

Contextual Bayesian optimization of congestion pricing with day-to-day dynamics

Renming Liu^{a,*}, Yu Jiang^a, Ravi Seshadri^a, Moshe Ben-Akiva^b, Carlos Lima Azevedo^a

^a Department of Technology, Management and Economics, Technical University of Denmark, Denmark

^b CEE, Massachusetts Institute of Technology, USA

ARTICLE INFO

Keywords:

Congestion pricing
Demand management
Day-to-day dynamics
Contextual Bayesian optimization

ABSTRACT

Congestion pricing is a common approach to alleviate urban traffic congestion. The design of second-best congestion pricing schemes is typically formulated as non-linear programming and bi-level optimization problems, where the lower-level problem involves either a static or dynamic network equilibrium model. The complexity of these bi-level toll optimization problems increases considerably when incorporating day-to-day dynamic models of travel behavior and dynamic models of network congestion. These models are often operationalized using simulation, and consequently, the toll design problem is a computationally challenging simulation-based optimization problem where the evaluation of a single candidate pricing scheme involves simulating the day-to-day model until convergence. In order to circumvent this issue, we propose a contextual Bayesian optimization (BO) framework, where the BO scheme is embedded within the day-to-day dynamic model by using temporal contextual information. The framework implicitly incorporates the relationship between the objective function across days and uses past days' observations (function evaluations) as weak priors when constructing the Gaussian process underlying the BO algorithm for the current day's toll optimization problem, resulting in gains in computational efficiency.

The contextual BO approach is applied to the design of distance-based pricing schemes for the morning commute problem. We demonstrate numerically that the scheme converges to the system optimum, and moreover, utilizes a significantly smaller number of simulation evaluations (ten-fold reduction) than the standard approach wherein each function evaluation involves simulating the day-to-day model until convergence. From a policy perspective, we find that the distance-based schemes yield significant welfare gains relative to area-based schemes and show that the design of the distance-based tariff scheme can significantly affect distributional impacts. A suitably designed two-part tariff structure can partially offset the relatively large welfare losses of travelers with longer commute distances while maintaining overall welfare. The proposed contextual BO scheme is also extended to incorporate context specific demand and supply information, which can be of value to policy-makers when evaluating optimal toll design schemes under a wide range of scenarios in a computationally tractable manner.

* Corresponding author.

E-mail address: liure@dtu.dk (R. Liu).

<https://doi.org/10.1016/j.tra.2023.103927>

1. Introduction

Road traffic congestion is a critical problem affecting urban mobility worldwide and its severity continues to increase, causing significant costs at the individual and societal level (Eurostat and E.U. Commission, 2018; Schrank et al., 2015). While a significant agenda has been put forward on the transport supply side, mostly driven by vehicle technology (automation and electrification), demand shifts are often considered a hard-to-reach but effective means to reduce the social and environmental costs associated with transport. Demand management has thus become an increasingly important focus of the policy agenda in many metropolitan areas.

As a promising demand management policy, congestion pricing has been widely investigated in the literature and successfully implemented in practice (Langmyhr, 1999; Yang and Huang, 2005; Lindsey, 2006; Vonk Noordegraaf et al., 2014; Gu et al., 2018). Since the seminal study by Pigou (1913), numerous first-best pricing models have been proposed, where exact system optimum (SO) solutions are analytically derived given the explicit demand and cost functions (Smith, 1979; Arnott et al., 1990; Yin and Yang, 2004; Yang and Huang, 2005). However, in practice these functions can be empirically difficult to estimate. Consequently, several studies have attempted to determine the congestion pricing toll rates via a trial-and-error approach, wherein tolls are dynamically updated and imposed without knowledge of the demand functions. For example, Li (1999) proposed an iterative bisection algorithm to optimize tolls using observable traffic count data and travel cost functions for a single road link, and analytically proved its convergence to SO (Li, 2002). Yang et al. (2004) then extended this trial-and-error approach based on the method of successive averages (MSA) to the first-best pricing problem in a general network.

However, first-best pricing is difficult to implement in practice due to several complications as summarized in Small et al. (2007), de Palma and Lindsey (2011); hence, various second-best pricing schemes have been explored in both theory and practice. Typical second-best pricing strategies include facility-based schemes, e.g., high occupancy vehicle lanes; cordon-based schemes (Meng and Liu, 2012; Zheng et al., 2012, 2016); area-based schemes (Simoni et al., 2015; Ye et al., 2015); and distance-based schemes (Meng et al., 2012; Daganzo and Lehe, 2015). Distance-based schemes improve fairness and efficiency by charging based on distance, thus capturing the congestion externality that a vehicle imposes more accurately than cordon and area-based schemes (Lentzakis et al., 2020, 2023). Accordingly, this study focuses on distance-based time-of-day pricing (TODP) schemes.

The design of second-best pricing schemes is typically formulated as non-linear programming and bi-level optimization problems (see for example Zhang and Yang (2004), Meng et al. (2004)). These problems are in general non-convex and computationally difficult to solve, and hence, researchers typically resort to meta-heuristics, heuristics, or approximations (Zhang and Yang, 2004; Ekström et al., 2012; Verhoef, 2002; Gupta et al., 2020). Trial-and-error methods also have been applied. For instance, Meng et al. (2005) proposed a method to determine effective link tolls without knowledge of demand and link travel time functions. Yang et al. (2010) further considered flow interactions, while Meng and Liu (2011) considered probabilistic route choices using a stochastic user equilibrium formulation. Other extensions of the trial-and-error method related to second-best pricing were proposed in Wang and Yang (2012), Xu et al. (2013), Zhou et al. (2015).

Congestion pricing has also been studied in the context of day-to-day dynamics and several studies adopt a continuous time approach, formulating the toll design problem as an optimal control problem (Friesz et al. (2004), Wie and Tobin (1998); see also (Sandholm, 2002) for an application of evolutionary game theory). In contrast, Rambha and Boyles (2016) propose a dynamic day-to-day pricing mechanism that computes the optimal link tolls to reduce the expected total system travel time in a discrete time setting. Yang et al. (2007) derived dynamic marginal cost tolls based on the steepest descent direction of the system cost and proved convergence to the system optimum. Trial and error schemes of the type discussed earlier have also been proposed within the day-to-day dynamic setting (Ye et al., 2015; Guo et al., 2016). More recently, Liu et al. (2017) proposed a robust optimization framework (solved using an artificial bee colony algorithm) for the determination of a day-to-day distance-based toll considering network performance on each day.

In contrast with the above approaches which largely adopt analytical models of network congestion, when utilizing simulation-based models, solution of the toll design problem requires computationally intensive objective function evaluations that rely on a simulator. Given these computational constraints, recently, surrogate-based optimization approaches have been applied as they can approximate the relationship between decision variables and corresponding objective functions values using a small number of function evaluations (Osorio and Bierlaire, 2013). For example, Chen et al. (2014, 2016) used the Kriging model for link-based toll optimization, where the objective functions are assessed through a stochastic simulator for a freeway network. Gu et al. (2019) applied a surrogate-based optimization with expected improvement sampling to a time- and distance-based toll to reduce travel time and heterogeneity in the distribution of congestion. Liu et al. (2021) proposed a Bayesian optimization method for the time-of-day distance-based pricing optimization problem and demonstrated its effectiveness numerically. Zhong et al. (2021) also adopted the Bayesian optimization framework to design an area-based flat toll considering long-term land use effects. Furthermore, surrogate-based optimization methods have been widely applied to various traffic and signal control problems (Chong and Osorio, 2018), transit scheduling (Zhang et al., 2017), and day-to-day model calibration (Cheng et al., 2019).

Although promising, surrogate-based methods still pose computational challenges in the context of day-to-day dynamic models, where each evaluation of the objective function requires simulating the day-to-day model until stationarity or convergence. This paper contributes to the literature by proposing a Contextual Bayesian Optimization approach that circumvents this issue by embedding and integrating the optimization procedure within the day-to-day dynamic model. This is done by incorporating temporal contextual information (the ‘day’) into the Gaussian process (GP) underlying the BO (in other words the surrogate model maps the objective function value to the toll design variables and ‘day’). Thus, we implicitly model the day-to-day dynamic evolution of flows and perform the optimization for each ‘day’, while transferring information of the GP across days. The proposed contextual BO approach is applied to the design of distance-based time-of-day pricing (TODP) schemes using a day-to-day modeling framework

consisting of a logit-mixture departure time choice model on the demand side and a simulation-based trip-based macroscopic fundamental diagram (MFD) model (Lamotte and Geroliminis, 2018) on the supply side. We demonstrate that the scheme efficiently evolves to a near system-optimal state while achieving a ten-fold reduction in the number of single-day simulations. Further the scheme is extended to incorporate context specific demand and supply information, which can be of value to policy-makers when evaluating optimal toll design schemes under a wide range of scenarios in a computational tractable manner. Finally, from a policy perspective, we find that the distance-based schemes yield significant welfare gains relative to area-based schemes and show that the design of the distance-based tariff scheme can significantly affect distributional impacts. A suitably designed two-part tariff structure can partially offset the relatively large welfare losses of travelers with longer commute distances while maintaining overall welfare gains.

