgoing links; one-step rewards $u(s,a)$ are typically the negatives of link costs. Let Δ_s be the probability simplex on $\mathcal{A}_s$, and for any value function $V \in \mathbb{R}^{|\mathcal{S}|}$ define the state-action value $Q_s(V) = u(s,\cdot) + P(\cdot|s,\cdot)V$, where $P(\cdot|s,\cdot)$ collects nonterminal transition probabilities.

At each state s, a *surplus* (choice-probability generating) function $H_s : \mathbb{R}^{|\mathcal{A}_s|} \to \mathbb{R}$ maps Q-values to the maximum expected perturbed utility; its gradient returns the conditional choice probabilities. We impose the following conditions, which extend the classical CPGF conditions of Fosgerau, McFadden, & Bierlaire (2013) to the Markovian setting so that the induced choice map is well defined on the full simplex and can produce corner solutions.

Assumption 1 (Surplus function)**.** For every $s \in \mathcal{S}$:

(A1) *Convexity and smoothness.* H_s is convex and continuously differentiable (C^1) on $\mathbb{R}^{|\mathcal{A}_s|}$.

(A2) *Simplex gradient.* $\nabla H_s(Q) \in \Delta_s$ for every $Q \in \mathbb{R}^{|\mathcal{A}_s|}$.

(A3) *Translation equivariance.* $H_s(Q + \alpha\mathbf{1}) = H_s(Q) + \alpha$ for all $Q \in \mathbb{R}^{|\mathcal{A}_s|}, \alpha \in \mathbb{R}$.

(A4) *Optional smoothness.* H_s is twice continuously differentiable (C^2).

Condition (A2) is the key modeling difference from ARUM-based Markovian models: the gradient lies in Δ_s but need not be interior, so $\nabla H_s(Q)$ may assign exactly zero probability to dominated actions. The log-sum-exp surplus $H_s(Q) = \mu_s \log \sum_a \exp(Q_a/\mu_s)$ satisfies (A1)–(A4) and yields the recursive logit model (Fosgerau, Frejinger, & Karlstrom, 2013; Mai et al., 2015). The sparsemax surplus $H_s(Q) = \max_{\pi \in \Delta_s} \{\pi^\top Q - (\mu_s/2)\|\pi\|_2^2\}$ satisfies (A1)–(A3) but not (A4) and produces corner solutions via Euclidean projection onto Δ_s Martins & Astudillo (2016); the α-entmax family interpolates between logit and sparsemax Peters et al. (2019).

The surplus and its conjugate *perturbation* $H_s^*(\pi) \coloneqq \sup_Q \{\pi^\top Q - H_s(Q)\}$ on $\pi \in \Delta_s$ are linked by Fenchel–Young duality.

Lemma 2 (Fenchel–Young duality)**.** *Under Assumption* (A1)–(A3)*, for every $s \in \mathcal{S}$ and every $Q \in \mathbb{R}^{|\mathcal{A}_s|}$ the following hold.*

(a) *Weak duality. For every $\pi \in \Delta_s$,*

$$\pi^\top Q - H_s^*(\pi) \;\le\; H_s(Q).$$

(b) *Strong duality. Equality is attained at $\pi^* = \nabla H_s(Q) \in \Delta_s$; that is,*

$$\nabla H_s(Q)^\top Q \;-\; H_s^*\big(\nabla H_s(Q)\big) \;=\; H_s(Q).$$

(c) *Optimality. The variational representation*

$$H_s(Q) \;=\; \max_{\pi \in \Delta_s} \big\{\pi^\top Q - H_s^*(\pi)\big\}$$

holds, and the maximizer $\pi^* = \nabla H_s(Q)$ is unique.

Writing $U_\pi(s) \coloneqq \pi(\cdot|s)^\top u(s,\cdot) - H_s^*(\pi(\cdot|s))$ for the policy-induced stage reward and $P_\pi \in \mathbb{R}^{|\mathcal{S}|\times|\mathcal{S}|}$ with $[P_\pi]_{s,s'} = \sum_a \pi(a|s) P(s'|s,a)$ for the induced transition kernel, the policy Bellman operator is

$$T_\pi V = U_\pi + P_\pi V, \tag{1}$$

and the Bellman optimality operator reads

$$T_* V(s) = H_s(Q_s(V)). \tag{2}$$

Because the model is undiscounted, the policy value V_π is well defined only when the destination is reached almost surely. We formalize this definition

Definition 3 (Proper policy, Bertsekas (2012))**.** A policy π is *proper* if the terminal state is reached with probability one from every nonterminal state; equivalently, $\rho(P_\pi) < 1$, so $(\mathbb{I} - P_\pi)^{-1}$ exists and is elementwise nonnegative.

