Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

BAYESIAN RANKING AND SELECTION MODEL FOR THE SECOND-BEST NETWORK PRICING PROBLEM

Zhen Tan
H. Oliver Gao

Department of Civil and Environmental Engineering
Cornell University
Hollister Hall
Ithaca, NY 14853, USA

ABSTRACT

We adopt a Bayesian ranking and selection (R&S) model to solve the Second-best Network Pricing Problem (SNPP) in transportation. The objective of SNPP is to find an optimal subset of links and toll levels so as to minimize the total travel time on the network. It is an NP-hard problem with a large number of candidate solutions. We consider every combination of tollable link(s) and toll levels as an "alternative", and the problem's objective function value is regarded as a "reward", with uncertainties modeled by normal perturbations to the travel demand. We use a linear belief based Knowledge Gradient sampling policy to maximize the expected reward, with Monte Carlo sampling of the hyperparameters used to reduce the choice set size. Simulation experiments for a benchmark network show the effectiveness of the proposed method and its superior performance to a Sample Average Approximation based Genetic Algorithm.

1 INTRODUCTION

Network pricing has been widely recognized as an important countermeasure for traffic congestion (Yang and Huang1998). One well-known first-best pricing policy involves tolls set at marginal external costs on each link in the network and has been discussed in many studies (e.g., Sheffi 1985, Yang and Huang 1998). This policy has been regarded as merely a theoretical construct but impractical for real-world implementation. Under this first-best pricing scheme, total travel time on the network per modeling interval is minimized. Bergendorff et al. (1997) and Hearn and Yildirim (2002) solved the problem of finding the first-best pricing scheme that tolls the smallest number of links in a network. They showed that the first-best toll may not be unique. Their optimal "toll set", however, does not account for restrictions for the location of the tolled links (e.g., restricting tolled links within a certain cordon). Out of practical considerations, most of the recent studies have shifted to solving the second-best network pricing problem (SNPP), e.g., tolling only a subset of links that are tollable. There are generally two branches of research on SNPP: toll level design for a given set of links and optimizing toll rates and link selection simultaneously.

For solving the toll level design problem, derivative-based mathematical programing methods (e.g., Verhoef 2001; Lawphongpanich and Hearn 2004) and meta-heuristics such as genetic algorithms (GA) and simulated annealing (SA) (e.g., Yang and Zhang 2003; Shepherd and Sumalee 2004) have been proposed. Due to path selection assumptions and assignment convergence errors, the derivative-based approaches encounter deficiencies in finding global optimums (Shepherd and Sumalee 2004). It is also known that global optimum is not ensured in meta-heuristics. Ekström and Quttineh (2014) used surrogate optimization method for cordon-based SNPP and achieved close-to-optimal toll solutions with only tens of function evaluations. For joint optimization of toll rates and link selection, Verhoef (2001) proposed three iterative heuristic strategies based on a "link index" I_a, which represents the welfare gain from implementing a toll

on link *a* alone. This strategy fully accounts for interactions among tolled links but requires calculation of "location indices" for all possible combinations of candidate links. Shepherd et al. (2001) observed that such an index-based approach has two practical problems: the potential for negative toll predictions and the likelihood of poor initial predictions for parallel links. It was also found that the index-based approach could miss out on toll locations that can yield a high benefit if they are tolled simultaneously (Ekström et al. 2012). Shepherd and Sumalee (2004) linked "location indices" with GA to optimize toll locations and used GA to design toll levels given the toll links, due to the "location indices" used, GA often suggested solutions with more toll links that were in fact less optimal. Yang and Zhang (2003) used GA for the selection of toll locations and SA for optimal toll level design. Several subsequent studies also applied such heuristics for SNPP (e.g., Zuo et al. 2009). However, as is the case for the toll design problem, none of these methods guarantees global optimality (Ekström et al. 2012). Ekström et al. (2012) instead approximated discrete toll location decision variables with a continuous function and formulated a mixed integer linear program that can be solved for its global optimum. The method only gives a lower bound of the original SNPP, and its computational cost grows rapidly as the accuracy requirement for approximation increases.

All of the aforementioned studies are for deterministic SNPP. Accounting for inherent uncertainty in travel demand, some recent studies (e.g., Gardner et al. 2010; Li et al. 2012; Sumalee and Xu 2011) started to develop methods that also consider demand uncertainty in SNPP. Gardner et al. (2010) demonstrated better performance of a multiple point inflation/deflation solution method in comparison to that of single point approximation using GA and SA, in terms of computational time versus solution quality. The study was for first-best toll design. Li et al. (2012) considered demand uncertainty and environmental externalities for toll-design problems and used sample average approximation (SAA) and sensitivity analysis to solve for the optimal toll levels, given the tolled links. The multi-point approximation method or SAA with a local derivative based method requires a large number (depending on the sample size) of objective function evaluations; and this computational overhead will be costly when the network size is large. In addition, these studies focused on first-best tolling without addressing the more practical second-best toll design, not to mention the consideration of toll location selection. When these toll level design methods are incorporated into heuristics-based toll location optimization problems, the overhead of expensive objective evaluations (simulations) will increase dramatically, making it computationally intractable.

In summary, we feel that there are several important aspects of existing methods for SNPP that need to and can be significantly improved. First, due to the multi-modal nature of the objective functions in SNPP, derivative-based mathematical programming approaches for toll level design can only achieve local optimum. These methods are not suitable for discrete toll location optimization, either. Secondly, heuristic methods, although frequently used for simultaneous toll location and toll level design in SNPP, do not take advantage of the underlying system correlation structure in guiding their search process. Such heuristics generally cannot approach a global optimum within a limited computational budget. In fact, due to the combinatorial nature of toll location plus toll level alternatives, we would expect significant correlations of system performance across candidate solutions that share a common subset of links or that include links on parallel paths for some Origin-Destination (O-D) pairs. Thirdly, in those very limited studies that attempted to also deal with uncertainty in SNPP, the number of scenarios/repetitions used for simulation of candidate solutions was pre-determined and fixed. This leads to under- or over-simulation, since it foregoes the opportunity of adjusting the sampling budget dynamically according to the solution quality and the associated uncertainty. Therefore, if we can formulate an SNPP model and design a solution procedure that efficiently leverages upon the underlying correlation structure among different toll levels/locations combinations and their uncertainties, we would be able to improve both the capability and efficiency in approaching or finding the global optimum within limited computational constraints.

With this motivation, in this study we adopt a Bayesian Ranking and Selection (R&S) model and design new solution algorithms to address SNPP for joint optimization of toll locations and toll levels. R&S models (Kim and Nelson 2007) have shown superior performance, particularly under a limited sampling budget, in analyzing stochastic outcomes across various alternatives. In the Bayesian R&S formulation, we view each candidate solution to the SNPP as an alternative, and the objective function values are brained by taking "samples" using a Knowledge Gradient policy with Correlated Belief (KGCB) (Frazier et al. 2009). The

Bayesian R&S model fits nicely with SNPP due to its discrete formulation, flexible characterization of correlation structures, and capability to incorporate prior knowledge and the good performance of its sampling policies (e.g., Frazier et al. 2009; Ryzhov and Powell 2009). To adapt the original KGCB sampling strategy to SNPP (which typically has a very large number of alternatives), we further develop the Monte-Carlo-Linear-Belief-KG algorithm (MCLB-KG) based on a non-perfect additive linear belief model to reduce the computational cost for practical SNPP applications.

The paper is organized as follows. Section 2 introduces the mathematical model of SNPP. Section 3 formulates SNPP as a Bayesian R&S problem and describes the construction of the MCLB-KG policy for SNPP. Section 4 presents results and discusses computational examples. Section 5 concludes.