The rest of this paper is organized as follows. Section 2 presents the traffic flow model and day-to-day dynamic model. Section 3 introduces the basic components of BO, and then proposes a contextual BO-based approach to integrate the optimization of the TODP tariffs within the day-to-day model, as well as two alternative BO-based benchmark optimization strategies. Section 4 consists of numerical experiments on a reservoir network, which are focused on examining the convergence of the proposed framework, effectiveness of the BO-based approach and the distance-based TODP scheme, potential of transferability that the approach provides, and distributional effects of different tariff schemes. Finally, Section 5 summarizes our findings and discusses possible avenues for future research.

2. Background and modeling context

We consider a morning commute problem where a fixed demand of N travelers (indexed by $i = 1 \dots N$) wish to travel during the morning peak. Travelers are assumed to choose their departure time intervals based on travel time, schedule delay (the difference between the actual arrival time and desired arrival time), and toll cost, in line with the classic Vickrey model (Vickrey, 1969). The trip-based MFD model, which considers individual traveler attributes, serves as the supply model. A day-to-day modeling framework (Cantarella and Cascetta, 1995), where perceived travel times and schedule delay of travelers evolve from day to day through a learning process, determines the evolution of traffic flows. Finally, under the proposed distance-based TODP scheme, which is used to manage peak-period congestion and achieve peak spreading, traveler i 's toll cost depends on a time-of-day pricing scheme denoted by $g(t|\theta)$ and the trip length L_i . θ represents the parameters that define the time-dependent tariff structure.

In this section, we present (1) the supply model based on the trip-based MFD, (2) the demand and day-to-day dynamic model, and (3) toll design optimization formulation.

2.1. Supply model

As defined and investigated in Fosgerau (2015), Daganzo and Lehe (2015), Mariotte et al. (2017), *trip-based MFD* models treat each traveler individually and hence, allow for heterogeneity in terms of trip length, value of time (VOT) and schedule delay penalties, thus providing a more detailed description of traffic dynamics. The general principle of the *trip-based MFD* is that the trip length of traveler i can be computed as the integral of the speed from the entering time t_i^{dep} to the exiting time $t_i^{dep} + T_i(t_i^{dep})$, which is written as follows,

$$L_i = \int_{t_i^{dep}}^{t_i^{dep} + T_i(t_i^{dep})} V(n(t)) dt \quad (1)$$

The speed $V(n(t))$ in the network is computed by the aggregated speed-MFD (Daganzo, 2007; Geroliminis et al., 2007), which is assumed to be the same for all travelers in the network and only changes with an event (departure or arrival). Hence, Eq. (1) can be decomposed into p parts (where p is the number of arrivals and departures occurring while vehicle i is on the network) and computed within an event-based simulation framework. This iterative updating is executed until each traveler finishes her/his trip. Note that this event-based simulation provides more accuracy and realism but makes the model computationally demanding.

2.2. Demand model and day-to-day dynamic learning

The money-metric utility for a traveler i departing in a time interval $t \in TW_i$ (TW_i is the set of feasible departure time intervals for individual i) on day d is given by:

$$U_{i,d}(t) = C_{i,d}(t) + \epsilon_i, \quad (2)$$

where ϵ_i is an identically and independently distributed error term; and $C_{i,d}(t)$ is the systematic utility (defined in Eq. (4)) for traveler i departing in time interval t on day d . The probability of choosing departure time interval t is calculated using the logit model (Ben-Akiva et al., 1985) as follows:

$$Pr_{i,d}(C_{i,d}(t)) = \frac{\exp(\mu \cdot C_{i,d}(t))}{\sum_{s \in TW_i} \exp(\mu \cdot C_{i,d}(s))}, \quad (3)$$

where $\mu > 0$ is the scale parameter, which determines the variance of the unobserved utility.

Let $\tilde{c}_{i,d}(t)$ denote the perceived utility associated with the time components (travel time, schedule delay early, schedule delay late) for traveler i on day d departing in time interval t ; the systematic utility $C_{i,d}(t)$ is defined as:

$$C_{i,d}(t) = \tilde{c}_{i,d}(t) - g(t|\theta) \cdot L_i \cdot w, \quad (4)$$

where the tariff payment of traveler i is the product of tariff $g(t|\theta)$, trip length L_i and a scaling factor w .

Within the day-to-day modeling framework, at the end of each day d , travelers update their perception of the time components for day $d + 1$, $\tilde{c}_{i,d+1}(t)$, combining the initially perceived $\tilde{c}_{i,d}(t)$ on day d with the experienced (chosen alternatives) and estimated (unchosen alternatives) travel time, schedule delays on day d , $c_{i,d}(t)$ as follows:

$$\tilde{c}_{i,d+1}(t) = \omega \cdot \tilde{c}_{i,d}(t) + (1 - \omega) \cdot c_{i,d}(t), \quad (5)$$

where $0 < \omega < 1$ is a learning parameter, which represents the relative weight given to historical experience (previous perceived time components) versus current experience (Horowitz, 1984).

The experienced (or estimated) money-metric utility associated with time components for traveler i on day d departing in time interval t is given by

$$\begin{aligned} c_{i,d}(t) = & -\alpha_i \cdot T_{i,d}(t) - \delta_i \cdot \beta_i \cdot (T_i^* - t - T_{i,d}(t)) \\ & - (1 - \delta_i) \cdot \gamma_i \cdot (t + T_{i,d}(t) - T_i^*), \end{aligned} \quad (6)$$

where α_i is the VOT for traveler i , $T_{i,d}(t)$ is the travel time for traveler i on day d departing in time interval t , β_i and γ_i are the schedule delay penalty parameters for early and late arrival for traveler i , δ_i is a binary variable that equals 1 if traveler i arrives early and 0 otherwise, and T_i^* is the preferred arrival time of traveler i .

Moreover, to estimate the travel time for all unchosen departure time intervals in the choice set TW_i , we use the concept of fictional travelers who are assumed to choose these departure time intervals without influencing the accumulation of the network (Lamotte and Geroliminis, 2015). The choice set of feasible departure time intervals or departure time window TW_i is individual-specific and defined as $TW_i = \{t_{i,0}^{dep} - \tau, t_{i,0}^{dep} - (\tau - 1), \dots, t_{i,0}^{dep} + \tau\}$, where τ is a parameter and $t_{i,0}^{dep}$ represents the initial departure time interval on day 0, which is computed from the preferred arrival time T_i^* and the perceived travel time on day 0. Thus, the departure time window TW_i consists of 2τ time intervals centered around the preferred departure time on day 0, $t_{i,0}^{dep}$.

Note that the day-to-day model described above is a stochastic process if the underlying day-to-day stochasticity in departure flows is explicitly treated (Cantarella and Cascetta, 1995; Watling, 1996). Due to the heterogeneity in trip lengths, which makes it difficult to derive travel times as a function of aggregate flows (Lamotte and Geroliminis, 2018), the stochastic process is solved using simulation, in which we assume that the error terms for a given individual are perfectly correlated across days. For more on the convergence and stationarity properties of this day-to-day dynamic model, we refer the reader to Liu et al. (2022).

2.3. Toll optimization

The social welfare per capita W computed at the equilibrium state (convergence of the day-to-day model) is used to measure the performance of scenarios with and without the distance-based TODP, and we wish to determine the time-dependent tariff structure that maximizes W . First, in the no toll case (or NTE, the scenario without distance-based TODP), the social welfare per capita W_N^d is the average consumer surplus (CS) per traveler, i.e., the average of observed travel utilities $U_{i,d}(t_{i,d}^{dep})$, including travel time cost, schedule delay penalty and the random component, given by

$$\begin{aligned} W_N^d &= \frac{CS}{N} = \frac{1}{N} \sum_{i=1}^N U_{i,d}(t_{i,d}^{dep}) \\ &= \frac{1}{N} \sum_{i=1}^N [c_{i,d}(t_{i,d}^{dep}) + \epsilon_i(t_{i,d}^{dep})] \\ &= \frac{1}{N} \sum_{i=1}^N U'_{i,d}(t_{i,d}^{dep}), \end{aligned} \quad (7)$$

where d is the day when the system converges to the equilibrium.

For the TODP scenario, the social welfare is computed as the average of the CS and the regulator revenue (RR) from toll collection where $RR = \sum_{i=1}^N g(t_{i,d}^{dep}|\theta) \cdot L_i \cdot w$. The social welfare per capita W_P^d is calculated as follows:

$$\begin{aligned} W_P^d &= \frac{1}{N} [CS + RR] \\ &= \frac{1}{N} \sum_{i=1}^N [c_{i,d}(t_{i,d}^{dep}) - g(t_{i,d}^{dep}|\theta) \cdot L_i \cdot w + \epsilon_i(t_{i,d}^{dep})] \\ &\quad + \frac{1}{N} \left[\sum_{i=1}^N g(t_{i,d}^{dep}|\theta) \cdot L_i \cdot w \right] \\ &= \frac{1}{N} \sum_{i=1}^N U'_{i,d}(t_{i,d}^{dep}). \end{aligned} \quad (8)$$

Note that the TODP welfare measure is equivalent to the NTE case and equals the sum of travel time cost, schedule delay penalty, and the random utility component.

