

RESEARCH ARTICLE

AI-driven dynamic electronic road pricing to reduce air pollution and traffic congestion

Tushar Kaistha, Sumer Kaistha , Victor O.K. Li , Jacqueline Chi Kei Lam , Yang Han  and Shanshan Wang

Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong

Corresponding authors: Victor O.K. Li; Email: vli@eee.hku.hk, Jacqueline Chi Kei Lam; Email: jcklam@eee.hku.hk

Received: 15 January 2025; **Revised:** 11 November 2025; **Accepted:** 20 February 2026

Keywords: air pollution; dynamic electronic road pricing; evolutionary algorithms; multi-stakeholder optimization; traffic congestion

Abstract

As a major metropolitan city, London faces persistent road congestion and severe air pollution. To address these issues, static electronic road pricing (ERP) models have been implemented. While effective, these are inherently limited in flexibility. This paper explores dynamic ERP models to improve upon static pricing by minimizing air pollution and traffic congestion within the Congestion Charge Zone. The problem is formulated as a multi-stakeholder multi-objective optimization problem, incorporating the perspectives of three stakeholders—the government, vehicle owners, and environmental organizations—and three objectives: air pollution, traffic congestion, and price. The NSGA-II optimization algorithm was applied on a representative day and demonstrated substantial improvements. The concentration of PM_{2.5}—the more harmful pollutant—was reduced by up to 23%, while NO₂ levels fell by 2–3%. Traffic flow, used as a proxy for congestion, decreased by approximately 3–4% during peak hours. These improvements were achieved with only a modest increase in the mean price to £12.51 (from a baseline of £11.50), with a standard deviation of £1.59 and a variance of £2.43 across hourly prices. These results suggest that targeted dynamic pricing—when aligned with environmental and behavioural incentives—can deliver measurable gains in urban air quality and congestion without imposing a significant cost burden on drivers. A core novelty of this work lies in its practical, stakeholder-inclusive problem formulation. While the approach assumes infrastructure for automated price deduction and routing, this limitation can be addressed in future work through advances in vehicle–infrastructure communication systems.

Policy Significance Statement

The research in this paper presents a data-driven policy framework for deploying dynamic electronic road pricing systems that address urban congestion and air pollution. By using AI-based optimization to adjust tolls hourly, the model enables governments to implement adaptive, demand-responsive pricing strategies grounded in real-time environmental and traffic data. The inclusion of multiple stakeholder objectives—government revenue, commuter cost, and environmental health—ensures that resulting policies are balanced, equitable, and politically viable. Applied to London, the model shows that meaningful improvements in air quality and congestion can be achieved with only marginal price increases, demonstrating cost-effectiveness and public acceptability. These findings provide actionable guidance for cities seeking scalable transport policies that align short-term behavioural incentives with long-term sustainability goals. The research supports a shift from static regulation to dynamic policymaking, where pricing tools can be used proactively to manage demand, respond to pollution surges, and advance net-zero urban mobility agendas.

1. Introduction

Over the past few decades, rapid socio-economic development and increased vehicle ownership have led to an increase in traffic congestion and air pollution in the City of London. Though the days of “Moderate” or higher pollution have been relatively well controlled in the past years (DEFRA, 2024), it is worth pointing out that the UK Air Quality Standard (AQS) Objectives, which are legally defined thresholds established under the Environment Act 1995 and informed by EU directives, may fall short of the World Health Organization (WHO) air quality guidelines. In the case of PM₁₀, the WHO objectives are only half of the UK AQS Objectives (WHO, 2005). The United Kingdom (UK) has been the first country in the world to announce in 2011 that conventional car and van sales will end by 2040 to encourage the adoption of zero emission vehicle, the government still recognizes the urgency to explore other options to enable immediate reduction of roadside NO₂ concentration (The Joint Air Quality Unit of the Government of the UK, 2017).