2 SECOND-BEST NETWORK PRICING PROBLEM (SNPP)

The SNPP is generally modeled by a bi-level program (e.g., Zuo et al. 2010; Ekström et al. 2012). Let's consider a bi-level program of SNPP with uncertainty. The upper-level problem models the decision maker's objective while the lower-level problem models the network users' travel behavior. We assume the expected traffic demand is known and inelastic to travel cost. The total travel time per time interval on the network is a commonly-adopted measure of traffic efficiency (e.g. Yang and Zhang 2003). The decision maker's objective is to minimize total travel time per unit time over the network by identifying and tolling a subset of links (selected from predefined candidate links) at appropriate toll levels. Each network user chooses the route with minimum cost to travel from her origin to her destination. Assuming homogenous unit "value of time" (VOT) among users, the formulation of SNPP is given as:

$$\text{(Upper-level)} \quad \max_{\boldsymbol{d}} \ \mathrm{E}(T_{\boldsymbol{d}}) = \sum_{\omega\in\Omega} p(\omega)\cdot\left[T_0(\omega) - \sum_{a\in A} z_a^*(\omega)t_a(z_a^*(\omega))\right], \tag{1}$$
$$s.t. \quad \boldsymbol{d} = [d_1 \ldots d_l], \ d_i \in \{0,1,\ldots,m\}, \ \forall i \in A'$$

$$\text{(Lower-level)} \quad \min_{z} \ Z(\boldsymbol{d}) = \sum_{a\in(A-A')} \int_0^{x_a} t_a(v)dv + \sum_{a\in A'} \int_0^{x_a} [t_a(v)+u_a]dv$$
$$s.t. \ \sum_k f_k^{rs} = q_{rs}, \ f_k^{rs} \ge 0, \quad \forall (r,s)\in D, \forall k \in K_{rs} \tag{2}$$
$$z_a = \sum_{(r,s)\in D}\sum_{k\in K_{rs}} f_k^{rs}\cdot\delta_{ak}^{rs},$$

In the upper level problem, we maximize $E(T_d)$, the expected difference between the total travel time of the no-toll scenario, T_0, and the total travel time of the tolled scenario, $\Sigma_{a\in A} z_a^* t_a(z_a^*)$. Link set $A=\{a\}$ is the set of all the directed links in the network; z_a^* is the traffic volume per unit time on link a under the optimal solution of the lower level problem; t_a is the corresponding travel time on link a, which is a function of z_a^*. In the objective function, uncertainty is generally considered by a countable scenario space Ω, the probability of each scenario $\omega\in\Omega$ is $p(\omega)$ (e.g., Gardner et al. 2010). Model inputs and parameters such as traffic demand can take on different values for different scenarios. Traffic demand plays a fundamental role determining network performance. Without loss of generality we use demand scenario space Ω for uncertainty consideration in SNPP. A subset of road links, $A'\subseteq A$ ($|A'|=l$) is the set of candidate links for pricing. A' is usually chosen empirically according to congestion level, toll facility installation and operation, existing ITS facilities, etc. (Yang and Zhang 2003; Zuo et al. 2009). $\boldsymbol{d} = (d_1,\ldots,d_l)^{\mathrm{T}}$ is an integer-valued decision vector, i.e., $d_i\in\{0,1,\ldots, m\}$ ($\forall i = 1, ..., l$) indicating the possible toll levels to be applied to each candidate link. For example, if m=3, then d_i = 0, 1, 2, 3 represent no toll, low, medium and high toll levels, respectively. Link travel time $t_a(z_a)$ is a non-decreasing convex function of traffic volume z_a, and the popular BPR formula (Bureau of Public Roads 1964) is used here: $t_a=t_a^0[1+\alpha(z_a/c_a)^\beta]$, where $\alpha>0$, $\beta>1$ are empirical parameters, t_a^0, c_a are free-flow travel time and capacity of link a, respectively. $u_a = e\cdot d_i$ / VOT (i corresponds to a) is the equivalent time cost of the toll on a candidate link $a\in A'$, e is the unit toll level.

Given candidate link set A', number of toll levels m and incremental unit e across toll levels, the key inputs to this upper-level maximization problem of SNPP is the traffic flow assignments $\boldsymbol{z}^* = \{z_a^*\}$ to all the directed links on the network resulting from the solution of the lower-level problem. The lower-level

problem is to find a user equilibrium (UE) flow pattern with potential equivalent time cost u_a considering link travel time and toll; under UE no user has incentive to change route. The UE problem has well-established solution methods like Frank-Wolfe algorithm (Sheffi 1985), but the computational cost grows significantly with the network size due to the shortest path subroutine involved. Since the lower- and upper-level problems in SNPP are hard to be integrated as one objective due to intrinsic difficulty of the problem (e.g., Yang and Zhang 2003), the SNPP can be regarded as a "black-box" discrete optimization problem. This nature of SNPP is at the heart of its Bayesian R&S formulation.

3 SNPP AS A BAYESIAN R&S PROBLEM WITH LINEAR BELIEFS

3.1 Bayesian R&S Formulation of SNPP

In a Bayesian R&S framework, we have M alternative decisions $\boldsymbol{X}=\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_M\}$ whose rewards (e.g., the values of objective functions for different pricing schemes in SNPP) are random with unknown mean $\boldsymbol{\theta}=(\theta_1,\ldots,\theta_M)^{\mathrm{T}}$ and unknown variance $\boldsymbol{\lambda}=(\lambda_1,\ldots,\lambda_M)^{\mathrm{T}}$. Our goal is to identify the alternative with the maximum expected reward through limited sample measurements. We have a prior belief about $\boldsymbol{\theta}$ with mean $\boldsymbol{\mu}^0=\{\mu^0{}_1\ldots\mu^0{}_{\mathrm{M}}\}$ and covariance $\boldsymbol{\Sigma}^0$ (an $M\times M$ positive semi-definite covariance matrix). For SNPP, we have network performance under different pricing schemes as $\boldsymbol{\theta}=(\mathrm{E}(T_{d1}),\ldots,\mathrm{E}(T_{dM})]^{\mathrm{T}}\sim N(\boldsymbol{\mu}^0,\boldsymbol{\Sigma}^0)$. Assume that we can evaluate N sample decisions, $x^0, x^1,\ldots, x^{N-1}$. At stage n, we make a measurement or evaluation of decision x^n, with measurement noise, $\varepsilon^n \sim N(0, \lambda_{x^n})$, independent across samples conditional on x^n. This yields sample observation (i.e., objective function evaluation in SNPP) $y^{n+1}=\theta_{x^n}+\varepsilon^n$. Let F^n be the sigma-algebra generated by $\{x^0\ldots x^{n-1}\}$ and $\{y^1\ldots y^n\}$. It is a well-known result that the conditional posterior distribution of $\boldsymbol{\theta}$ is also multivariate normal. Let $\boldsymbol{\mu}^n=\mathrm{E}(\boldsymbol{\theta}|F^n)$ and $\boldsymbol{\Sigma}^n=Cov(\boldsymbol{\theta}|F^n)$ be stage-n conditional expectation and covariance of $\boldsymbol{\theta}$, respectively, then $\boldsymbol{\mu}^n$ and $\boldsymbol{\Sigma}^n$ can be calculated recursively using standard results based on Bayes' Rule (Gelman et al. 2004):

$$\boldsymbol{\mu}^{n+1}=\boldsymbol{\mu}^n+\frac{y^{n+1}-\mu^n_{x^n}}{\Sigma^n_{x^nx^n}+\lambda_{x^n}}\boldsymbol{\Sigma}^n\boldsymbol{e}_{x^n};\ \boldsymbol{\Sigma}^{n+1}=\boldsymbol{\Sigma}^n-\frac{\boldsymbol{\Sigma}^n\boldsymbol{e}_{x^n}\boldsymbol{e}^T_{x^n}\boldsymbol{\Sigma}^n}{\Sigma^n_{x^nx^n}+\lambda_{x^n}}, \tag{3}$$

where $\boldsymbol{e}_x{}^n$ is a column vector of 0s with a 1 at position x^n. $\boldsymbol{\Sigma}^n{}_{x^nx^n}$ is the (x^n, x^n) diagonal entry of matrix $\boldsymbol{\Sigma}^n$. $\lambda_x{}^n$ is the performance variance corresponding to decision alternative x^n.