We can thus formulate the simulation-based toll optimization problem with the objective of maximizing social welfare W_p as follows:

$$\begin{aligned} \max_{\theta} \quad & W_p^d \\ \text{s.t.} \quad & W_p^d = SM(g(t|\theta), \Phi, \Psi) \\ & g(t|\theta) \geq 0 \\ & t \in TW, \end{aligned} \quad (9)$$

where $g(\cdot)$ is the TODP toll function and θ are the associated parameters. In the case of distance-based scheme, $g(\cdot)$ is in the unit of [DKK/meter], while in the case of standard area-based scheme, it is in the unit of [DKK]. $\Phi = \{N, \alpha, \beta, \gamma, T^*, TW, L\}$ contains all input data: number of travelers, value of time, value of schedule delays, preferred arrival time, departure time window and trip length. $\Psi = \{n_{jam}, v_f, V(n), \omega\}$ represents model parameters: network capacity, free flow speed, speed function and learning rate. Function $SM(\cdot)$ is the system model described in Sections 2.1 and 2.2.

3. Methodology

In this section, we first provide an introduction to Bayesian optimization. Next we propose three different Bayesian optimization approaches to solve the toll optimization problem (9). This includes the novel contextual Bayesian optimization approach designed specifically for the day-to-day model, and two alternative benchmark BO-based optimization strategies.

3.1. Bayesian optimization

A BO framework essentially consists of two main steps (Frazier, 2018): (1) update a Bayesian statistical model that approximates a complex map from the decision variables (i.e., the TODP scheme parameters θ ; also called inputs or input variables in the remainder of this section) to the output or objective function value (i.e., the social welfare W_p); (2) determine the next trial decision variable by optimizing an acquisition function (in other words, select the next vector of candidate TODP parameters to observe or evaluate). These two steps are discussed in detail in Sections 3.1.1 and 3.1.3, respectively.

3.1.1. Gaussian process

The BO framework typically relies on a Gaussian Process (GP), which assumes the objective function values and input variables are jointly distributed. Let the input variables be denoted by $\mathbf{x} = \theta$; then the GP is fully specified by its mean function $\mu(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$ as follows:

$$W(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (10)$$

Here W stands for the social welfare per capita under a TODP scheme, dropping the subscript P for notational convenience.

As is standard practice, we assume the mean function $\mu(\cdot) = 0$ (Williams and Rasmussen, 2006). Further, assume we have evaluated m different TODP scheme parameters $\mathbf{x}$ according to a space-filling experimental design given by $D_m = \{\mathbf{x}_{1:m}, W_{1:m}\}$, where $\mathbf{x}_{1:m} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m]^T$ are the input points and $W_{1:m} = [W_1, W_2, \dots, W_m]^T$ are the corresponding objective values. Then, the joint distribution of $W_{1:m}$ and a new candidate input point $\mathbf{x}_{m+1}$, W_{m+1} , is given by:

$$\begin{bmatrix} W_{1:m} \\ W_{m+1} \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K} & \mathbf{k} \\ \mathbf{k}^T & k(\mathbf{x}_{m+1}, \mathbf{x}_{m+1}) \end{bmatrix}\right), \quad (11)$$

where $\mathbf{k} = [k(\mathbf{x}_{m+1}, \mathbf{x}_1), k(\mathbf{x}_{m+1}, \mathbf{x}_2), \dots, k(\mathbf{x}_{m+1}, \mathbf{x}_m)]^T$, and $\mathbf{K}$ is the covariance matrix with entries $K_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$ for $i, j \in \{1, 2, \dots, m\}$.

The posterior distribution of W_{m+1} can be computed using Bayes' theorem,

$$W_{m+1} | W_{1:m} \sim \mathcal{N}(\mu(\mathbf{x}_{m+1}), \sigma^2(\mathbf{x}_{m+1})), \quad (12)$$

where $\mu(\mathbf{x}_{m+1}) = \mathbf{k}^T \mathbf{K}^{-1} W_{1:m}$ and $\sigma^2(\mathbf{x}_{m+1}) = k(\mathbf{x}_{m+1}, \mathbf{x}_{m+1}) - \mathbf{k}^T \mathbf{K}^{-1} \mathbf{k}$. For the covariance function, we choose the commonly used Matern kernel (Matérn, 2013),

$$k(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} (\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|)^\nu H_\nu(\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|), \quad (13)$$

where $\Gamma(\cdot)$ is the Gamma function and H_ν is the modified Bessel function. In the numerical experiments, we assume $\nu = 5/2$.

3.1.2. Contextual Gaussian process

To take into account environmental variables (termed the 'context') in addition to the action points (TODP scheme parameters in our case), Krause and Ong (2011) proposed the contextual GP (CGP), which allows for fitting the GP over the action-context space.




Fig. 1. Flowchart of Bayesian optimization for equilibrium social welfare.

Under the day-to-day framework, we consider time, i.e., the index of day, as the context. Therefore, the input point $\mathbf{x}$ is augmented as $\hat{\mathbf{x}} = (\mathbf{x}, d)$, associated with the CGP as $W(\hat{\mathbf{x}}) \sim \text{CGP}(\hat{\mu}(\hat{\mathbf{x}}), \hat{k}(\hat{\mathbf{x}}, \hat{\mathbf{x}}'))$. Note that the chosen Matern kernel has the property that the product of two Matern kernels with smoothness parameters ν is still a Matern kernel with smoothness parameters ν (Krause and Ong, 2011). We use the product form for the composite kernel $\hat{k}$, which is written as $\hat{k}\{(\mathbf{x}, d), (\mathbf{x}', d')\} = k(\mathbf{x}, \mathbf{x}') \cdot k_d(d, d')$.

3.1.3. Acquisition function

After updating the posterior distribution over W , an acquisition function is constructed that measures the value of candidate input points by using the inferred objective function value and variance of the prediction. The upper confidence bound (UCB) (Srinivas et al., 2009) is a commonly used acquisition function, written as follows:

$$\eta_{UCB}(\mathbf{x}; \rho) = \mu(\mathbf{x}) + \rho \sigma(\mathbf{x}), \quad (14)$$

where ρ is a hyperparameter which controls balance between exploration and exploitation, such that a larger ρ will lead to more exploration. The next point to be evaluated can be determined by maximizing the UCB function (14),

$$\mathbf{x}_{m+1} = \arg \max_{\mathbf{x}} \eta_{UCB}(\mathbf{x}). \quad (15)$$

We can derive the generalized contextual UCB function by using the augmented input points $\hat{\mathbf{x}}$ and mean function and standard deviation of the associated CGP.

3.1.4. Experiment design

When applying the BO algorithm, the initial set of input points is usually randomly generated from the input space. This process may influence solution quality and algorithm efficiency. Thus, a space-filling experimental design is useful for providing a good set of initial input points.

We apply one of the most popular sampling methods, Latin Hypercube Sampling (LHS) (McKay et al., 2000), to generate the initial set of sample points. LHS stratifies each variable of $\mathbf{x}$ into m equal intervals, and draws sample points from each sub-interval uniformly. Compared to the Monte Carlo method, LHS has the advantage that the sampled points are independent without overlap, which is representative of the real variability.

3.2. TODP optimization strategies based on BO

In this subsection, we introduce three BO-based optimization strategies for the proposed area- and distance-based TODP scheme: solve optimization problem (9) for (1) the equilibrium state, (2) each day independently, and (3) each day considering temporal contextual information and the day-to-day evolution of departure flows.

3.2.1. Optimize the social welfare at equilibrium

As shown in Fig. 1, the first toll optimization strategy (denoted by S_1) is to find a toll that maximizes the social welfare at equilibrium. This is analogous to the standard black-box optimization and surrogate-based optimization techniques. A single GP is constructed using the sample points consisting of the toll parameters and social welfare. Under this strategy, the evaluation of each input tolling scheme requires a complete day-to-day simulation wherein the value of the objective function is computed after equilibrium is reached. Thus, the training of the GP essentially treats the day-to-day simulation as a black box and calls for a




Fig. 2. Flowchart of Bayesian optimization of daily social welfare.

large number of ‘within-day’ (or single-day) simulations. Specifically, for S_1 , let M_1 represent the number of days needed for the day-to-day process to converge, Q_1 represent the number of iterations for BO, and $\|D_1\|$ be the number of initial sample points, then the total number of within-day simulations, denoted by ζ , is given by $\zeta_1 = (\|D_1\| + Q_1) \cdot M_1$.