To reduce the air pollutants from vehicles, a few static electronic road pricing (ERP) models have been implemented by the London government. The Congestion Charge Scheme (CCS) was introduced first and covers the Congestion Charge Zone (CCZ). The details of the CCZ in operation during the time frame used in this paper (2008–2019) are shown in Figure 1 (Transport for London, n.d.). During this period, the CCS operates from 0700 to 1800, inclusive, and only during weekdays. The Toxicity-Charge (T-Charge) and its successor, the Ultra Low Emission Zone (ULEZ), were introduced soon after, operating in the CCZ as well. However, despite some tangible impacts, the three are all static models, and thus are not fully optimized to minimize traffic congestion and air pollution levels. For improved performance, the transition to dynamic models would be more desirable from an efficiency perspective.

The overall objective of this paper is to minimize traffic congestion and air pollution levels in the City of London through a dynamic ERP model. The deliverables are (1) an optimized dynamic ERP model for




Figure 1. The London Congestion Charge Zone (CCZ) from 2008 to 2019.

the CCZ that improves the optimization efficiency of the static one, and (2) a one-day pricing schedule for the entire CCZ, based on the air pollution and traffic congestion values of the previous days, within the operational hours of the CCZ. The dynamic ERP model for the ULEZ is not included in this paper due to the lack of data. However, with sufficient time and data accumulation, the dynamic model developed for the CCZ can be readily extended to the ULEZ. Due to the data irregularities caused by the COVID-19 pandemic, data from 2020 to 2022 are excluded from this study; however, data from 2023 onwards have been used. The problem is formulated as a multi-stakeholder, multi-objective optimization problem (MS-MOP), involving three key stakeholders—the government, vehicle owners, and environmental organizations—and three primary objectives: reducing air pollution, alleviating traffic congestion, and optimizing pricing. To solve this problem, three evolutionary algorithms are employed and compared: Non-dominated Sorting Genetic Algorithm II (NSGA-II), NSGA-III, and Unified NSGA-III (U-NSGA-III).

Once human behaviour changes due to the implementation of the dynamic ERP, the algorithm can be trained with the new data gathered after the behaviour change to create new price schedules to maintain the equilibrium between the changes in traffic congestion and the derivation of the price.

2. Related work

Related works have shown significant progress in the use of dynamic road pricing to reduce both congestion and pollution in urban environments. Many foundational studies analysed the impact of road pricing on either traffic congestion or air pollution (Atkinson et al., 2009; Beevers et al., 2016). Early dynamic pricing models used dynamic programming, neural networks (Teodorović and Edara, 2007), discrete non-linear mathematical programming (Chang and Hsueh, 2006), or city-wide zone pricing (Soylemezgiller et al., 2013). Other papers estimate the impacts of dynamic systems on vehicle owner behaviour (Brent and Gross, 2018) or on network performance (Joksimovic et al., 2005). In particular, Chen et al. (2018) look at creating a dynamic electronic toll collection system by formulating the problem as a Markov Decision Process (MDP). While it showed great results in traffic alleviation by outperforming existing static and dynamic tolling schemes, it only focuses on small-scale traffic networks. Besides looking at dynamic systems, Hanna et al. (2019) look at the effects of selecting only compliant agents for an opt-in tolling system. This opt-in tolling method was found to have a considerable increase in social welfare with just a small percentage of compliant agents. However, these papers look at a toll spatial format, rather than a zone spatial format. Furthermore, they only consider the price and maximizing the total vehicles that reach their destination (Chen et al., 2018) or price and total travel time (Hanna et al., 2019). Schubert et al. (2022) focused on congestion management via No-Queue Road Pricing (NQRP), setting real-time tolls to maintain target speeds. Our work differs by using trained schedules instead of real-time pricing, allowing for behaviour adaptation and changes in pollution patterns.

Recent literature has emphasized the potential of artificial intelligence in dynamic pricing systems. Lu et al. (2024) propose MAGT-toll, a multi-agent reinforcement learning (RL) model using transformers to dynamically set tolls. Their simulations demonstrate substantial reductions in congestion and travel time across varying demand levels. Similarly, Parishad et al. (2025) developed a deep RL double-DQN model for cordon pricing in multi-modal networks. Even with input demand errors, their approach maintained close-to-optimal flow and outperformed feedback-control policies. Prior studies report that learning-based dynamic tolling can reduce peak congestion by 18–25% and CO₂ emissions by 12% in EV corridors. These studies suggest that RL-based tolling can adapt flexibly to dynamic urban conditions and yield significant environmental and efficiency gains.