After N measurements through a sampling policy $\pi=\{x^1{}_\pi\ldots x^N{}_\pi\}$ from the policy space Π, we choose the alternative that yields the largest posterior mean of objective function value (rewards) as the optimal solution: $\sup_{\pi\ \Pi}\mathrm{E}^\pi[\max_x(\mu^0{}_x)]$, where E^π denotes the conditional expectation under π. In Bayesian R&S SNPP, as we evaluate pricing alternatives $\boldsymbol{d}_x{}^1,\ldots\boldsymbol{d}_x{}^N$, we obtain measurements of the random "rewards" that represent total network travel time reductions $T_d{}^1,\ldots,T_d{}^N$ in comparison to the no-toll scenario.

3.2 KGCB Sampling Policy

The Knowledge Gradient policy with Correlated Belief (KGCB policy) is originally introduced in Frazier et al.(2009). It samples alternative x that maximizes the incremental value (knowledge gradient) of the objective function:

$$x^{KG,n}(s)=\arg\max_x v^{KG,n}(x)=\arg\max_x(\mathrm{E}_n[\max_i \mu_i^{n+1}\mid S^n, x^n=x]-\max_i \mu_i^n), \tag{4}$$

where $S^n=(\boldsymbol{\mu}^n, \boldsymbol{\Sigma}^n)$ is the state of our posterior beliefs at measurement n. The KG-factor $v^{KG,n}(x)$ represents the incremental value (i.e., the expected improvement in posterior optimal value) obtained from measuring x at stage n. It is shown that the KGCB policy is almost-surely optimal for N=1 or $N\to\infty$, and has sub-optimality bounds when N is finite (Frazier et al. 2009).

By calculating the conditional predictive expectation of $\max_i\mu_i^{n+1}$, we can forecast the performance of all alternatives without taking actual samples of them. Therefore, one key step in KGCB policy is to compute conditional predictive distribution of μ_i^{n+1} given information at stage n. This conditional distribution is multivariate normal, with mean $\mathrm{E}_n[\mu^{n+1}]=\mu^n$ and covariance $\tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}^n, x^n)\tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}^n, x^n)^T$, where $\tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}, x)=\boldsymbol{\Sigma}^n e_x/(\lambda_x+\Sigma^n{}_{xx})^{0.5}$, details of this calculation can be found in Frazier et al. (2009). Thus the stage-

n conditional distribution of $\boldsymbol{\mu}^{n+1}$ is the same as $\boldsymbol{\mu}^{n+1} = \boldsymbol{\mu}^n + \tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}^n, x^n)Z$, where Z is scalar standard normal random variable. This allows us to compute $x^{KG,n}$ in (4) as:

$$x^{KG,n}(s) = \arg\max_x (\mathrm{E}_n[\max_i(\mu_i^n + \tilde{\sigma}(\Sigma^n, x^n)Z \mid S^n, x^n = x] - \max_i \mu_i^n) = \arg\max_x h(\mu^n, \tilde{\sigma}(\Sigma^n, x^n)), \quad (5)$$

where function h: $\mathrm{R}^M \times \mathrm{R}^M \rightarrow \mathrm{R}$ is defined as $h(\boldsymbol{p}, \boldsymbol{q}) = \mathrm{E}[\max_i p_i + q_i Z] - \max_i p_i$, $\boldsymbol{p}$ and $\boldsymbol{q}$ are deterministic M-dimensional vectors. Frazier et al. (2009) provides a method to compute $h(\boldsymbol{p}, \boldsymbol{q})$, where the entries of $\boldsymbol{p}$ and $\boldsymbol{q}$ are firstly sorted and then only the distinct ones retained to define a vector $\boldsymbol{c}$ with $c_j = (p_j - p_{i+1})/(q_{i+1} - q_j)$. These quantities are then used to calculate $h(\boldsymbol{p}, \boldsymbol{q}) = \sum_{i=1\sim|p|-1} (q_{i+1} - q_i) f(-|c_i|)$, where $f(z) = \varphi(z) + z\Phi(z)$, φ and Φ are standard normal PDF and CDF, respectively. We call this method "Subroutine-h", by which we can compute $h(\mu, \tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}, x))$ for any prior belief $\boldsymbol{\mu}$ and $\tilde{\boldsymbol{\sigma}}(\boldsymbol{\Sigma}, x)$. Then we are able to compute $v^{KG,n}(x)$ for each alternative x, and the largest $v^{KG,n}(x)$ gives the measurement of decision $x^{KG,n}$.

In the standard setting of KGCB, the dimensions of $\boldsymbol{q}$, $\boldsymbol{p}$ and $\boldsymbol{\mu}^n$ are the number of alternatives, M. Therefore, the Subroutine-h is executed M times for obtaining $x^{KG,n}$. Since the sorting step dominates the computational cost of Subroutine-h, so the overall complexity of the standard KGCB algorithm is $O(M^2 \log M)$. Thus for large number of alternatives, say $M > 10^5$ (which is usually the case for SNPP due to large number of link/toll combinations), the computational demand of the standard KGCB is prohibitive. This leads us to the modification of KGCB as follows.

3.3 Linear Belief Model for SNPP

In SNPP, the compounding effect of tolls on multiple links at their respective toll levels is not simply additive, i.e., summation of individual "link indices" of links a and b, $I_a + I_b$, will not simply be equal "location index" I_{ab}, the effect of simultaneous tolls on links a and b (Verhoef 2001). In fact, the interaction effects among tolled links, although hard to quantify, can be remarkable, especially among links on parallel paths for the same OD pair (Shepherd et al. 2001). Therefore, we propose a non-perfect additive linear belief model to consider such joint effect from tolling multiple links. This approach is inspired by the linear belief model used by Negoescu et al. (2011) in their study of sequential experimenting for drug discovery.

3.3.1 Model Structure and Priors on Model Coefficients

For our Bayesian R&S SNPP, we assume the marginal effect (on the final objective value) from one unit increase in toll rate on a link varies significantly across different toll levels. Thus we have $m \times l$ attributes (l candidate links, each with m toll levels). This leads to a new binary column decision vector $\boldsymbol{d}^e$ of size ml, expanded from the original l-dimension decision vector $\boldsymbol{d}$. By assigning m entries for each candidate link j in set A' and placing these m-entries across the links in the order of $j = 1,2,\ldots,l$, we have:

$$d^e{}_i = \begin{cases} 1 & [i - m(j-1)]^{\text{th}} \text{ toll level on link } j \\ 0 & \text{otherwise} \end{cases}, \quad \forall j \in (1,\ldots,l), \forall i \in [m(j-1), mj].$$

For example, consider a toy example with only two candidate links (dimensions) and two toll levels (attributes), i.e., $l = |A'| = 2$, $m = 2$, $j \in \{1,2\}$. Using the notation above, the no-toll alternative can be represented by $\boldsymbol{d}^e_{(00)} = (0,0,0,0)^T$, where the first two zeros correspond to first link and the last two for the second link. The alternative of applying toll level 2 on link 1 and toll level 1on link 2 is represented by $\boldsymbol{d}^e_{(21)} = (0,1,1,0)^T$.