We next impose two conditions that jointly ensure well-posedness of the undiscounted Bellman operators.

Assumption 4 (Reachability and strict stage negativity)**.** The MDP and surplus function jointly satisfy:

(i) *Reachability.* There exists a proper policy π_0 with $H_s^*\big(\pi_0(\cdot|s)\big) < \infty$ for every $s \in \mathcal{S}$.

(ii) *Strict stage negativity.* $H_s\big(u(s,\cdot)\big) < 0$ for every $s \in \mathcal{S}$.

Strict stage negativity penalizes every additional period in the network and, together with reachability, replaces the role typically played by discounting. In routing with strictly positive link costs, it holds automatically for common perturbations; when it fails, translation equivariance (A3) provides a state-wise shift that enforces it without altering choice probabilities.

Theorem 5 (Automatic properness)**.** *Suppose Assumptions* (A1)–(A3) *and 4 hold. Then every policy π that satisfies the Bellman consistency equation*

$$V \;=\; T_\pi V \qquad \text{for some } V \in \mathbb{R}^{|\mathcal{S}|}$$

is proper, i.e. its transition kernel obeys $\rho(P_\pi) < 1$. In particular, at any fixed point V_* of the optimality operator T_*, the greedy policy $\pi_* = \nabla H(Q(V_*))$ is proper.

Automatic properness is the key structural property of PUMCM: every Bellman-consistent policy induces an absorbing chain and policy evaluation is well defined even under corner solutions, without relying on interiority Mai & Frejinger (2022) or topological restrictions Oyama et al. (2019).

Proposition 6 (Optimal value and induced demand)**.** *Suppose Assumptions* (A1)–(A3) *and 4 hold. Then:*

(i) *The Bellman optimality operator T_* admits a unique fixed point $V_* \in \mathbb{R}^{|\mathcal{S}|}$, and for every proper policy π,*

$$V_\pi \;\le\; V_* \;\le\; \mathbf{0}.$$

(ii) *The optimal policy $\pi_* = \nabla H\big(Q(V_*)\big)$ is proper and unique.*

(iii) *For any nonnegative load $q \in \mathbb{R}_+^{|\mathcal{S}|}$, the action demand map*

$$x^*(u) \;=\; \big[(\mathbb{I} - P_{\pi_*})^{-1}\pi_*\big]^\top q$$

is continuous and monotone nondecreasing in u. If in addition Assumption (A4) *holds, then x^* is continuously differentiable with positive semidefinite Jacobian.*

Perturbed utility Markovian equilibrium

We now couple PUMCM demand with a supply-side mapping through a cost-space equilibrium condition. Let $\mathcal{L}$ be a finite set of cost components and $\Omega \subseteq \mathbb{R}_+^L$ a closed convex cost domain. Demand segments $k \in \mathcal{K}$ (destinations or traveler classes) share cost vector c but have their own state-action structure, load q^k, and perturbation. Every state-action pair (s, a) in segment k maps to a cost component $\ell^k(s, a) \in \mathcal{L}$, encoded by the aggregation matrix $A^k \in \{0, 1\}^{L \times |\mathcal{A}^k|}$ with $A^k_{\ell,(s,a)} = \mathbf{1}[\ell^k(s, a) = \ell]$.

Assumption 7 (Utility–cost mapping)**.** For each segment $k \in \mathcal{K}$ and pair (s, a), the one-step utility is $u^k(s, a) = -c_{\ell^k(s,a)}$; strict stage negativity (Assumption 4(ii)) holds for every $c \in \Omega$.

The aggregate cost-space demand is

$$x(c) = \sum_{k \in \mathcal{K}} A^k x^{*,k}(u^k(c)) \in \mathbb{R}_+^L \tag{3}$$

.