3.2.2. Optimize the social welfare at each day

As shown in Fig. 2, the second toll optimization strategy (denoted by S_2) aims to obtain the optimal tolls for each day within the day-to-day process independently. The evaluation of a candidate toll scheme uses the social welfare computed from the current within-day simulation, and the BO algorithm is applied to the optimization of the tolls for the ‘current day’ only. At an arbitrary intermediate ‘day’, the current perceived travel costs are not (necessarily) equal to realized costs, or in other words, the system is not yet in equilibrium. The hypothesis underlying this approach is that despite performing the toll optimization independently for each day, the day-to-day dynamic process (with the tolls in effect) eventually converges to an equilibrium. If convergence is not significantly worse than the day-to-day model with a fixed toll scheme, and if the number of iterations for the BO algorithm for each day is significantly smaller than that for S_1 , computational efficiency is likely to be improved. Observe also that the evaluation of a single candidate toll scheme now involves only a single day simulation rather than the simulation of the entire day-to-day model until convergence. In Section 4.2, we investigate numerically whether this approach yields the same optimal toll scheme and equilibrium state as approach S_1 , and whether it does so using a fewer number of single-day simulations.

Note that this approach may be viewed from two perspectives. First, one may treat it as simply a computational method to determine an optimal toll scheme that is not day-to-day adaptive. The regulator simply implements the optimal within-day dynamic toll scheme, which does not vary across days. The second is to in fact view the approach as a form of day-to-day dynamic pricing, where the toll schemes computed on each day are implemented in practice. Such schemes have been widely studied in the literature (Friesz et al., 2004; Rambha and Boyles, 2016) and afford an attractive approach to pricing that can accommodate day-to-day variability and the likely absence of equilibria in the real-world. However, they are also more difficult to implement in practice (in Singapore for example tolls are only revised once every three months or so) due to the complexity of communicating a complex tariff structure to travelers every day. For the remainder of the paper, we focus on the first perspective, and defer the study of day-to-day dynamic pricing to future research.

In this approach, multiple GPs are constructed, one for each day d , using the toll parameters and social welfare on the given day. Each GP is trained independently, and therefore, the initial sample points are generated separately. If $Q_{2,d}$ denote the number of BO iterations needed on day d , we have $\zeta_2 = \|D_2\| \cdot M_2 + \sum_{d=1}^{M_2} Q_{2,d}$.

3.2.3. Optimize the social welfare at each day utilizing context

The third solution approach (denoted by S_3) is the novel Contextual BO, shown in Fig. 3, where once again, the optimization of the toll scheme is performed separately for each day within the day-to-day process using the BO algorithm. However, in contrast with approach S_2 where independent GPs are trained for each day d , here, a single underlying CGP is trained that uses the ‘day’ as an additional ‘context’ variable. Moreover, sample points are continuously augmented as the day-to-day process evolves. Compared to S_1 and S_2 , the constructed single CGP of S_3 implicitly incorporates information of the day-to-day dynamic evolution of flows through the temporal context. In other words, S_3 takes the relationship between objective functions from day to day into account and uses past observations as weak priors for the current day’s optimization problem (9), resulting in a more accurate prediction (this weak effect is enforced through the covariance structure, i.e., $k_d(d, d')$ will be larger when d and d' are farther apart, which makes $\hat{k}((x, d), (x', d'))$ larger as we use the product kernel). This can further enhance computational efficiency relative to S_2 since fewer function evaluations are required for each application of the BO algorithm.

Specifically, in approach S_3 , initial sample points are first generated on day 1. Then, at each day, given the fixed context (day), the BO algorithm is applied to optimize the toll parameters iteratively. The trial toll schemes are generated along with the GP’s




Fig. 3. Flowchart of contextual Bayesian optimization for daily social welfare.

Table 1

Numerical settings.

Parameters [unit]	Specification
Demand [traveler]	$N_1 = 3700$, $N_2 = 4500$
Initial departure time [min]	$t_{i,0}^{dep} = \mathcal{N}(80, 18)$, $t_{i,0}^{dep} \in [20, 150]$
Trip length [m]	$L_i = 4600 + \mathcal{N}(0, (0.2 \times 4600)^2)$, $L_i > 0$
Trip length scale factor	$w = 2 \times 10^{-4}$
Schedule delay penalty [DKK*/min]	$\beta_i \sim \text{Lognormal}(-1.9, 0.2^2) \times 4$
Value of time [DKK/min]	$\gamma_i = \beta_i \times e^1$
Time window parameter	$\alpha_i = \beta_i \times e^{0.5}$
Time interval [min]	$\tau = 90$
Time interval [min]	$\Delta t = 1$
Network capacity [vehicle]	$n_{jam} = 4500$
Free flow speed [m/s]	$v_f = 9.78$
Speed function [m/s]	$V(n) = v_f(1 - \frac{n}{n_{jam}})^2$
Learning parameter	$\omega = 0.9$
Tariff profile function	$g(t A, \xi, \sigma) = A \times e^{-\frac{(t-\xi)^2}{2\sigma^2}}$

* Danish krone.

predictive mean and variance, and these toll parameters and the current ‘day’ are augmented to the data, which is subsequently passed on to the next day. At the same time, the optimal toll is used for the current within-day simulation, and the corresponding perceived disutility is updated for each traveler within the day-to-day process. For S_3 , we thus have, $\zeta_3 = \|D_3\| + \sum_{d=1}^{M_3} Q_{3,d}$.

The contextual BO (approach S_3) is also appealing since demand and supply information can be added as context variables. Thus, from a regulator’s standpoint, different optimal tolling schemes can be computed efficiently by utilizing the same underlying CGP across a range of demand and supply scenarios. The addition of demand and supply information as contextual variables is also appealing in the context of day-to-day dynamic pricing, where tolls can be efficiently optimized conditional on revealed information about current or anticipated demand and supply conditions.

4. Numerical experiments

This section begins by introducing the design of the simulation experiments. Next, we present the results of (1) the comparative performance of the three BO-based optimization strategies (Section 4.2); (2) the performance of the contextual BO using contextual demand and supply information for optimization of toll schemes under varying travel demand and network capacity (Section 4.3); and (3) comparison of distance-based and area-based schemes and distributional effects (Section 4.4).

4.1. Experiment settings

The key simulation inputs and other parameters are presented in Table 1.

The experiments consider a single-reservoir network with a capacity of $n_{jam} = 4500$ travelers, with the speed function adopted from Lamotte and Geroliminis (2018) and other parameters (trip length, time window parameters, and time interval) used in Yildirimoglu and Ramezani (2020). Two demand scenarios, moderate congestion ($N_1 = 3700$ travelers) and high congestion ($N_2 = 4500$ travelers), are considered, where N_1 makes the accumulation at no tariff equilibrium just exceed the critical value of the accumulation (which equals $n_{jam}/3 = 1500$ in our settings), and N_2 is the largest possible value that will not trigger gridlock. The




Fig. 4. Day-to-day evolution in the moderate congestion scenario with optimized TODP for S_1 .

initial departure time $t_{i,0}^{dep}$ is generated from a truncated Gaussian distribution (note that hereafter $t_{i,0}^{dep}$ refers to a specific departure time rather than an interval). The desired arrival time T_i^* is then computed as $t_{i,0}^{dep} + L_i/v_f$ for all travelers, which is also normally distributed. Additionally, heterogeneous travelers are captured by drawing their trip lengths from another truncated Gaussian distribution and values of schedule delay from a lognormal distribution, respectively. The mean value of the generated value of time is 1 [DKK/min] (Fosgerau et al., 2007) and the standard deviation is 0.2 [DKK/min]. The values of schedule delay are assumed to vary proportionally to the value of time and satisfy the widely used trip timing preferences relationship, i.e., $\beta_i < \alpha_i < \gamma_i$ (Small, 2015). In both demand scenarios, the same distributions are used while all other parameters are constant (see Table 1). To set up a time-varying pricing scheme, the tariff function is assumed to take the form of a (positive) Gaussian function with three parameters, mean ξ , variance σ and amplitude A . Without loss of generality, the method described below can be extended to a Gaussian mixture function to allow for asymmetric and more flexible tariff profiles (Liu et al., 2021).

$$g(t|\theta = [\xi, \sigma, A]) = A \exp \frac{-(t - \xi)^2}{2\sigma^2} \quad (16)$$

Yang et al. (2007) refers to a day-to-day dynamic pricing scheme as strong if it improves the social welfare monotonically during the evolution process. This property guides the day-to-day process to the SO and provides a clear stopping criterion for the proposed optimization strategies S_2 and S_3 . However, due to the uncertainties in our simulation, the monotonic improvement of social welfare is not guaranteed. Instead, we let the BO terminate if the optimized social welfare on day d , $W_{p,d}$, is greater than $\max\{W_{p,d}, \dots, W_{p,d-4}\} - \hat{\sigma}$, where $\hat{\sigma}$ is the standard deviation of $W_N^{d,*}$, which is the equilibrium social welfare of the no tariff case. Therefore, the number of BO iterations at each day for S_2 and S_3 are adaptively determined.