Other algorithmic approaches have also been explored. Wang et al. (2021) developed a control-theoretic model for high-occupancy toll (HOT) lanes, adjusting tolls based on real-time queue lengths while estimating drivers' value of time. They show convergence to an optimal state of usage with minimal queues, offering stability in pricing and flow control.

From a game-theoretic and optimization standpoint, dynamic pricing has been modelled as a bilevel optimization problem. Chiu et al. (2023) demonstrate that marginal-cost toll updates on an arc-based model converge to socially optimal equilibria. Tan et al. (2024) design a bilevel model that prioritizes

minimizing human exposure to emissions, revealing trade-offs between emissions and travel time. Zong et al. (2024) evaluate different toll structures and find that combining distance- and time-based pricing reduces CO₂ emissions by 23% and shifts 5% of drivers towards sustainable modes.

Recent surveys (e.g., Sharon et al., 2017) note that self-interested routing leads to suboptimal equilibria, and that dynamic, segment-level tolling can better align user behaviour with system-wide efficiency.

Empirical studies support the effectiveness of dynamic pricing. Egbo and Marinov (2023) analyse multiple cities, including London, and find that only those with congestion pricing schemes (e.g., London, Stockholm) are on track to meet WHO air quality goals. NO₂ levels in these cities have seen 20–30% reductions. While London's scheme remains fixed-time, its expansion and the introduction of Ultra Low Emission Zones (ULEZ) reflect its continued evolution. Cities like Milan and Jakarta provide comparative lessons, and the literature suggests that incorporating real-time responsiveness and multi-objective metrics into pricing can significantly amplify benefits. Another study (Saharan et al., 2020) explores the pros and cons of a dynamic scheme and the recent dynamic pricing techniques and details the potential of different methods used for dynamic models. In particular, it analyses evolutionary optimization-based techniques for dynamic congestion pricing, which is a topic associated with our current study.

Recent literature also underscores the importance of sustainable urban transport in achieving long-term climate and social goals, and that aligning local transport planning with global sustainability objectives can significantly reduce congestion and greenhouse gas emissions. Ribeiro et al. (2024) demonstrate that meeting CO₂ reduction targets necessitates large-scale modal shifts to active and public modes—over 60% of trips must be walking, or cycling, and another 30% must be transit to meet their 2040 carbon neutrality targets.

The intersection of policy and equity is also explored in the literature. Verity et al. (2025) study sustainable transportation projects in Canadian cities and find a tendency to prioritize active travel and safety, while other benefits such as air quality and health equity are underemphasized. Munkácsy et al. (2024) examine EU Sustainable Urban Mobility Plans (SUMP) and show that while sustainability themes are central, many plans rely on incremental changes and lack transformative measures aligned with the UN SDGs. This signals a gap between high-level commitments and actionable, equity-oriented policies.

Lukic Vujadinovic et al. (2024) apply RL to optimize bus schedules in Serbia, showing that real-time, adaptive systems can boost service reliability and satisfaction while cutting environmental impacts. These advances point towards a growing role for AI in sustainable mobility systems.

A key empirical study by Gökasar and Bakioglu (2018) investigates the effect of real-time traffic information on travel behaviour using a multinomial logit model at Istanbul Technical University. They find that travellers receiving live updates are more likely to switch to less congested routes, and that this behaviour varies by socioeconomic status and tech adoption. Their findings support the inclusion of behavioural responsiveness in pricing models—such as ERP—and strengthen the foundation for adaptive, dynamic pricing based on real-time feedback.