Theorem 8 (Monotone demand in cost space)**.** *Suppose Assumptions 1*(A1)–(A3)*, 4, and 7 hold. Then:*

(i) *There exists a convex, continuously differentiable potential $\Phi : \Omega \to \mathbb{R}$ whose gradient returns the aggregate cost-space demand,*

$$\nabla\Phi(c) \;=\; -\,x(c), \qquad \forall\, c \in \Omega.$$

(ii) *Consequently the demand map x is continuous and monotone non-increasing on Ω:*

$$\langle x(c) - x(c'),\, c - c' \rangle \;\le\; 0, \qquad \forall\, c, c' \in \Omega.$$

(iii) *If in addition Assumption* (A4) *holds, then* Φ *is twice continuously differentiable,* x *is continuously differentiable, and its Jacobian is symmetric negative semidefinite:* $\nabla_c x(c) \preceq 0$ *with* $\nabla_c x(c) = \nabla_c x(c)^\top$.

Let $z : \Omega \to \mathbb{R}^L_+$ be a continuous supply function. Crucially, z may be nonseparable and its Jacobian need not be symmetric; this rules out a potential representation and covers asymmetric interactions (e.g., merge/diverge effects). Define the excess supply

$$E(c) \coloneqq z(c) - x(c). \tag{4}$$

Definition 9 (Perturbed utility Markovian equilibrium)**.** A cost vector $c^* \in \Omega$ is a *perturbed utility Markovian equilibrium* if

$$\langle E(c^*), c - c^* \rangle \geq 0, \qquad \forall\, c \in \Omega. \tag{5}$$

When $c^* \in \operatorname{int}(\Omega)$ this reduces to market clearance $z(c^*) = x(c^*)$.

The VI handles boundary constraints through Ω and requires only monotonicity of E, not integrability of z. Since x is monotone (Theorem 8), E is monotone whenever z is. We impose the following conditions on supply.

Assumption 10 (Coercive supply)**.** There exists $\hat{c} \in \Omega$ such that $\langle z(c), c - \hat{c} \rangle / \|c - \hat{c}\| \to +\infty$ as $\|c\| \to \infty$ with $c \in \Omega$.

Standard separable supply functions (e.g., inverse BPR $z_\ell(c_\ell) = \kappa_\ell c_\ell^{\beta_\ell}$, $\beta_\ell \geq 1$) are coercive. Monotonicity of x rules out cancellation and yields coercivity of E from coercivity of z alone, which ensures existence of a PUME.

Theorem 11 (Existence and uniqueness of PUME)**.** *Suppose Assumptions* (A1)–(A3)*, 4, 7, and 10 hold. Then:*

(i) *Existence. The variational inequality* (5) *admits at least one solution* $c^* \in \Omega$.

(ii) *Interiority and market clearance. If there exists* $\bar{c} \in \operatorname{int}(\Omega)$ *such that*

$$\langle E(c),\, c - \bar{c} \rangle \;>\; 0, \qquad \forall\, c \in \partial\Omega,$$

then every PUME lies in $\operatorname{int}(\Omega)$ *and satisfies the clearance condition* $z(c^*) = x(c^*)$.

(iii) *Uniqueness. If* z *is strictly monotone on* Ω*, the PUME is unique.*

The cost-space PUME admits a flow-space interpretation. Let $\mathcal{X}^k$ denote the nonnegative action flows of segment k satisfying flow conservation under the absorbing MDP, $\mathcal{X} = \{\sum_k A^k x^k : x^k \in \mathcal{X}^k\}$ the aggregate feasible set, and $R(x) = \Phi^*(-x)$ the flow-space aggregate perturbation obtained as the convex conjugate of the cost-space potential Φ from Theorem 8. Then PUMCM demand solves the perturbed best-response $x(c) \in \arg\min_{x \in \mathcal{X}} \{c^\top x + R(x)\}$. The pair (R, Φ) is thus a Fenchel conjugate pair linking a *primal* description on flow space and a *dual* description on cost space: the present paper develops the dual side, embedding PUMCM in a cost-space equilibrium formulation. In a companion paper, Fosgerau & Yao (2026) develop the primal side of the same conjugate pair, taking the generalized entropy R over link flows as the starting point and focusing on Markov decomposition of the flow problem and estimation using Fenchel–Young loss.