In addition, for S_1 , we generate 30 initial sample points via LHS and set the maximum number of iterations as 50 due to the limited computational budget; for S_2 , we generate 10 initial sample points at each day; while for S_3 , we only generate 10 initial sample points at the first day of implementing pricing. Moreover, we compute a benchmark solution using the differential evolution algorithm (DEA), for which, the population size is 20, and the number of iterations is 100. For all experiments, the simulated number of days is 50.

4.2. Comparative performance of the BO-based optimization strategies

First, we examine the convergence of the day-to-day dynamic system under the optimized tariff profiles obtained from the developed BO-based methods for both demand scenarios. Theoretically, when the day-to-day evolution reaches an equilibrium, the vector of the perceived travel cost of all travelers, c_d , should be equal to the vector of the experienced travel cost of all travelers, $\bar{c}_d$ (recall also that we assume that the individual specific error terms are fixed across days and not redrawn). Thus, the inconsistency between c_d and $\bar{c}_d$ is used as a measure of convergence towards the equilibrium. Specifically, the normalized L1 norm of the difference between them, $\|(c_d - \bar{c}_d)/c_d\|_1 \times 100\%$, is computed to represent the inconsistency. It is worth noting that the equilibrium states of the base cases are used as the starting states (i.e., day 0) of the TODP cases. Fig. 4 presents the evolution processes of the social welfare gains per capita, the departure rates for every 5-minute interval (referred to as flow pattern hereafter), and the states of accumulation on different days for the moderate congestion scenario with optimized TODP by S_1 . We find that the social welfare gains per capita stabilises after 40 days with an inconsistency gap of 0.01%, and the flow pattern and accumulation on day 40 overlap with those of day 49. These observations imply that the day-to-day evolution reaches an equilibrium. We also observe similar patterns in the high congestion scenario, although the plots are omitted here.

Similarly, Fig. 5 shows the evolution processes of the three aforementioned performance indicators for the moderate congestion scenario with day-to-day dynamic optimized TODP for S_3 . We find that an equilibrium state is also reached with an inconsistency gap of 0.02%.

To provide a straightforward comparison among the three proposed BO-based optimization strategies S_1 , S_2 , and S_3 , we plot the evolution profiles of social welfare gains and the profiles of flow patterns and accumulations at the equilibrium states together in Fig. 6. Note that the time-dependent toll rates are fixed across days under S_1 ; we therefore use a horizontal line to represent the optimal welfare gains in Fig. 6(a). It can be observed that all three strategies attain a near identical social welfare and the system converges to the same equilibrium.




Fig. 5. The evolution process in the moderate congestion scenario with optimized day-to-day dynamic TODP for S_3 .




Fig. 6. The evolution process in the moderate congestion scenario with optimized TODPs for S_1 , S_2 and S_3 .

Table 2

Average performance gains under TODP optimized by S_1 , S_2 , S_3 and DEA.

Unit: [DKK/cap]	Social welfare	Travel time cost	Schedule delay cost	Tariff payment	ζ
$N_1 = 3700$					
DEA	3.20	5.74	2.86	10.96	10^5
S_1	3.20	5.61	2.72	10.71	4000
S_2	3.20	5.78	2.90	10.97	3323
S_3	3.20	5.71	2.82	10.77	337
$N_2 = 4500$					
S_1	5.95	8.89	3.08	13.27	4000
S_3	5.95	8.96	3.16	12.60	278

Next, we investigate the effectiveness of optimization strategies S_1 , S_2 , and S_3 in terms of optimality, computational efficiency, welfare and network performance. Table 2 lists monetary gains of the various performance measures at equilibrium for both demand scenarios under time-dependent distance-based tolls optimized by the different approaches. In the moderate congestion scenario, it can be seen that all the optimization approaches S_1 , S_2 , and S_3 yield an optimal social welfare that is identical to the benchmark DEA algorithm. In other words, the optimization of the tolls separately for each day does result in the system converging to the system optimum for both S_2 and S_3 . All the BO-based approaches require far fewer single-day simulations than the meta-heuristic method, and moreover, S_3 requires only 1/10th the number of single-day simulation runs ζ compared to S_1 and S_2 . This clearly demonstrates the benefit of embedding temporal contextual information. It is worth noting that under S_1 the welfare gain reaches 3.19 DKK per capita after $\zeta = 3350$ single-day simulations and 3.20 after $\zeta = 3950$. Note also that the improvement of S_2 over S_1 is marginal.

In the high congestion scenario, we see similar results in that S_1 and S_3 have a near identical solution and S_3 performs significantly better than S_1 in terms of computational efficiency (once again, a more than ten-fold reduction in ζ). Note that S_1 obtains the best solution at $\zeta = 3650$. Taking the moderate congestion scenario as an example, in Fig. 5, we can observe the change in departure flows and flattening of the accumulation due to the tolls, leading to an improvement in the social welfare per capita by 3.20 [DKK].

4.3. Contextual demand and supply information

The usefulness of the Contextual BO approach is that in principle, a range of different contextual variables can be incorporated into the CGP model. In this section, we consider two additional contextual variables, the total travel demand and capacity. To assess the capabilities of the contextual BO method to effectively 'transfer knowledge' of the underlying mapping between tolls and




Fig. 7. Day-to-day evolution for the scenario with moderate demand, lower capacity, and optimized TODPs by S_3 with and without transferred knowledge.




Fig. 8. Day-to-day evolution for the high congestion scenario under optimized TODPs by S_3 with and without transferred knowledge.

welfare across different demand and supply scenarios, we investigate how well a GP that is trained in one scenario can facilitate the optimization of toll schemes for other scenarios. Specifically, we select the pre-trained GP from the moderate congestion scenario, denoted as $\mathcal{GP}_{N_1}^{n_{jam}}$, and use it as the starting GP for (1) a scenario with the same demand but a lower network capacity $n'_{jam} = 4000$, wherein the resulting GP is denoted as $\mathcal{GP}_{N_1}^{n_{jam}, n'_{jam}}$, and (2) the high demand scenario. We further adopt $\mathcal{GP}_{N_1}^{n_{jam}, n'_{jam}}$ as the starting GP for the high demand scenario. In these tests, extra contextual information (the capacity/demand) is embedded into the pre-trained GP and the current optimization process.

First, we examine the performance of this method by comparing the equilibria reached via S_3 with and without transferred information from the GP for both scenarios, which are presented in Figs. 7 and 8, respectively. The social welfare gains shows different evolution processes but finally converge to the same solutions, where the inconsistency gaps are both smaller than 0.03%. In addition, the accumulation and flow pattern profiles from the two optimization strategies are nearly identical in both scenarios. Thus, the contextual BO with a pre-trained GP is effective.

Table 3 further summarizes the performance when utilizing the pre-trained GP for the two designed scenarios. The results show that the contextual BO with transferred knowledge yields an identical solution with an interesting superiority in computational efficiency over the method that learns from scratch. In particular, the number of iterations needed for this approach is around 20% lower than the one with no prior experience (S_3^* versus S_3), and it is further reduced when provided with more prior experience. This can provide significant computational gains when dealing with a large number of scenarios for large size networks. From a practical perspective, the experiments highlight the advantages of incorporating demand and supply contextual information and a pre-trained GP. Practical applications include cases where a regulator would like to determine optimal pricing schemes under a range of different non-recurring and recurring demand and supply scenarios including special events, weather, construction, incidents, etc.


4.3.1. Performance under demand uncertainty

In this section, we assess the capability of the contextual BO method considering demand uncertainty during the day-to-day process. Specifically, we assume there are 5000 travelers in total and generate individual characteristics (value of time, value of schedule delay etc.) as before based on the distributions from Table 1. The demand N_d on a given day d is now assumed to be uniformly distributed between [4000, 5000]. It is worth noting that we maintain the perception of the time components $\tilde{c}_{i,d}(t)$ for all travelers at the individual level using Eq. (5). For travelers who do not travel on day d , the travel time for all departure time intervals is estimated using the 'fictional user' method (Lamotte and Geroliminis, 2015).

Note that both 'day' and 'demand' are treated as contextual variables when applying S_3 . In this case, however, we do not determine the number of BO iterations on each day adaptively, and instead terminate the BO when the maximum number of iterations (30 in this experiment) is reached on each day. To compare the base case and the optimized case, we first run the base case until convergence, indicated by the normalized $L1$ norm, which reaches 5%. Starting from this state, we then run the base

Table 3Average performance gains under TODP optimized by S_3 with and without transferred knowledge.