We now assume a non-perfect linear additive model for the effect of a SNPP pricing scheme $\boldsymbol{d}_x$:

$$\theta_x = \eta_0 + \sum_{i=1}^{m \cdot l} \eta_i d^e_{x\,i} + \zeta_x \ , \quad (6)$$

where η_0 is the value for the no toll case (i.e., all entries in $\boldsymbol{d}^e_x$ are 0); coefficient $\eta^e{}_i$ represents the marginal effect per unit change in attribute $d^e{}_i$ (here it is the $[i - m(j-1)]^{\text{th}}$ toll level on link j); ζ is the deviation term from perfect additive structure, which is alternative specific as labeled by subscript x.

This non-perfect linear additive model is similar to that in Negoescu (2011). It is generalized from the perfect linear-additive model, Free-Wilson model (Free and Wilson 1964), by adding the deviation term ζ. Since only one toll level is to be implemented for each candidate link (as we focus on static network

optimization), i.e., $\sum_{m(j-1).i\le mj} d^e_i \le 1$, so the requirement of the Free-Wilson model that "at most one attribute is associated with each dimension" is automatically satisfied. Based on this linear belief model, if we sample $\boldsymbol{d}^e_x$ (corresponding to $\boldsymbol{d}_x$), the sample value would be:

$$T_{\boldsymbol{d}_x} = y_x = \eta_0 + \sum_{i=1}^{m\cdot l} \eta_i d^e_{x\,i} + \zeta_{\boldsymbol{x}} + \varepsilon_{\boldsymbol{x}} \quad , \tag{7}$$

where $\varepsilon_x \sim N(0, \lambda_x)$ is an independent measurement noise for $\boldsymbol{d}_x$. λ_x=0 models the deterministic SNPP, and λ_x>0 addresses SNPP with uncertainty (due to stochastic demand in this study).

Suppose we have independent priors of normal distributions η_0 and η_i, $\eta_0 \sim N(\mu_{\eta 0}, \sigma_{\eta 0}^2)$, $\eta_i \sim N(\mu_{\eta i}, \sigma_{\eta i}^2)$. We can also use independent normal distributions with mean 0 and variance σ_ζ^2 as priors for $\zeta_1, \ldots, \zeta_M$ (for any ζ with non-zero mean, it can be added to η_0). $\zeta_1,\ldots, \zeta_M$ are independent from other model coefficients. Then the prior belief about the mean value of decision $\boldsymbol{d}_x$ is:

$$\mu^0_{\boldsymbol{x}} = \mu_{\eta_0} + \sum_{i=1}^{m\cdot l} \mu_{\eta_i} d^e_{x\,i} \cdot \tag{8}$$

The prior belief of the covariance between the performance of $\boldsymbol{d}_x$ and $\boldsymbol{d}_{x'}$ is (Negoescu et al. 2011):

$$\Sigma^0(x,x') = \mathrm{cov}\Big(\eta_0 + \sum_{i=1}^{m\cdot l} \eta_i d^e_{x\,i} + \zeta_{\boldsymbol{x}}, \eta_0 + \sum_{i=1}^{m\cdot l} \eta_i d^e_{x'i} + \zeta_{\boldsymbol{x}'}\Big) = \sigma^2_{\eta_0} + \sum_{i=1}^{m\cdot l} d^e_{x\,i} d^e_{x'i} \sigma^2_{\eta_i} + \sigma^2_\zeta \cdot \mathbf{1}_{\{x=x'\}} . \tag{9}$$

3.3.2 Posterior Distributions on Model Coefficients

Maintaining and updating posteriors on linear belief model coefficients (marginal effects of different attributes) is a key step in solving the Bayesian R&S SNPP. Let the column coefficient vector $\boldsymbol{\eta}$ denote $(\eta_0, \eta_{1,\ldots,}\eta_{ml})^{\mathrm{T}}$ and $\boldsymbol{D}_{M\times(ml+1)}$ be a matrix comprised of rows each representing the attribute values of an alternative plus a "1" in the first entry corresponding to baseline (no-toll scheme) constant η_0. Thus a row in $\boldsymbol{D}$ is a "1" followed by the attribute values of $\boldsymbol{d}^e$. We also use a column vector $\boldsymbol{\zeta}$ to denote all the deviation terms (ζ_x). With these notations, the true value vector is $\boldsymbol{\theta} = \boldsymbol{D}\cdot\boldsymbol{\eta}+\boldsymbol{\zeta}$ by (6). Even though the number of ζ_x is M, which is generally very large in SNPP, we only need to maintain a mean vector and covariance matrix of ζ_x for alternatives that have already been measured. If we have not measured a alternative x by stage n, then the posterior of ζ_x will stay the same as its prior. ζ_x remains independent of η_i's, and all other deviation terms.

Toward this end, we define column vector $\boldsymbol{\eta}_n$ that contains $\boldsymbol{\eta}$ and ζ_x terms for an alternative x in ($x^0,\ldots, x^{n-1}$). Let $\boldsymbol{a}^n$ and $\boldsymbol{C}^n$ be the mean and covariance of our stage-n posterior of $\boldsymbol{\eta}_n$. Note that before the first measurement, the initial values are $\eta_0 = \eta$, $\boldsymbol{a}^0{}_{(ml+1)\times 1} = \{\mu_{\eta i}\}$, and diagonal matrix $\boldsymbol{C}^0{}_{(ml+1)\times(ml+1)}$ with diagonal entries $\{\sigma_{\eta i}^2\}$. There is a recursive expression for $\boldsymbol{a}^n$ and $\boldsymbol{C}^n$ (Negoescu et al. 2011):

$$\boldsymbol{a}^{n+1} = \tilde{\boldsymbol{a}}^n + \frac{y^{n+1} - (\tilde{\boldsymbol{a}}^n)^T \tilde{\boldsymbol{d}}_{x^n}}{\lambda_{x^n} + (\tilde{\boldsymbol{d}}_{x^n})^T \tilde{\boldsymbol{C}}^n \tilde{\boldsymbol{d}}_{x^n}} \tilde{\boldsymbol{C}}^n \tilde{\boldsymbol{d}}_{x^n};\ \tilde{\boldsymbol{C}}^{n+1} = \tilde{\boldsymbol{C}}^n - \frac{\tilde{\boldsymbol{C}}^n \tilde{\boldsymbol{d}}_{x^n} (\tilde{\boldsymbol{d}}_{x^n})^T \tilde{\boldsymbol{C}}^n}{\lambda_{x^n} + (\tilde{\boldsymbol{d}}_{x^n})^T \tilde{\boldsymbol{C}}^n \tilde{\boldsymbol{d}}_{x^n}}, \tag{10}$$

where $\tilde{\boldsymbol{a}}^{n-1}$and $\tilde{\boldsymbol{C}}^{n-1}$ are defined as below: if x^n has been previously measured by stage n, $\tilde{\boldsymbol{a}}^{n-1}=a^n, \tilde{\boldsymbol{C}}^{n-1}=C^n$; otherwise, let $\tilde{\boldsymbol{a}}^{n-1}$ be the column vector formed by adding a 0 to $\boldsymbol{a}^{n-1}$, and $\tilde{\boldsymbol{C}}^{n-1}$ be the matrix formed by adding a row and a column after the last row and column of $\boldsymbol{C}^{n-1}$, where all the entries of the new row and new column are 0 but the diagonal entry is σ_ζ^2. Then the posterior of $\boldsymbol{\eta}_n$ in stage n-1 is $N(\tilde{\boldsymbol{a}}^{n-1}, \tilde{\boldsymbol{C}}^{n-1})$. $\tilde{\boldsymbol{d}}_{x^n}$ is a column vector consisting of 1's at indices of $\boldsymbol{\eta}_{n+1}$ for which alternative x^n contains the corresponding baseline term, toll level attributes and deviation term, and 0's elsewhere. (10) is a linear square recursive model (e.g., Powell and Ryzhov 2012) modified by incorporating the deviation terms from our non-perfect linear additive model for SNPP. These updating equations allow us to track and update our beliefs about η_n in a computationally efficient way. The prior beliefs of the model coefficients, parameterized by $\mu_{\eta 0}$, $\sigma_{\eta 0}^2$, $\mu_{\eta i}$, $\sigma_{\eta i}^2$, σ_ζ^2 can be estimated from initial sampling or prior information.