Definition 12 (Variational Nash equilibrium)**.** Assume the supply z is injective with continuous, monotone inverse $t = z^{-1}$ on $\mathcal{X}$. A flow $x^* \in \mathcal{X}$ is a *variational Nash equilibrium* (VNE) if there exists $\rho^* \in \partial R(x^*)$ with $\langle t(x^*) + \rho^*, x - x^* \rangle \geq 0$ for every $x \in \mathcal{X}$.

Theorem 13 (PUME induces a VNE)**.** *Assume the conditions of Theorem 11 hold with* $c^* \in \operatorname{int}(\Omega)$*, and that the supply* z *is injective with continuous, monotone inverse* $t = z^{-1}$ *on* $\mathcal{X}$*. Then* $x^* \coloneqq x(c^*) \in \mathcal{X}$ *is a variational Nash equilibrium in the sense of Definition 12: the vector* $\rho^* = -t(x^*)$ *belongs to* $\partial R(x^*)$*, and*

$$\langle t(x^*) + \rho^*,\, x - x^* \rangle \;\geq\; 0, \qquad \forall\, x \in \mathcal{X}.$$

Network loading via modified policy iteration

Computing a PUME is a nested problem. For a fixed cost c, the inner loop evaluates the demand $x(c)$ by solving the PUMCM loading problem $V_* = T_* V_*$; the outer loop updates c by any globally convergent monotone-VI method for (5). We focus here on the inner loop, which is intrinsic to PUMCM and is shared across every outer solver.

Given V_n, modified policy iteration (Puterman & Shin, 1978) alternates greedy policy improvement with $m \geq 1$ evaluation sweeps: $\pi_n = \nabla H(Q(V_n))$, $V_{n+1} = T_{\pi_n}^m V_n$. The depth m interpolates between value iteration ($m = 1$) and exact policy iteration ($m \to \infty$).

Algorithm 1 Modified policy iteration (MPI) for the PUMCM fixed point $V_* = T_* V_*$

Input: $V_0 = V_{\pi_0}$ with π_0 proper; depth $m \geq 1$; tolerance $\delta > 0$

for $n = 0, 1, 2, \ldots$ **do**
  $\pi_n(\cdot|s) \leftarrow \nabla H_s(Q_s(V_n)), \forall s$ ▷ greedy
  $W \leftarrow V_n$; **for** $h = 1, \ldots, m$ **do** $W \leftarrow U_{\pi_n} + P_{\pi_n} W$ **end for**
  $V_{n+1} \leftarrow W$; **if** $\|T_* V_{n+1} - V_{n+1}\|_\infty < \delta$ **then return** V_{n+1}
end for

Theorem 14 (Convergence of modified policy iteration)**.** *Suppose Assumptions* (A1)–(A3) *and 4 hold. Fix the evaluation depth $m \geq 1$ and initialize $V_0 = V_{\pi_0}$ for some proper policy π_0. Let (V_n) be the sequence generated by Algorithm 1.*

(i) *Global convergence. The sequence is monotone nondecreasing and bounded above by V_*:*

$$V_0 \;\leq\; V_1 \;\leq\; \cdots \;\leq\; V_n \;\leq\; V_* \;\leq\; \mathbf{0}, \qquad \forall\, n \geq 0,$$

and $V_n \to V_*$ componentwise as $n \to \infty$.

(ii) *Local linear rate. Let $G_* = P_{\pi_*}$ and $\mathcal{M}_m(V) := T_{\pi(V)}^m V$, where $\pi(V)$ denotes the greedy policy at V. For any matrix norm with $\|G_*^m\| < 1$ and any $\varepsilon \in \big(0,\, 1 - \|G_*^m\|\big)$, there exists a neighborhood $\mathcal{U}$ of V_* on which*

$$\|\mathcal{M}_m(V) - V_*\| \;\leq\; \big(\|G_*^m\| + \varepsilon\big)\, \|V - V_*\|, \qquad \forall\, V \in \mathcal{U}.$$