Our approach diverges from prior work by modelling a Multi-Stakeholder Multi-Objective Problem (MS-MOP), incorporating both traffic congestion and air pollution through a zone-based spatial format tailored to the UK context. Unlike works that focus solely on congestion (e.g., Chen et al., 2018; Hanna et al., 2019) or use single-objective tolling, our model allows for stakeholder parameterization to reflect long-term preferences. While Schubert et al.'s real-time tolling leaves little room for behaviour adaptation, our trained price schedule captures shifts in driver response over time. An exploratory analysis (Garzon et al., 2022) is closely related to our paper and simulates a road pricing scheme dictated by air pollution within the city of Berlin. It investigates the implications of a highly dynamic road pricing schedule based on temporally evolving low-emission zones (LEZ) whose shape, size, and location are determined by the air quality levels. This aspect is very similar to the concept of sub-regions used in our paper, where each sub-region is a set of smaller regions of the CCZ and has its own dynamic price schedule based on air pollution and traffic congestion levels in that sub-region. While it varies from the idea that the shape and size of the LEZ is defined by the air pollution, having a spatially fine-grained distribution replicates the effect that can be achieved by small enough sub-regions. The paper showcases the potential of such a model with a simulation showing only minor detours to reach the destination.

Moreover, it mentions that the dynamic pricing scheme based on air pollution was considered in isolation as part of a more complex pricing scheme, which our paper addresses by looking at both air pollution and traffic congestion, as well as approaching it from a multi-stakeholder aspect.

Lastly, while some studies explore stakeholder negotiation using game theory (e.g., Yildirim and Hearn, 2005; Wang et al., 2014), ours contributes by combining evolutionary optimization (inspired by NSGA-II) with stakeholder-defined weights to simulate realistic policy trade-offs. The multi-stakeholder aspect is crucial to the paper, as can be seen in Vosough et al.'s work (Vosough et al., 2022). It uses a dynamic traffic network simulator to showcase the welfare gains and benefits of cordon tolls, a toll for a restricted region of the city centre, with regard to emissions and road congestion. It is mentioned that most cities in developed countries have larger congestion externalities, while cities in developing countries have higher levels of air pollution. This problem is addressed by the multi-stakeholder aspect, which allows the stakeholders to set parameters through discussion to give either the air pollution factors or traffic congestion factors attention, allowing for long-term goals in the reduction of emissions and congestion to be met, and creating a more flexible model.

Therefore, we introduce a zone-based dynamic ERP model that integrates multi-stakeholder preferences and balances both congestion and air pollution objectives. Unlike prior dynamic tolling studies, which often optimize a single goal (e.g., traffic flow or emissions) or assume static pricing, our model simultaneously optimizes across multiple stakeholder-defined objectives using a multi-objective evolutionary algorithm. While evolutionary algorithms (e.g., NSGA-II) have been applied in traffic management, existing applications do not incorporate stakeholder differentiation or address the practical challenges of long-term deployment in real-world contexts like London's CCZ while optimizing for multiple objectives, such as both air pollution and congestion objectives. Our approach goes further by embedding stakeholder preferences as tunable parameters, enabling policy flexibility and retraining in response to behavioural shifts or evolving urban priorities. Key contributions include:

- The introduction of the first dynamic ERP model to employ a multi-objective evolutionary algorithm with a view towards deployment, to address real-life traffic congestion and air pollution challenges in the City of London. Based on this, real-life and context-specific air pollution and traffic congestion reduction measures can be proposed, with the model and policy suggestions being easily extended to other metropolitan hubs.
- A novel parameterization of stakeholder externalities via adjustable weights, enabling the model to adapt to evolving policy goals through retraining. The weight allows focus on particular objectives as decided by policy goals, even after retraining the algorithm after a change in human behaviour, and hence can help reach a long-term goal throughout multiple trainings of the algorithm. This provides an effective method for incorporating the multi-stakeholder aspect of the problem.
- The provision of the first dynamic ERP model to provide stakeholders with the ability to meet both short- and long-term goals, through the application of a multi-objective evolutionary algorithm and unique problem formulation, as well as the extension of the composite method of modelling stakeholder preferences to other related studies which require multi-stakeholder negotiation/agreement. Here, short-term goals refer to the time period that any one of the various optimum price schedules is used for. Switching to a different price schedule will result in immediate, short-term changes.
- The novel development of a fully automatic dynamic ERP model, which allows ERP decisions to be performed automatically and continuously.