Based on the updated beliefs of the hyperparameters, we can construct the posteriors of the alternatives' values. Noting that any multivariate normal belief on η_n induces a multivariate normal belief on $\boldsymbol{\theta}^n$ (Negoescu et al. 2011), we have $\boldsymbol{\theta}^n \sim N(\boldsymbol{\mu}^n, \boldsymbol{\Sigma}^n)$ from $\boldsymbol{\eta}_n \sim N(\boldsymbol{a}^n, \boldsymbol{C}^n)$. $\boldsymbol{\mu}^n$ and $\boldsymbol{\Sigma}^n$ are calculated from $\boldsymbol{a}^n$ and $\boldsymbol{C}^n$

using the same method as (8) and (9). However, to use KGCB algorithm, we only need to retrieve partial information without computing the whole $\boldsymbol{\Sigma}^n$ matrix, which is prohibitive in SNPP.

3.4 KGCB Algorithm for SNPP with Linear Beliefs and MC Sampling for the Hyperparameters

Now we can compute the KG factors from a belief on $\boldsymbol{\eta}_n$ parameterized by $\boldsymbol{a}^n$ and $\boldsymbol{C}^n$. By (5), we can obtain $v^{KG,n}(x)$ using Subroutine-*h* when $\boldsymbol{\mu}^n$ and $\tilde{\sigma}(\boldsymbol{\Sigma}^n, x)$ are available. Independent of x, $\boldsymbol{\mu}^n = \boldsymbol{D}^n\boldsymbol{a}$, where $\boldsymbol{D}^n$ is a $M\times|\boldsymbol{\eta}_n|$ indexing matrix of 0's and 1's, each row corresponds to an alternative and has 1's for the baseline term, toll attribute terms and the deviation terms from $\boldsymbol{a}^n$ that are contained in the alternative. To compute $\tilde{\sigma}(\boldsymbol{\Sigma}^n, x)$, we set $x^n = x$ and get the corresponding $\boldsymbol{\eta}_{n+1}$ and $\widetilde{\boldsymbol{C}}^n$. Let $\widetilde{\boldsymbol{D}}^n$ be a $M\times|\boldsymbol{\eta}_{n+1}|$ matrix that is similar to $\boldsymbol{D}^n$, except that it maps alternatives to component of $\boldsymbol{\eta}_{n+1}$ instead of $\boldsymbol{\eta}_n$. Note that the beliefs on those ζ_x terms that are not included in $\boldsymbol{\eta}_{n+1}$ will not change as a result of measuring x^n. In addition, $\tilde{\sigma}(\boldsymbol{\Sigma}^n, x^n)$ is not affected by such deviation terms. So we can ignore the left-out deviation terms, by Frazier et al. (2009), $\tilde{\sigma}(\boldsymbol{\Sigma}^n, x) = \boldsymbol{\Sigma}_{x,\cdot}{}^n/(\lambda_x+\Sigma^n{}_{xx})^{0.5}$, where the x^{th} column of $\boldsymbol{\Sigma}^n$ is:

$$\boldsymbol{\Sigma}^n_{x,\cdot} = (\widetilde{\boldsymbol{D}}^n\widetilde{\boldsymbol{C}}^n)(\widetilde{\boldsymbol{D}}^n_{x,\cdot})^T. \tag{11}$$

However, when M is large, computing $v_{\boldsymbol{d}}^{KG,n}$ for all decision vectors $\boldsymbol{d}$ as required in standard KGCB algorithm is very expensive. Inspired by Ryzhov and Powell (2009), we propose a Monte Carlo (MC) sampling step to substantially reduce the size of the choice set. But instead of sampling $\boldsymbol{\theta}$ from $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as used in Ryzhov and Powell (2009), we directly sample from hyperparameter space and generate realizations of $\boldsymbol{\theta}$ according to the linear belief model, which in the first place permits significant savings. Suppose we generate K sample realizations of $\boldsymbol{\theta}'^n$ based on the non-perfect linear additive models and the posterior beliefs of the model coefficients at stage n. Let $\boldsymbol{\eta}'_n(\omega_k)$ be the k^{th} sample realization of model coefficients from the posterior distribution $N(\boldsymbol{a}^n, \boldsymbol{C}^n)$. The M-dimensional column vector $\boldsymbol{\zeta}'^n(\omega_k)$ has entries of 0's for sampled alternatives, and each entry of $\zeta_x'^n(\omega_k)$ corresponds to unmeasured alternatives is separately drawn from the prior distribution $N(0, \sigma_\zeta{}^2)$. Then the mean $\boldsymbol{\theta}'^n(\omega_k)$ of the k^{th} sample realization will be an M-dimensional column vector for each k:

$$\boldsymbol{\theta}'^n(\omega_k) = \boldsymbol{D}^n\boldsymbol{\eta}'_n(\omega_k) + \boldsymbol{\zeta}'^n(\omega_k). \tag{12}$$

Let $t_k = \text{argmax}_t\, \theta'^n_t(\omega_k)$ be the toll alternative that appears to be the best from sample k and let K_0 be the number of such distinct alternatives from all K samples. The number of alternatives in SNPP is much larger than that encountered in Ryzhov and Powell (2009), so as a remedy, in iteration n, we propose to randomly sample K_1 distinct alternatives $s_1, \ldots, s_{K1}$ from the complete alternative space and use the final choice set $S=\{t_1, \ldots, t_{K0}\} \cup \{s_1, \ldots, s_{K1}\}$. Then $v^{KG,n}$ can be computed over set S. We call this the Monte Carlo linear belief KG policy (MCLB-KG), which is adopted for our challenging SNPP. So the complexity of MCLB-KG algorithm in v^{KG} calculation becomes $O(|S|^2\log|S|)$, much less than $O(M^2\log M)$ in the standard KGCB policy. Also we only need to sample a $|\boldsymbol{\eta}_n|$-dimensional vector ($|\boldsymbol{\eta}_n| \le n+lm+1$) from multivariate normal distribution $N(\boldsymbol{a}^n, \boldsymbol{C}^n)$ at iteration n, the complexity of the MC sampling step is $O(K|\boldsymbol{\eta}_n|^3)$ when the Cholesky factorization of $\boldsymbol{C}^n$ is used (which is very efficient in modern computing package). So this implementation is $<< O(M^2\log M)$, Note that recognizing those unmeasured alternatives by stage n can be done efficiently by keeping a list of sampled alternatives rather than looping over tags for M alternatives. So the other overhead of the MCLB-KG algorithm mainly comes from the multiplications of high-dimension matrices in (11) and (12), which have only linear dependency on M. Therefore, the computational cost can be significantly reduced compared to the standard KGCB policy.