Consequently,

$$\limsup_{n \to \infty} \frac{\|V_{n+1} - V_*\|}{\|V_n - V_*\|} \;\leq\; \|G_*^m\|.$$

The outer loop is a monotone variational inequality (5) whose operator $E(c) = z(c) - x(c)$ is monotone under the conditions of Theorem 11, with PUMCM demand oracle. Any globally convergent method for monotone VIs applies; natural choices that require only evaluations of E and projections onto Ω include the Solodov–Tseng modified projection method Solodov & Tseng (1996) and the adaptive golden-ratio algorithm (aGRAAL) Malitsky (2020). Each outer iteration invokes MPI (Algorithm 1) once per demand segment to evaluate $x(c)$, so Theorem 14 governs the inner cost. A full analysis of safeguarded outer-loop acceleration is developed in a companion work; the numerical results below use its aGRAAL-M (Malitsky, 2020) with Anderson acceleration (Anderson, 1965).

3 Results and discussion

Corner solution in PUMCM

We isolate the central modeling consequence of PUMCM – admission of corner solutions – on a three-action simplex example, which captures in its simplest form the routing primitive that occurs at every state of the Markovian choice process. Consider an ordered state-action value vector

$$(Q_1, Q_2, Q_3) \;=\; (0, -\Delta, -2\Delta), \qquad \Delta \in [0, 2],$$

on which we evaluate the α-entmax choice map $\nabla H_\alpha^{\text{ent}}(Q)$ for $\alpha \in \{1, 1.2, 1.5, 2\}$, with $\alpha = 1$ recovering logit and $\alpha = 2$ recovering sparsemax Martins & Astudillo (2016); Peters et al. (2019). Figure 1 plots the three choice probabilities π_1, π_2, π_3 as Δ grows.

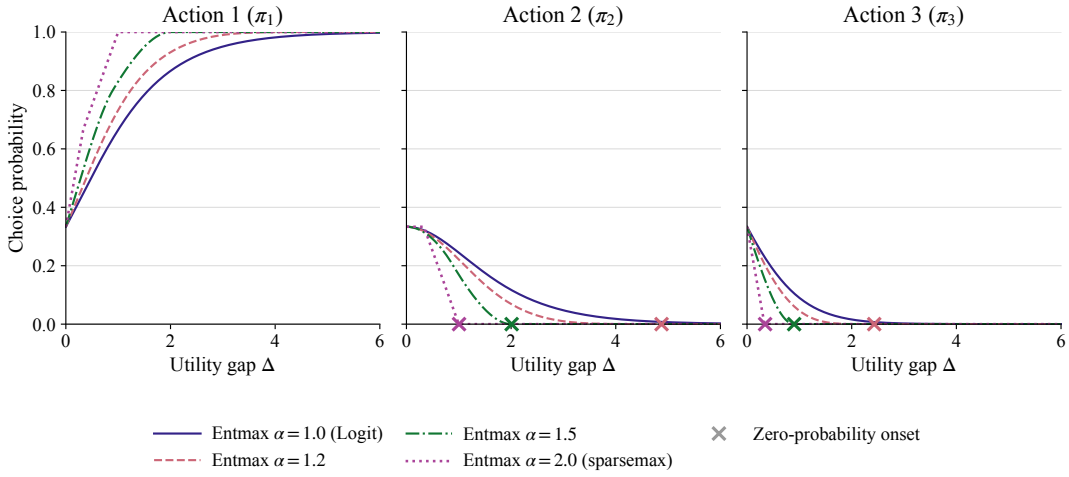


Figure 1: Choice probabilities on the simplex for the α-entmax family.

Three observations cleanly separate interior and corner behavior. First, logit ($\alpha = 1$) preserves full support for every Δ: even when Q_3 is severely dominated, it retains strictly positive probability, so in a routing context any detour action carries flow. Second, $\alpha = 2$ (sparsemax) eliminates Q_3 already at $\Delta \approx 0.345$ and concentrates the entire mass on Q_1 by $\Delta = 1$, reaching the boundary with visible kinks exactly at support-change thresholds where Assumption (A4) fails. Third, $\alpha \in (1, 2)$ interpolates: $\alpha = 1.5$ removes the dominated action at $\Delta \approx 0.9$ while remaining continuously differentiable, so corner solutions coexist with the smoothness required by Proposition 6 and Theorem 8. Embedded in PUMCM, this mechanism propagates to the network: dominated links and detour routes receive exactly zero flow without ex-ante exclusion, and automatic properness (Theorem 5) ensures that the resulting policies still reach the destination and that MPI remains globally convergent (Theorem 14).