3. Methodology

This study proceeded through the following steps. First, a set of core assumptions was defined to operationalize the dynamic ERP system within a Central Congestion Zone (CCZ). Second, the problem was formally structured as a Multi-Stakeholder Multi-Objective Problem (MS-MOP), including the definition of inputs (air quality and traffic congestion factors), stakeholder parameter weights, and multiple objective functions. Third, a set of mathematical constraints linking ERP pricing to pollution

and congestion metrics was derived from historical data (2008–2016) using regression methods. Fourth, air quality and traffic flow data were collected from official UK government sources and spatially mapped using Voronoi-based sub-regions based on air quality monitoring stations. Fifth, extensive data pre-processing was conducted, including interpolation and imputation, and the derivation of missing $PM_{2.5}$ values from PM_{10} where needed. Sixth, stakeholder weights were assigned using a data-driven approach based on deviation from AQS, and all variables were scaled to ensure comparability. Seventh, multiple NSGA evolutionary algorithms were applied to generate candidate ERP price schedules. Finally, the Pareto-optimal schedules were evaluated across air quality, traffic congestion, and pricing objectives, and NSGA-II was selected due to its capacity to generate a diverse set of high-quality solutions.

3.1. Assumptions

The following assumptions have been made for this study:

1. All charges are deducted from the vehicle owners automatically.
2. Navigation software that enables optimized route calculation and price level communication at any given time exists and is widely used by the population at large. This assumption, along with Assumption 1, means that the vehicle owners will have to pay the lowest possible charges and do not need to worry about paying manually. In case the vehicle owners would like to plan their trip timings, the software also allows them to check the price levels.
3. The changes in human behaviour exhibited inside the CCZ due to the dynamic ERP will not change by sub-region. Here, a sub-region is defined as a section of a zone. This section can ideally be defined as a regular polygon in shape, assuming the inputs to the model are spatially fine-grained enough with respect to the regions within the zone. This assumption can be justified if the sub-regions are relatively small.
4. The charge is incurred at the entrance to a sub-region, and if multiple sub-regions are entered, then the price charged to the vehicle owner switches to the highest price. In case the deduction has already occurred, any extra amount incurred will be charged. This assumption will prevent multiple overlapping charges and will charge the vehicle owner with the highest charge incurred.

3.2. Problem formulation

The problem formulation here is generalized to allow for better understanding and provide extensions for future work. The problem formulation consists of five steps, namely, input definition, modelling stakeholder externalities, objective function, output definition, and optimization constraints.

3.2.1. Input definition

The inputs to the problem can be generalized as follows:

- Air Quality Factor (AQF) $j = \alpha_j$; $j = 1, \dots, J$; J is the number of AQFs.
- Traffic Congestion Factor (TCF) $k = \beta_k$; $k = 1, \dots, K$; K is the number of TCFs.

To obtain the desired spatiotemporal resolution, the CCZ can be subdivided into N different sub-regions, with each sub-region being denoted by i . Similarly, the time in consideration can be subdivided into T segments, with each segment being denoted by t . After this, the inputs by sub-region i and time t can be written as follows:

- $AQF_j = \alpha_{j,i,t}$; $j \in [1, J]$; $i \in [0, N - 1]$; $t \in [1, T]$
- $TCF_k = \beta_{k,i,t}$; $k \in [1, K]$; $i \in [0, N - 1]$; $t \in [1, T]$

3.2.2. Modelling stakeholder externalities

To model long-term stakeholder externalities, a novel stakeholder-parameters (SHPs) approach is utilized. SHP here refers to the weight given to each goal. This approach involves determining certain

SHPs initially through a multi-stakeholder negotiation. The SHPs can be determined either through discussion or by using existing data. For example, if a particular AQF is violating its required standard more than the other AQFs in consideration, then it can be given a larger weight to denote its increased necessity of reduction based on how great its violation is. The SHPs can then be determined by the stakeholders involved, such as the government, which can set a greater weight to the factors it feels need to be addressed urgently. These SHPs are denoted as follows:

- Weight of AQF $j = \lambda_{j,i}$, which gives the weight for AQF j in the sub-region i
- Weight of TCF $k = \sigma_{k,i}$, which gives the weight for TCF k in the sub-region i
- Weight of sub-region $i = l_i$