3.5 Updating the Unknown Variances

For the SNPP under demand uncertainty, the variance for each alternative x, λ_x (i.e., the variance of measurement noise ε_x in (7)) is usually unknown. With very limited sampling budget, this variance affects the belief update of the hyperparameters in (10) and the characterization of conditional distribution of $\boldsymbol{\mu}^{n+1}$. Therefore, an estimation updating procedure of λ_x is needed to improve the performance of MCLB-KG policy. We use an approach inspired by the Bayesian normal model with known mean and unknown

variance (Gelman et al. 2004). We start with a prior belief λ_x^0 that is constant or varying across alternatives, it can be simply the best guess based on the information available. As the learning progresses in implementing the solution algorithm, we can collect more samples for a certain $\boldsymbol{d}_x$ and update our estimate of that λ_x. In iteration n where alternative x is sampled, we use:

$$\lambda_x^n = \frac{\lambda_x^{n-1}(w+ns_x^n-3)+(y_x^n-\mu_x^n)^2}{w+ns_x^n-2} = \frac{w\lambda_x^0+\sum_{i=1}^{ns_x^n}(y_x^{n_x(i)}-\mu_x^{n_x(i)})^2}{w+ns_x^n-2} \approx \frac{w\lambda_x^0+\sum_{i=1}^{ns_x^n}(y_x^{n_x(i)}-\theta_x)^2}{w+ns_x^n-2}, x\in\{x^1\ldots x^N\}, \quad (13)$$

where ns^n_x is the x^{th} entry in the M-dimensional vector $\boldsymbol{ns}^n$ used to record how many times each alternative has been sampled up to stage n; iteration $n_x(i)$ is the iteration when alterative x is sampled for the i^{th} time ; $\mu_x^{n(i)}$ is the posterior mean for x at iteration $n_x(i)$; and weight $w\geq 0$. The idea behind (13) is to estimate λ_x as a weighted average of the prior belief and the information observed by the samples. To marginalize the impact of inaccuracy from posterior means, we require $n \geq 3$ before (13) is applied. As the number of samples increases, the variance estimates will gradually converge to their true values.

4 NUMERICAL EXPERIMENT AND DISCUSSION

We apply the method to the benchmark Sioux Falls network, which is used in recent SNPP studies (e.g., Ekström et al., 2012). It has 24 nodes and 76 links and 576 OD pairs (see Fig. 1(a)) with detailed network date given in Bar-Gera (2013). Due to budget constraints, 10 candidate links A'={16, 19, 29, 39, 48, 49, 52, 66, 74, 75} based on initial congestion levels are of interest, 3 toll levels are proposed with unit toll level e=\$2. The homogenous VOT = \$1/min. The total travel time under base demand is $T_0 = 8\times10^6$ min per unit time. In the implementation of MCLB-KG policy, we set the number of random samples K=100 and K_1=200. We use non-informative priors for most of the parameters in the belief model: $\sigma_\zeta^2 = 10^5$, $\sigma_{\eta i}^2 = 4\times10^5$, $\mu_{\eta i} = 400, 800, 1200$ for toll levels 1, 2 and 3, respectively, and $\sigma_\zeta^2 = 10^6$. Although the prior means of the toll attributes' marginal effects are positive, large uncertainties are attached to these coefficients as well as to the deviation terms. We have almost complete information about the baseline no-toll alternative, so we set $\mu_{\eta 0} = 0$ and $\sigma_{\eta 0}^2 = 10$ for deterministic tests and $\sigma_{\eta 0}^2 = 10^4$ for stochastic case. We use a non-informative prior to demonstrate the effective learning capability of the MCLB-KG policy for SNPP.

We examine the performance of Bayesian R&S SNPP model solved via MCLB-KG algorithms in comparison to the GA (which is usually used for SNPP) for deterministic setting (λ_x=0) and SAA-GA for stochastic (λ_x>0) setting. We use the standard GA (Deep et al. (2009)) with population size $|A'|$ and elite size 1 (optimized by grid search). SAA is used for GA to evaluate individual solution (e.g., Gardner et al.(2010)) with sample size 5 for stochastic setting (performs best among 2~6). The simulation budget N is 100 and 300 for the deterministic and stochastic tests, respectively. Because in stochastic case, evaluation of one solution contains 2 simulations under the same demand realization, one for the toll alternative and the other for the non-toll one, and the SAA sample size is 5 for GA, so this means 100 and 150 iterations in R&S and $\lceil 100/|A'|\rceil$ and $\lceil 30/|A'|\rceil$ generations in GA (or SAA-GA) for the deterministic and stochastic cases, respectively. In stochastic test, each OD demand q_{rs} in (2) is subject to a common p% ~$N(0, v^2)$ perturbation, q_{rs} is set to 0 if it drops below 0. Two cases v=0.01 and 0.05 are tested, with $\lambda_x^0=10^5$ and 4×10^5, respectively. We run 10 independent sample paths for each algorithm in both deterministic and stochastic tests. The main performance measure is the Relative Opportunity Cost (RelOC), defined as the relative difference between the objective value of the true optimal solution (the best solution possibly known) and the objective value of the "best" alternative proposed by the algorithm.

Note that normally distributed perturbation in traffic demand does not necessarily results in normally distributed objective value T_d, as shown in Fig. 1(b). We can see the travel time reductions under two toll alternatives are not affine in demand with markedly different patterns. This is due to the nonlinear function $t_a(z_a)$ and complicated system response of underlying UE flow assigment. We use this deviation from normality to show the robustness of the normality based Bayesian R&S algorithm for practical problems such as SNPP. This also indicates that the objective value evaluated at base demand may not be the true value of an toll alternative. So the mean of 1000 Monte Carlo samples is used as this "true" value.

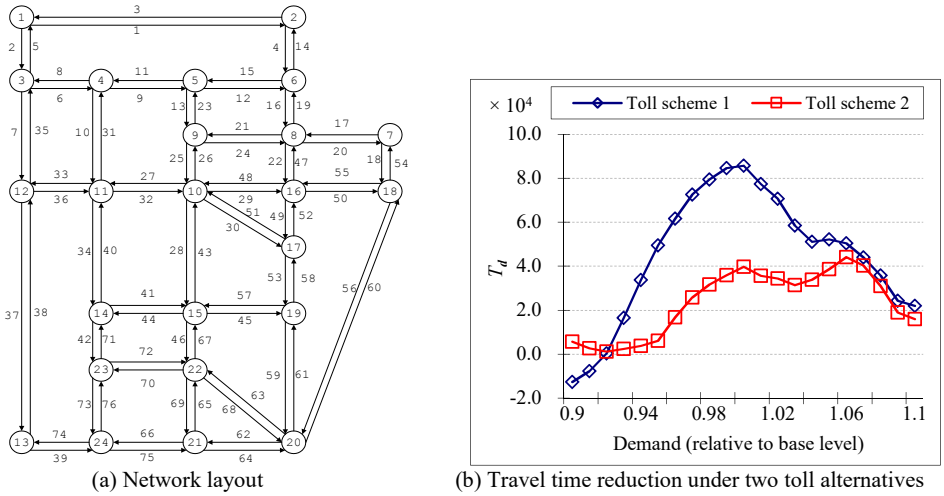

(a) Network layout (b) Travel time reduction under two toll alternatives

Figure 1: Sioux Falls test network and non-normality of alternative values.

Fig. 2(a), (b) and (c) compare the RelOC between Bayesian R&S SNPP solved by MCLB-KG and solved by GA or SAA-GA (point estimate of each RelOC with its ~95% confidence interval (CI) plotted). In all three cases, the MCLB-KG algorithms outperform GA or SAA-GA, approaching the best solution within fewer simulations and has a constantly better RelOC within the simulation budget. In fact, in deterministic case, MCLB-KG finds the true optimum within 80~100 iterations in most sample paths, while GA often stays in local optimum with RelOC values above 0.1 after reaching the 100 sampling budget. Fig. 2(d) shows how many times of alternatives in each region j are measured (j=1,…,10) by each algorithm in three typical sample paths. As can be seen, GA tends to spend most time around certain area (near local optimum) while the MCLB-KG algorithm explores across the decision space more evenly. In fact, The global exploration of the MCLB-KG algorithm happens in earlier iterations accounting for larger uncertainties in the hyperparameters and then the algorithm quickly identifies promising regions to spend more iterations in. In the stochastic setting, our algorithm also explores across the decision space while SAA-GA's searching is much more localized, similar to the observation in the deterministic case.