Unit: [DKK/cap]	Social welfare	Travel time cost	Schedule delay cost	Tariff payment	ζ
$N_1 = 3700, n_{jam} = 4000$					
S_3	5.11	8.15	3.21	11.94	329
$\hat{S}_3^*$	5.13	8.12	3.17	11.67	267
$N_2 = 4500, n_{jam} = 4500$					
S_3	5.95	8.96	3.16	12.60	278
$\hat{S}_3^*$	5.96	8.98	3.16	13.00	201
$\hat{S}_3^{**}$	5.96	8.97	3.16	12.94	168

* Using S_3 with transferred knowledge from $\mathcal{GP}_{N_1}^{n_{jam}}$.** Using S_3 with transferred knowledge from $\mathcal{GP}_{N_1}^{n_{jam}, n'_{jam}}$.**Fig. 9.** Day-to-day evolution for the variant demand scenario under optimized TODPs by S_3 .

case for another 35 days, recording the IDs of travelers traveling on each day. Next, we run the TODP case starting from the same state with the same groups of travelers who make trips on the corresponding day in the base case. Fig. 9 presents the welfare gains per capita (the red solid line) and the number of travelers (the gray dashed line) across days. As shown, there are positive welfare gains across all days under the optimized TODPs, and the gain in welfare is found to be relatively larger when there are more travelers (i.e., when the network is more congested). Note that the large magnitude of welfare gains on some days are due to the fact that the demand of 5000 travelers leads to extreme levels of congestion. The results also suggest that the proposed contextual BO method can effectively yield welfare gains (comparable with the fixed demand case) in scenarios where the day-to-day process is not convergent. We can also expect it to be more efficient than the standard BO method (S_2) from the findings in Sections 4.2 and 4.3.

4.4. Comparison of distance-based and area-based schemes and distributional effects

In this section, we compare the performance of the distance-based and area-based schemes from the standpoint of both efficiency (welfare) and equity. Vertical equity is a key concern with congestion pricing, which has long been recognized as being regressive, as it is often the case that the willingness to pay for gains in travel time is smaller for lower-income travelers (de Palma and Lindsey, 2020). In the analysis that follows, we assume a simple uniform redistribution of the regulator revenue to all travelers (see for example Kockelman and Kalmanje (2005)) and investigate travelers' benefits across different VOT (proxy for income) and trip length groups. Under the assumption that regulator revenue is uniformly redistributed, traveler's benefit (TB) is the sum of social welfare gain under the optimized tariff and the average toll revenue, i.e., $TB_i = W_{P,i} - W_{N,i} + \frac{RR}{N}$. All travelers are divided into four groups based on the quartiles of VOT and trip length, respectively.

4.4.1. Comparison between distance-based and standard area-based mechanisms

First, we examine the (trip agnostic) standard area-based (zonal) mechanism, wherein the unit of $g(t)$ is [DKK]. The box-plots of traveler benefits (grouped by VOT and trip-length) under high demand are shown in Fig. 10. Note that similar patterns are observed in the moderate demand scenario although the detailed plots are omitted.




Fig. 10. Box-plot of the individual traveler's benefit according to VOT and trip length under area-based tariff.




Fig. 11. Box-plot of the individual traveler's benefit according to VOT and trip length under distance-based tariff.

In Fig. 10, the line inside the box represents the median, and the lower and upper lines are 25th (q_1) and 75th (q_3) percentiles, respectively, the whiskers extend from the edges of box by $1.5 * (q_3 - q_1)$. It can be seen in Fig. 10(a) that most travelers are better off under the optimized congestion tariff with uniform revenue distribution, and the traveler benefits increase as the VOT increases, which is in line with intuition and the literature (see for example de Palma et al. (2018), Seshadri et al. (2022), Chen et al. (2023), Jing et al. (2023)). Note that in the absence of revenue redistribution, a significantly larger percentage of users would experience negative net benefits. In Fig. 10(b), the traveler benefits also increase along with the trip length, since travelers with longer trip length benefit more from the reduction in travel time. Next, we examine the distance-based tariff, wherein the unit of $g(t)$ is [DKK/meter].

First, we observe that the overall social welfare gain under the distance-based scheme (5.95 DKK per capita) is higher than that of the standard zonal-based scheme (4.85 DKK per capita). This is expected as the former scheme better internalizes congestion externalities by charging directly based on the contribution to congestion, or in other words, travelers pay in proportion to their benefits (Levinson, 2010).

The box-plots shown in Fig. 11 summarize distributional impacts. Again, traveler benefits increase with the VOT as expected. However, contrary to the pattern in Fig. 10(b), it is observed in Fig. 11(b) that the traveler benefits decrease as trip length increases. Though travelers with larger trip length have larger travel time reduction, they suffer from larger schedule delay penalties at the same time as they tend to depart much earlier or later so as to avoid the large tariffs. In addition, they still pay more than those with smaller trip lengths (given that the toll tariff is proportional to distance with no fixed charge). Consequently, travelers with large trip lengths have significantly lower gains in benefits. Although this is an outcome of the distance-based charge by design, it may be problematic from the standpoint of vertical equity if the structure of the urban area is such that lower income travelers reside in the suburbs and have longer commutes. While inequities of this nature can be remedied via sophisticated revenue redistribution schemes or tradable credit schemes (see Levinson (2010) and Grant-Muller and Xu (2014) for more discussion on this), they can also be partially addressed through the design of the distance-based scheme, which we examine in the next section.




Fig. 12. Box-plot of the individual traveler's benefit according to VOT and trip length under optimized $Toll_1$.

4.4.2. Two-part distance-based tariffs

Under the proposed distance-based tariff, travelers are charged proportionally to their trip lengths, resulting in gains in efficiency, but a potential loss in equity if VOT is negatively correlated with trip length (which is possible for certain city structures). Travelers with short trips might underpay relative to costs they impose, particularly if one considers emissions and environmental impacts. In this case, a two-part tariff, which is composed of a fixed toll and a distance-based toll, can help in achieving a more equitable distribution of benefits.

We propose two simple two-part designs to demonstrate how the tariff structure can offset the relatively large welfare losses of travelers with longer commute distances. In the first tariff, the fixed toll is a constant and the distance-based toll is a linear function of trip length:

$$Toll_1(t, L|\theta) = \begin{cases} g(t|\theta) \cdot L \cdot w + \lambda, & L \in [0, \hat{L}] \\ g(t|\theta) \cdot L \cdot w, & \text{otherwise,} \end{cases} \quad (17)$$

where λ (unit: DKK) is the fixed toll and $\hat{L}$ is a predetermined constant trip length such that all travelers with trip lengths smaller than $\hat{L}$ need to pay the extra toll λ . In this experiment, $\hat{L}$ is set as the median of all trip lengths.

The results indicate that the overall social welfare gain under the optimized two-part tariff is almost identical to the value before (5.95 DKK per capita) but the distribution of benefits is significantly different, as shown in Fig. 12.

In Fig. 12(a), the distributional effect in terms of VOT is in line with expectation that higher VOT groups benefit more from the congestion pricing scheme. In Fig. 12(b), we observe that travelers with trip lengths longer than $\hat{L}$ have larger traveler benefits than those who have trip lengths smaller than $\hat{L}$. This is because the fixed toll λ imposes extra costs to the latter travelers and leads to a change in the distribution of social welfare gains. In addition, within the first (or the last) two trip length groups, traveler benefits still decrease with the trip length as is expected with the distance-based toll. Note also that the median user benefits (in each trip-length quartile) is not monotonic with trip length as was the case with the area-based scheme and original proportional distance-based tariff.

Note that these results are not intended to suggest this tariff structure is necessarily superior (this requires more detailed experiments that explicitly consider emissions and environmental impacts in the objective function). Rather, they highlight the fact that the distribution of benefits with respect to trip-length can be significantly altered through the tariff structure without adversely affecting overall welfare. Moreover, due to the discontinuity with respect to trip length, $Toll_1$ is not a practically effective toll scheme. A more practical alternative would be to revise the second part to be a linear function of trip length which reduces from $\hat{\lambda}$ to 0 as trip length increases from 0 to the maximum trip length. Although such a scheme solves the discontinuity issue and is easy to implement, the results indicate that it does not significantly affect the benefits distribution, which shows a similar pattern to the original case shown in Fig. 11. In order to address this issue, we next consider a second two-part tariff where the second-part of the tariff is continuous in trip length with an exponential form:

$$Toll_2(t, L|\theta) = g(t|\theta) \cdot L \cdot w + \lambda \cdot e^{-w \cdot L} \quad (18)$$

Again, we observe a similar social welfare gain under the optimized $Toll_2$ (5.95 DKK per capita) but a different benefits distribution as shown in Fig. 13.

As expected, higher VOT groups have higher traveler benefits as demonstrated in Fig. 13(a). The slope of the negative exponential function is relatively steep when trip length is smaller than the first quartile and becomes gentle as trip length increases; in Fig. 13(b), we find that traveler benefits of the first group is smaller than the second group, and decrease as trip length increases for groups 2–4. It is worth noting that the overall social welfare distributes more evenly across the trip length groups compared to that under $Toll_1$ thanks to the continuous fixed toll.




Fig. 13. Box-plot of the individual traveler's benefit according to VOT and trip length under optimized $Toll_2$.