Algorithm performance on benchmark networks

We next verify that the framework is computationally effective on standard benchmarks. Using aGRAAL-M and Anderson acceleration as the outer solver and MPI (depth $m = 10$) for network loading, we solve PUME on three networks of increasing size under both symmetric and asymmetric supply, with α-entmax ($\alpha = 1.5$) demand. Table 1 reports wall-clock runtimes to relative natural-residual tolerance 10^{-5} on an Apple M4 Pro laptop (48 GB).

Table 1: Runtime on benchmark networks (seconds).

Network	Links	ODs	Symmetric supply	Asymmetric supply
SiouxFalls	76	552	1.5	1.6
Anaheim	914	1,406	4.1	5.8
Chicago Sketch	2,950	149,382	101	125

All six instances converge under the same solver parameters: the corner-solution demand requires effort comparable to ARUM models, and asymmetric supply adds at most a modest overhead.

To assess scaling over problem sizes, we construct synthetic grid networks indexed by an integer $k \geq 1$: each grid has $(4k + 1) \times (4k + 1)$ nodes connected to their four cardinal neighbors by bidirectional BPR links with free-flow time $t_0 = 1$ and capacity $\kappa = 2{,}500$,

a fixed subset of nodes designated as OD zones, and zone-to-zone demand generated by a gravity model $q_{od} \propto \exp(-0.25 \cdot \mathrm{SPT}(o,d))$ with total per-origin demand $q = 10{,}000$. Sweeping $k \in \{2, 4, \ldots, 20\}$ spans roughly a 90-fold range in problem size (Figure 2). The outer iteration count remains bounded between 265 and 450, and total runtime grows approximately linearly with problem size on a log–log scale, indicating that per-iteration cost is the dominant driver.

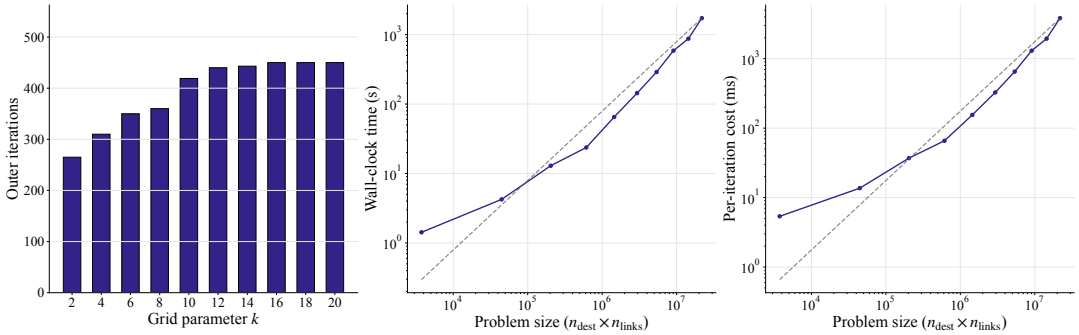


Figure 2: Scalability on synthetic grid networks ($k = 2, \ldots, 20$).

4 Conclusions

The perturbed utility Markovian equilibrium framework integrates two components: a choice model that admits both interior and corner solutions, and a cost-space variational inequality that accommodates both symmetric and asymmetric supply. Automatic properness makes the undiscounted Bellman problem well posed without discounting or topology-dependent assumptions, and propagates through the optimal value function, the induced demand map, and the global convergence of modified policy iteration. The stylized α-entmax example shows that corner solutions are a strictly non-ARUM feature: dominated actions can receive exactly zero probability at a finite utility gap. Benchmark experiments demonstrate that the resulting equilibrium problem is solvable at scale, with runtime growing approximately linearly in network size.

References

Akamatsu, T. (1997). Decomposition of path choice entropy in general transport networks. *Transportation Science*, *31*(4), 349–362.