Take case v=0.01 as an example, Figs. 3(a) shows the entries of posterior mean vector $\boldsymbol{a}^N$ and diagonal entries of covariance matrix $\boldsymbol{C}^N$ resulting from MCLB-KG. We see that the absolute values of the posterior means of model coefficients for most attributes are well above that of the sampled deviation terms, and the posterior variances of the deviation terms are smaller than those for the coefficients of toll attributes. This explains why the non-perfect additive linear models are useful for SNPP. They are also true for v=0.05 (although the absolute values of $\boldsymbol{a}^N$ entries decrease) and the deterministic case. Under larger v=0.05, the relative ranking among the posterior means of different model coefficients remain almost unchanged , and posterior variances of the toll attribute effects increase, which is not surprising. Based on these posterior means and e=2$, we compute and plot the marginal effects of toll rates for each tollable link, as shown in Fig.3 (b). The results suggest that most links have positive expected marginal effect on travel time reduction at all toll levels, but interestingly, the expected marginal effect of link 66 and 75 are positive at toll level 1 and 2 but negative at toll level 3. Link 16 and 19 have negative expected marginal effect at all toll levels, indicating that the initial congestion level may not always be a good criteria for selecting candidate links. The notable variations of the marginal effects across links and toll levels justify the belief model used.

We also note in the test that measurement decisions $x^{KG,n}$ are usually from the set $\{t_1, \ldots, t_{K0}\}$ by MC sampling, which is a bigger set in earlier iterations ($n \leq 20$). However, in later iterations (after enough observations that make the belief upon those hyperparameters relatively stable and outweigh the effect of non-informative priors), the MC step often selects only one alternative t_1 (i.e., K_0=1). Interestingly, this t_1 then often stalls for several iterations before a change is invoked by a relatively significant refinement of the belief in the linear model coefficients after sampling a new "promising" alternative from the set $\{s_1, \ldots,$

$s_{K1}\}$. This shows the effectiveness of MC sampling as well as the necessity of including extra alternatives in the candidate set S in each iteration, particularly when the number of alternatives is large.

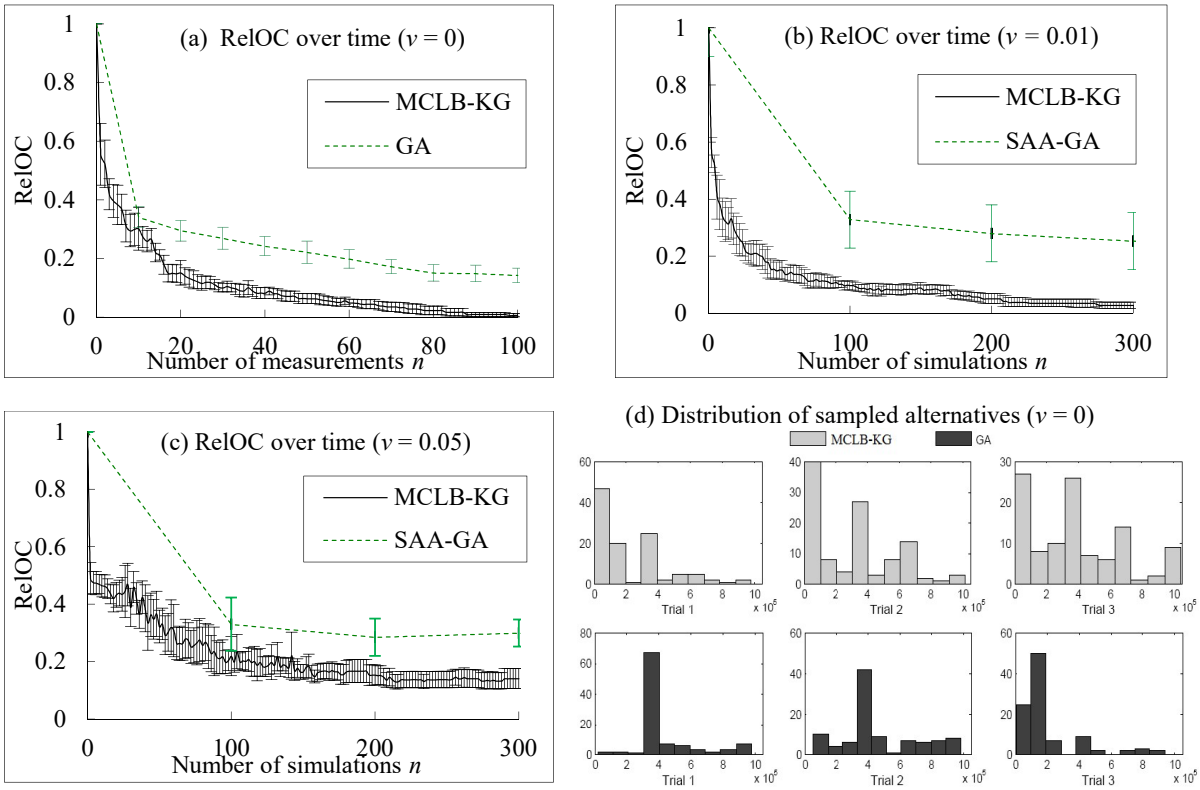

Figure 2: Performance comparison between two algorithms.

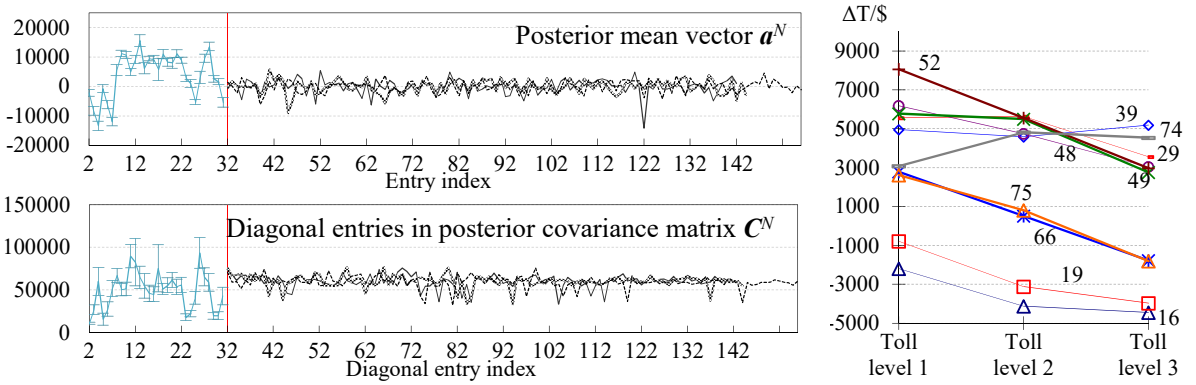

(a) Posteriors on model coefficients (the lines on the left: posteriors of η_i's with ~95% CI, the lines on the right: posterior means of ζ_x's under 3 typical trials)

(b) Marginal effect at three toll levels (the link number is beside the line)

Figure 3: Posterior distributions on hyperparameters and marginal effects of toll attributes (v = 0.01).