In summary, these experiments imply that the two-part tariff can alter the distribution of benefits and potentially improve equity (when VOT is inversely correlated with trip-length) while sustaining the social welfare gains. However, more detailed experiments are needed that examine different tariff structures considering environmental impacts, congestion and equity.

5. Conclusions

This paper proposed a contextual BO-based approach for the design of congestion pricing schemes under day-to-day flow dynamics. The approach is applied to optimize area- and distance-based time-of-day pricing schemes for social welfare maximization in the context of the morning commute problem. Specifically, the demand (logit mixture) and supply (trip-based MFD) models are operationalized using an agent-based simulation framework with heterogeneous travelers, which is associated with an expensive-to-evaluate objective function. The contextual BO-based approach, which uses a Gaussian process embedded with temporal contextual information in addition to the toll function parameters, takes the day-to-day dynamic evolution of flows into account implicitly, and therefore optimizes the tolls more efficiently.

The experimental results demonstrated the convergence of the day-to-day system under the proposed contextual-BO approach (where tolls are optimized for each day separately) to the system optimum. Importantly, the contextual BO-based approach yields a ten-fold reduction in the number of evaluations compared to benchmark approaches that do not use contextual information. Further, the scheme is extended to incorporate context specific demand and supply information. The results showed the superiority of the contextual BO with transferred knowledge in computational efficiency compared to the method that learns from scratch. Finally, from a policy perspective, we find that the distance-based schemes yield significant welfare gains relative to area-based schemes and show that the design of the distance-based tariff scheme can significantly affect distributional impacts. A suitably designed two-part tariff structure can partially offset the relatively large welfare losses of travelers with longer commute distances while maintaining overall welfare. More detailed experiments are warranted to examine the design of the distance-based, and potentially, travel time-based tariffs. Moreover, environmental and emission externalities have been ignored and should be considered in the design of the pricing schemes.

In summary, the above developments and findings provide a promising approach for the optimal design of congestion pricing schemes considering day-to-day dynamics. Applications include cases where the regulator wishes to determine optimal pricing schemes under a range of different non-recurring and recurring demand and supply scenarios including special events, weather, construction, incidents, etc. The addition of demand and supply information as contextual variables also has applications in the context of day-to-day dynamic pricing (in which tolls adapted from day to day), where tolls can be efficiently optimized conditional on revealed information about current or anticipated demand and supply conditions.

The current research can be extended in a few directions. First, to bring congestion pricing closer to practice, it is necessary to consider additional and combined choice dimensions (such as mode, route, and trip cancellation) and more detailed models that are calibrated using real-world data. Second, the design and performance of the congestion pricing schemes should be studied under day to day variability in demand and supply. Finally, more contextual variables could be incorporated into the underlying GP to enhance efficiency for more complex model structures.

CRedit authorship contribution statement

Renming Liu: Study conception and design, Data collection, Analysis and interpretation of results, Writing – original draft, Reviewed the results. **Yu Jiang:** Study conception and design, Analysis and interpretation of results, Writing – review and editing, Reviewed the results. **Ravi Seshadri:** Study conception and design, Analysis and interpretation of results, Writing – review and

editing, Reviewed the results. **Moshe Ben-Akiva**: Study conception and design, Analysis and interpretation of results, Reviewed the results. **Carlos Lima Azevedo**: Study conception and design, Analysis and interpretation of results Writing – review and editing, Reviewed the results.

Acknowledgments

This research was carried out under the NEMESYS project funded by the DTU (Technical University of Denmark)-NTU (Nanyang Technical University) Alliance. Part of this research was funded by the U.S. National Science Foundation (Award CMMI-1917891). All authors approved the final version of the manuscript.

References

- Arnott, R., De Palma, A., Lindsey, R., 1990. Economics of a bottleneck. *J. Urban Econ.* 27 (1), 111–130.
- Ben-Akiva, M.E., Lerman, S.R., Lerman, S.R., 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*, Vol. 9. MIT Press.
- Cantarella, G.E., Cascetta, E., 1995. Dynamic processes and equilibrium in transportation networks: Towards a unifying theory. *Transp. Sci.* 29 (4), 305–329.
- Chen, S., Seshadri, R., Azevedo, C.L., Akkinipally, A.P., Liu, R., Araldo, A., Jiang, Y., Ben-Akiva, M.E., 2023. Market design for tradable mobility credits. *Transp. Res. C* 151, 104121.
- Chen, X.M., Xiong, C., He, X., Zhu, Z., Zhang, L., 2016. Time-of-day vehicle mileage fees for congestion mitigation and revenue generation: A simulation-based optimization method and its real-world application. *Transp. Res. C* 63, 71–95.
- Chen, X., Zhang, L., He, X., Xiong, C., Li, Z., 2014. Surrogate-based optimization of expensive-to-evaluate objective for optimal highway toll charges in transportation network. *Comput.-Aided Civ. Infrastruct. Eng.* 29 (5), 359–381.
- Cheng, Q., Wang, S., Liu, Z., Yuan, Y., 2019. Surrogate-based simulation optimization approach for day-to-day dynamics model calibration with real data. *Transp. Res. C* 105, 422–438.
- Chong, L., Osorio, C., 2018. A simulation-based optimization algorithm for dynamic large-scale urban transportation problems. *Transp. Sci.* 52 (3), 637–656.
- Daganzo, C.F., 2007. Urban gridlock: Macroscopic modeling and mitigation approaches. *Transp. Res. B* 41 (1), 49–62.
- Daganzo, C.F., Lehe, L.J., 2015. Distance-dependent congestion pricing for downtown zones. *Transp. Res. B* 75, 89–99.
- de Palma, A., Lindsey, R., 2011. Traffic congestion pricing methodologies and technologies. *Transp. Res. C* 19 (6), 1377–1399.
- de Palma, A., Lindsey, R., 2020. Tradable permit schemes for congestible facilities with uncertain supply and demand. *Econ. Transp.* 21, 100149.
- de Palma, A., Proost, S., Seshadri, R., Ben-Akiva, M., 2018. Congestion tolling-dollars versus tokens: A comparative analysis. *Transp. Res. B* 108, 261–280.
- Eklström, J., Sumalee, A., Lo, H.K., 2012. Optimizing toll locations and levels using a mixed integer linear approximation approach. *Transp. Res. B* 46 (7), 834–854.
- Eurostat, E.U. Commission, 2018. *Energy, Transport and Environment Indicators — 2018 Edition*, Vol. 2. Office for Official Publications of the European Communities.
- Fosgerau, M., 2015. Congestion in the bathtub. *Econ. Transp.* 4 (4), 241–255.
- Fosgerau, M., Hjorth, K., Lyk-Jensen, S.V., 2007. *The Danish Value of Time Study*. The Danish Transport Research Institute.
- Frazier, P.I., 2018. A tutorial on bayesian optimization. *arXiv preprint*.
- Friesz, T.L., Bernstein, D., Kydes, N., 2004. Dynamic congestion pricing in disequilibrium. *Netw. Spat. Econ.* 4 (2), 181–202.
- Geroliminis, N., Daganzo, C.F., et al., 2007. Macroscopic modeling of traffic in cities. In: *Transportation Research Board 86th Annual Meeting*, no. 07-0413. No. 07-0413.
- Grant-Muller, S., Xu, M., 2014. The role of tradable credit schemes in road traffic congestion management. *Transp. Rev.* 34 (2), 128–149.
- Gu, Z., Liu, Z., Cheng, Q., Saberi, M., 2018. Congestion pricing practices and public acceptance: A review of evidence. *Case Stud. Transp. Policy* 6 (1), 94–101.
- Gu, Z., Waller, S.T., Saberi, M., 2019. Surrogate-based toll optimization in a large-scale heterogeneously congested network. *Comput.-Aided Civ. Infrastruct. Eng.* 34 (8), 638–653.
- Guo, R.-Y., Yang, H., Huang, H.-J., Tan, Z., 2016. Day-to-day flow dynamics and congestion control. *Transp. Sci.* 50 (3), 982–997.
- Gupta, S., Seshadri, R., Atasoy, B., Prakash, A.A., Pereira, F., Tan, G., Ben-Akiva, M., 2020. Real-time predictive control strategy optimization. *Transp. Res. Rec.* 2674 (3), 1–11.
- Horowitz, J.L., 1984. The stability of stochastic equilibrium in a two-link transportation network. *Transp. Res. B* 18 (1), 13–28.
- Jing, P., Seshadri, R., Sakai, T., Shamshirpour, A., Alho, A.R., Lentzakis, A., Ben-Akiva, M.E., 2023. Evaluating congestion pricing schemes using agent-based passenger and freight microsimulation. *arXiv preprint arXiv:2305.07318*.
- Kockelman, K.M., Kalmanje, S., 2005. Credit-based congestion pricing: A policy proposal and the public's response. *Transp. Res. A* 39 (7–9), 671–690.
- Krause, A., Ong, C.S., 2011. Contextual Gaussian process bandit optimization. In: *Nips*. pp. 2447–2455.
- Lamotte, R., Geroliminis, N., 2015. Dynamic traffic modeling: Approximating the equilibrium for peak periods in urban areas.
- Lamotte, R., Geroliminis, N., 2018. The morning commute in urban areas with heterogeneous trip lengths. *Transp. Res. B* 117, 794–810.
- Langmyhr, T., 1999. Understanding innovation: The case of road pricing. *Transp. Rev.* 19 (3), 255–271.
- Lentzakis, A.F., Seshadri, R., Akkinipally, A., Vu, V.-A., Ben-Akiva, M., 2020. Hierarchical density-based clustering methods for tolling zone definition and their impact on distance-based toll optimization. *Transp. Res. C* 118, 102685.
- Lentzakis, A.F., Seshadri, R., Ben-Akiva, M., 2023. Predictive distance-based road pricing—Designing tolling zones through unsupervised learning. *Transp. Res. A* 170, 103611.
- Levinson, D., 2010. Equity effects of road pricing: A review. *Transp. Rev.* 30 (1), 33–57.
- Li, M.Z., 1999. Estimating congestion toll by using traffic count data—Singapore's area licensing scheme. *Transp. Res. E* 35 (1), 1–10.
- Li, M.Z., 2002. The role of speed-flow relationship in congestion pricing implementation with an application to Singapore. *Transp. Res. B* 36 (8), 731–754.
- Lindsey, R., 2006. Do economists reach a conclusion? *Econ. J. Watch* 3 (2), 292–379.
- Liu, R., Chen, S., Jiang, Y., Seshadri, R., Ben-Akiva, M., Azevedo, C., 2022. Managing network congestion with a trip- and area-based tradable credit scheme. *Transportmetrica B* <http://dx.doi.org/10.1080/21680566.2022.2083034>.
- Liu, R., Jiang, Y., Azevedo, C.L., 2021. Bayesian optimization of area-based road pricing. In: *2021 7th International Conference on Models and Technologies for Intelligent Transportation Systems. MT-ITS, IEEE*, pp. 1–6.
- Liu, Z., Wang, S., Zhou, B., Cheng, Q., 2017. Robust optimization of distance-based tolls in a network considering stochastic day to day dynamics. *Transp. Res. C* 79, 58–72.
- Mariotte, G., Leclercq, L., Laval, J.A., 2017. Macroscopic urban dynamics: Analytical and numerical comparisons of existing models. *Transp. Res. B* 101, 245–267.
- Matérn, B., 2013. *Spatial Variation*, Vol. 36. Springer Science & Business Media.
- McKay, M.D., Beckman, R.J., Conover, W.J., 2000. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 42 (1), 55–61.