Anderson, D. G. (1965). Iterative procedures for nonlinear integral equations. *Journal of the ACM*, *12*(4), 547–560.

Baillon, J.-B., & Cominetti, R. (2008). Markovian traffic equilibrium. *Mathematical Programming*, *111*(1), 33–56.

Ben-Akiva, M. E., & Lerman, S. R. (1985). *Discrete choice analysis: Theory and application to travel demand*. MIT Press.

Bertsekas, D. P. (2012). *Dynamic programming and optimal control* (4th ed., Vol. 2). Athena Scientific.

Dafermos, S. (1980). Traffic equilibrium and variational inequalities. *Transportation Science*, *14*(1), 42–54.

Daganzo, C. F., & Sheffi, Y. (1977). On stochastic models of traffic assignment. *Transportation Science*, *11*(3), 253–274.

Fosgerau, M., Frejinger, E., & Karlstrom, A. (2013). A link based network route choice model with unrestricted choice set. *Transportation Research Part B: Methodological*, *56*, 70–80.

Fosgerau, M., McFadden, D., & Bierlaire, M. (2013). Choice probability generating functions. *Journal of Choice Modelling*, *8*, 1-18. doi: https://doi.org/10.1016/j.jocm.2013.05.002

Fosgerau, M., & McFadden, D. L. (2012). A theory of the perturbed consumer with general budgets. *NBER Working Paper*, *17953*.

Fosgerau, M., Paulsen, M., & Rasmussen, T. K. (2022). A perturbed utility route choice model. *Transportation Research Part C: Emerging Technologies*, *136*, 103514.

Fosgerau, M., & Yao, R. (2026). A markov perturbed utility model of route choice.

Hofbauer, J., & Sandholm, W. H. (2002). On the global convergence of stochastic fictitious play. *Econometrica*, *70*(6), 2265–2294.

Mai, T., Fosgerau, M., & Frejinger, E. (2015). A nested recursive logit model for route choice analysis. *Transportation Research Part B: Methodological*, *75*, 100–112.

Mai, T., & Frejinger, E. (2022). Undiscounted recursive path choice models: Convergence properties and algorithms. *Transportation Science*, *56*(6), 1469–1482.

Malitsky, Y. (2020). Golden ratio algorithms for variational inequalities. *Mathematical Programming*, *184*, 383–410.

Martins, A. F. T., & Astudillo, R. F. (2016). From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proceedings of the 33rd international conference on machine learning (icml)* (pp. 1614–1623).

McFadden, D. (1981). Econometric models of probabilistic choice. In C. F. Manski & D. McFadden (Eds.), *Structural analysis of discrete data with econometric applications* (pp. 198–272). MIT Press.

Oyama, Y., Hara, Y., & Akamatsu, T. (2019). Prism-based path set restriction for solving Markovian traffic assignment problem. *Transportation Research Part B: Methodological*, *122*, 528–546.

Oyama, Y., Hara, Y., & Akamatsu, T. (2022). Markovian traffic equilibrium assignment based on network generalized extreme value model. *Transportation Research Part B: Methodological*, *155*, 135–159.

Peters, B., Niculae, V., & Martins, A. F. T. (2019). Sparse sequence-to-sequence models. In *Proceedings of the 57th annual meeting of the association for computational linguistics (acl)* (pp. 1504–1519).

Puterman, M. L., & Shin, M. C. (1978). Modified policy iteration algorithms for discounted Markov decision problems. *Management Science*, *24*(11), 1127–1137.

Sheffi, Y. (1985). *Urban transportation networks: Equilibrium analysis with mathematical programming methods.* Prentice-Hall.

Solodov, M. V., & Tseng, P. (1996). Modified projection-type methods for monotone variational inequalities. *SIAM Journal on Control and Optimization*, *34*(5), 1814–1830.

Tan, H., Xu, X., & Chen, A. (2024). On endogenously distinguishing inactive paths in stochastic user equilibrium: A convex programming approach with a truncated path choice model. *Transportation Research Part B: Methodological*, *183*, 102940.

Yao, R., Fosgerau, M., Paulsen, M., & Rasmussen, T. K. (2024). Perturbed utility stochastic traffic assignment. *Transportation Science*, *58*(4), 876–895.