Finally, Table 1 shows the average per iteration computation time of each algorithm over 10 sample paths, the MCLB-KG spends most time on sampling decision (MC step included), almost 10^3 times as that of the GA (or SAA-GA). This is mainly due to computing the KG-factor over the whole choice set especially during the earlier stages when candidate alternatives $\{t_1, \ldots, t_{K0}\}$ are more diversified with larger K_0. Besides the doubled simulation time per iteration (due to evaluating $T_0(\omega)$ in addition to $T_x(\omega)$), another

significant difference of the stochastic case compared to the deterministic case is that the average time spent on the MC sampling step increased by ~30% under v=0.05. This is because during earlier iterations more candidate alternatives are generated due to bigger uncertainty on the hyperparameters. Such uncertainty decreases significantly as measurements accumulate, but with K_0 drops in a slower rate compared to that in the case v=0 or 0.01. However, although the total computational time by MCLBKG is bigger in this test network, it considerably reduces the total number of simulations needed for reaching a satisfactory RelOC compared to the SAA-GA. This will bring us substantial time savings for large networks when each simulation takes hours even days, which is usually the case for SNPP in practice (e.g., Zuo et al. 2009).

Table 1: Average computational time per iteration. *: KG-factor computing (MC sampling)

v	Algorithm	Simulation (s)	Sampling decision (s)	Update (s)
0	GA	17.1	0.13	< 0.01
	MCLB-KG	16.5	41.4 (57.5)*	0.81
0.01	SAA-GA	33.6	0.07	< 0.01
	MCLB-KG	33.4	39.5 (62.4)	0.79
0.05	SAA-GA	34.0	0.05	< 0.01
	MCLB-KG	33.6	37.8 (76.4)	0.67

5 CONCLUSION

We have proposed a Bayesian R&S model for the Second-best Network Pricing Problem (SNPP) choosing toll locations and rates simultaneously. The large number of alternatives, combinatory nature and random demand make the problem challenging. We adopt a linear belief KG policy to solve the SNPP. As an extension of Ryzhov and Powell (2009) to the linear belief setting, MC-sampling of the hyperparameters is proposed to reduce the choice set. Experiment results on a SNPP with 4^{10} alternatives show good performance of the method and its superiority to the SAA-GA benchmark. We believe this is a promising tool for real-world SNPP under limited sampling budget. The successful application of the parameterized belief model tailored to SNPP also sheds lights on the underlying features of the problem itself.

REFERENCES

Bar-Gera, H. 2013. Transportation network test problems. http://www.bgu.ac.il/~bargera/tntp/.

Bergendorff, P., D. W. Hearn, M.B. Ramana. 1997. "Congestion Toll Pricing of Traffic Networks." *Lecture Notes in Economics and Mathematical Systems* 450: 51-71.

Bureau of Public Roads. 1964. *Traffic Assignment Manual*. U.S. Department of Commerce, Urban Planning Division, Washington D.C.

Deep, K., K. P. Singh, M.L. Kansal, C. Mohan. 2009. "A Real Coded Genetic Algorithm for Solving Integer and Mixed Integer Optimization Problems." *Applied Mathematics and Computation* 212(2): 505–518.

Ekström, J., A. Sumalee, H. Lo. 2012. "Optimizing Toll Locations and Levels Using a Mixed Integer Linear Approximation Approach." *Transportation Research Part B* 46 (7): 834-854.

Ekström, J., N. H. Quttineh. 2014. "Surrogate-Based Optimisation of Toll Levels in Congestion Pricing Schemes." In *Proceedings of the 19th International Conference of Hong Kong Society for Transportation Studies*, HKSTS.

Frazier, P. I., W. B. Powell, and S. Dayanik. 2009. "The Knowledge-Gradient Policy for Correlated Normal Rewards." *INFORMS Journal on Computing* 21(4): 599-613.

Free, S., J. Wilson. 1964. "A Mathematical Contribution to Structure-Activity Studies." *Journal of Medicinal Chemistry* 7 (4): 395–399.

Gardner, L. M., A. Unnikrishnan, S.T. Wallar. 2010. "Solution Methods for Robust Pricing of Transportation Networks under Uncertain Demand." *Transportation Research Part C* 18 (5): 656-667.

Gelman, A., B. Carlin and H. Stern et al. 2004. *Bayesian Data Analysis*. Boca Raton, Florida: Chapman & Hall/CRC.

Hearn, D. W., and M. B. Yildirim. 2001. "A Toll Pricing Framework for Traffic Assignment Problems with Elastic Demand." Research Report No. 99-8, Department of Industrial and Systems Engineering, University of Florida, Gainesville, Florida.

Kim, S. H., and B. L. Nelson. 2007. "Recent Advances in Ranking and Selection." In *Proceedings of the 2007 Winter Simulation Conference*, edited by S.G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J.D. Tew and R.R. Barton, 162–172. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Lawphongpanich, S., and D. W. Hearn. 2004. "An MPEC Approach to Second Best Toll Pricing." *Mathematical Programming* B (101): 33-55.

Li, Z., W. H. K. Lam, S.C. Wong, A. Sumalee. 2012. "Environmentally Sustainable Toll Design for Congested Road Networks with Uncertain Demand." *International Journal of Sustainable Transportation* 6 (3): 127-155.

Negoescu, D. M., P. I. Frazier, W. B. Powell. 2011. "The Knowledge-Gradient Algorithm for Sequencing Experiments in Drug Discovery." *INFORMS Journal on Computing* 23(3): 346-363.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*. Hoboken, New Jersey: John Wiley and Sons.

Ryzhov, I. O., and W. B. Powell. 2009. "A Monte Carlo Knowledge Gradient Method for Learning Abatement Potential of Emissions Reduction Technologies." In *Proceedings of the 2009 Winter Simulation Conference*, edited by M.D. Rosetti, R.R. Hill, B. Johansson, A. Dunkin, and R.G. Ingalls, 1492-1502. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Sheffi, Y. 1985. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

Shepherd, S. P., A. D. May, D.S. Milne, A. Sumalee. 2001. "Practical Algorithms for Finding the Optimal Road Pricing Location and Charges". In *Proceedings of European Transport Conference*, Cambridge.

Shepherd, S., and A. Sumalee. 2004. "A Genetic Algorithm Based Approach to Optimal Toll Level and Location Problems." *Networks and Spatial Economics* 4 (2): 161-179.

Sumalee, A., and W. Xu. 2011. "First-Best Marginal Cost Toll for a Traffic Network with Stochastic Demand." *Transportation Research Part B* 45(1): 41-59.

Verhoef, E. T., 2001. "Second-Best Congestion Pricing in General Static Transportation Networks with Elastic Demands." *Regional Science and Urban Economics* 32 (3): 281-310.

Yang, H., and H. Huang. 1998. "Principle of Marginal-Cost Pricing: How Does It Work in a General Road Network?" *Transportation Research Part A* 32 (1): 45-54.

Yang, H., and X. Zhang. 2003. "Optimal Toll Design in Second-Best Link-Based Congestion Pricing." *Transportation Research Record*, 1857 (2003): 85-92.

Zuo, Z., R. Kanamori, T. Miwa, T. Morikawa. 2010. "Comparison of Cordon and Optimal Toll Points Road Pricing Using Genetic Algorithm." In *Proceedings of American Society of Civil Engineers (ASCE) 2010 Traffic and Transportation Studies Conference*, 535-544.

AUTHOR BIOGRAPHIES

ZHEN TAN is a PhD student at Cornell University with major in Civil and Environmental Engineering and minor in Applied Mathematics. His research interests are operation research and machine learning applications in transportation and environmental systems. His email address is zt78@cornell.edu.

H. OLIVER GAO is an Associate Professor in the School of Civil and Environmental Engineering at Cornell University. He holds a Ph.D. in Civil and Environmental Engineering from University of California, Davis, CA, His research focuses on quantitative modeling and development of engineering systems solutions for sustainable and intelligent infrastructure and lifeline systems, low carbon and low emission transportation systems, and the closely related environment-energy systems. He is Editor-in-Chief of the international journal Transportation Research Part D: Transport and Environment. His email address is hg55@cornell.edu.