- Meng, Q., Lee, D.-H., Yang, H., Huang, H.-J., 2004. Transportation network optimization problems with stochastic user equilibrium constraints. *Transp. Res. Rec.* 1882 (1), 113–119.
- Meng, Q., Liu, Z., 2011. Trial-and-error method for congestion pricing scheme under side-constrained probit-based stochastic user equilibrium conditions. *Transportation* 38 (5), 819–843.
- Meng, Q., Liu, Z., 2012. Impact analysis of cordon-based congestion pricing on mode-split for a bimodal transportation network. *Transp. Res. C* 21 (1), 134–147.
- Meng, Q., Liu, Z., Wang, S., 2012. Optimal distance tolls under congestion pricing and continuously distributed value of time. *Transp. Res. E* 48 (5), 937–957.
- Meng, Q., Xu, W., Yang, H., 2005. Trial-and-error procedure for implementing a road-pricing scheme. *Transp. Res. Rec.* 1923 (1), 103–109.
- Osorio, C., Bierlaire, M., 2013. A simulation-based optimization framework for urban transportation problems. *Oper. Res.* 61 (6), 1333–1345.
- Pigou, A.C., 2013. *The Economics of Welfare*. Palgrave Macmillan.
- Rambha, T., Boyles, S.D., 2016. Dynamic pricing in discrete time stochastic day-to-day route choice models. *Transp. Res. B* 92, 104–118.
- Sandholm, W.H., 2002. Evolutionary implementation and congestion pricing. *Rev. Econom. Stud.* 69 (3), 667–689.
- Schrank, D., Eisele, B., Lomax, T., Bak, J., 2015. 2015 Urban Mobility Scorecard. Technical Report, The Texas AM and INRIX.
- Seshadri, R., de Palma, A., Ben-Akiva, M., 2022. Congestion tolling—Dollars versus tokens: Within-day dynamics. *Transp. Res. C* 143, 103836.
- Simoni, M.D., Pel, A.J., Waraich, R.A., Hoogendoorn, S.P., 2015. Marginal cost congestion pricing based on the network fundamental diagram. *Transp. Res. C* 56, 221–238.
- Small, K.A., 2015. The bottleneck model: An assessment and interpretation. *Econ. Transp.* 4 (1–2), 110–117.
- Small, K.A., Verhoef, E.T., Lindsey, R., 2007. *The Economics of Urban Transportation*. Routledge.
- Smith, M., 1979. The marginal cost taxation of a transportation network. *Transp. Res. B* 13 (3), 237–242.
- Srinivas, N., Krause, A., Kakade, S.M., Seeger, M., 2009. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*.
- Verhoef, E.T., 2002. Second-best congestion pricing in general networks. Heuristic algorithms for finding second-best optimal toll levels and toll points. *Transp. Res. B* 36 (8), 707–729.
- Vickrey, W.S., 1969. Congestion theory and transport investment. *Am. Econ. Rev.* 59 (2), 251–260.
- Vonk Noordegraaf, D., Annema, J.A., van Wee, B., 2014. Policy implementation lessons from six road pricing cases. *Transp. Res. A* 59, 172–191.
- Wang, X., Yang, H., 2012. Bisection-based trial-and-error implementation of marginal cost pricing and tradable credit scheme. *Transp. Res. B* 46 (9), 1085–1096.
- Watling, D., 1996. Asymmetric problems and stochastic process models of traffic assignment. *Transp. Res. B* 30 (5), 339–357.
- Wie, B.-W., Tobin, R.L., 1998. Dynamic congestion pricing models for general traffic networks. *Transp. Res. B* 32 (5), 313–327.
- Williams, C.K., Rasmussen, C.E., 2006. *Gaussian Processes for Machine Learning*, Vol. 2. MIT press Cambridge, MA.
- Xu, W., Yang, H., Han, D., 2013. Sequential experimental approach for congestion pricing with multiple vehicle types and multiple time periods. *Transportmetrica B* 1 (2), 136–152.
- Yang, H., Huang, H.-J., 2005. *Mathematical and Economic Theory of Road Pricing*. Elsevier, Oxford.
- Yang, H., Meng, Q., Lee, D.-H., 2004. Trial-and-error implementation of marginal-cost pricing on networks in the absence of demand functions. *Transp. Res. B* 38 (6), 477–493.
- Yang, H., Xu, W., He, B.-s., Meng, Q., 2010. Road pricing for congestion control with unknown demand and cost functions. *Transp. Res. C* 18 (2), 157–175.
- Yang, F., Yin, Y., Lu, J., 2007. Steepest descent day-to-day dynamic toll. *Transp. Res. Rec.* 2039 (1), 83–90.
- Ye, H., Yang, H., Tan, Z., 2015. Learning marginal-cost pricing via trial-and-error procedure with day-to-day flow dynamics. *Transp. Res. Procedia* 7, 362–380.
- Yildirimoglu, M., Ramezani, M., 2020. Demand management with limited cooperation among travellers: A doubly dynamic approach. *Transp. Res. B* 132, 267–284.
- Yin, Y., Yang, H., 2004. Optimal tolls with a multiclass, bicriterion traffic network equilibrium. *Transp. Res. Rec.* 1882 (1), 45–52.
- Zhang, X., Yang, H., 2004. The optimal cordon-based network congestion pricing problem. *Transp. Res. B* 38 (6), 517–537.
- Zhang, W., et al., 2017. Simulation-based robust optimization for the schedule of single-direction bus transit route: The design of experiment. *Transp. Res. E* 106, 203–230.
- Zheng, N., R  rat, G., Geroliminis, N., 2016. Time-dependent area-based pricing for multimodal systems with heterogeneous users in an agent-based environment. *Transp. Res. C* 62, 133–148.
- Zheng, N., Waraich, R.A., Axhausen, K.W., Geroliminis, N., 2012. A dynamic cordon pricing scheme combining the macroscopic fundamental diagram and an agent-based traffic model. *Transp. Res. A* 46 (8), 1291–1303.
- Zhong, S., Gong, Y., Zhou, Z., Cheng, R., Xiao, F., 2021. Active learning for multi-objective optimal road congestion pricing considering negative land use effect. *Transp. Res. C* 125, 103002.
- Zhou, B., Bliemer, M., Yang, H., He, J., 2015. A trial-and-error congestion pricing scheme for networks with elastic demand and link capacity constraints. *Transp. Res. B* 72, 77–92.