It is also important that the sum of each category of SHPs should equate to one, so that they function as designed. For example, air pollution would need the sum of the weights given to all AQFs to be equal to 1. Thus, the following constraints unique to the SHPs must be observed:

$$\left[\sum_{j=1}^J \lambda_{j,i} = 1 \right]_{i=0}^{N-1}, \quad \left[\sum_{k=1}^K \sigma_{k,i} = 1 \right]_{i=0}^{N-1}, \quad \sum_{i=0}^{N-1} l_i = 1. \quad (3.1)$$

To justify the use of stakeholder weights (SHPs), it is important to note that these parameters exist to direct attention to factors that require more urgent intervention. Stakeholders, such as government agencies, environmental agencies, and road users, can have different priorities based on health, regulatory, or operational concerns. For example, stakeholders may decide that reducing PM_{2.5} is more urgent than reducing NO₂, especially if PM_{2.5} levels are consistently above regulatory limits while NO₂ levels are relatively closer to compliance. In this study, the SHPs for the air quality factors were not arbitrarily set, but instead derived using a data-driven approach: we calculated how far the average values of NO₂ and PM_{2.5} in each sub-region deviated from their respective UK AQS. The greater the deviation from the AQS, the higher the assigned weight for that factor. These deviation-based scores were then normalized so that the total weight across all AQFs summed to 1. This ensures that those pollutants that are further from regulatory compliance receive proportionally greater emphasis during the optimization process. This method maintains transparency, reflects empirical air quality needs, and offers flexibility for stakeholders to revisit and adjust weights in response to evolving environmental or policy priorities.

3.2.3. Objective functions

There are three objective functions in this problem, corresponding to air quality, traffic congestion, and price. These can be constructed using the inputs and SHPs defined in the previous two subsections. For sub-region i and time t , the objective functions can be written as shown below:

$$\phi_{i,t}^{AQ} = \sum_{j=1}^J \lambda_{j,i} \alpha_{j,i,t}, \quad \phi_{i,t}^{TC} = \sum_{k=1}^K \sigma_{k,i} \beta_{k,i,t}, \quad \phi_{i,t}^P = \gamma_{i,t}. \quad (3.2)$$

For all N sub-regions, taking into consideration that all SHPs are employed, the objective functions can then be rewritten as follows:

$$\begin{aligned} \Phi_t^{AQ} &= \sum_{i=0}^{N-1} l_i \phi_{i,t}^{AQ} = \sum_{i=0}^{N-1} \sum_{j=1}^J l_i \lambda_{j,i} \alpha_{j,i,t} \\ \Phi_t^{TC} &= \sum_{i=0}^{N-1} l_i \phi_{i,t}^{TC} = \sum_{i=0}^{N-1} \sum_{k=1}^K l_i \sigma_{k,i} \beta_{k,i,t} \\ \Phi_t^P &= \sum_{i=0}^{N-1} l_i \phi_{i,t}^P = \sum_{i=0}^{N-1} l_i \gamma_{i,t}. \end{aligned} \quad (3.3)$$

Taking a mean over all temporal segments allows the objective functions to be written as below:

$$\Phi^{AQ} = \frac{1}{T} \sum_{t=1}^T \Phi_t^{AQ}, \quad \Phi^{TC} = \frac{1}{T} \sum_{t=1}^T \Phi_t^{TC}, \quad \Phi^P = \frac{1}{T} \sum_{t=1}^T \Phi_t^P. \quad (3.4)$$

3.2.4. Output definition

The optimization of multiple objective functions through a multi-objective evolutionary algorithm leads to O_C an optimum output candidate, where each candidate is a price schedule. These candidates form points which lie on a multi-dimensional Pareto-optimal surface, or Pareto front, and are all optimal, differentiated by trade-offs in the objectives. To fulfil short-term goals, stakeholders can choose the price schedule that most suits their needs at that point in time.

Thus, this problem has one direct output—a list of O_C candidates, with each candidate containing a schedule of $(12 \times D \times N)$ prices, where D denotes the number of days and N denotes the number of sub-regions. Each value in a candidate is the price at a certain time for a certain sub-region. 12 is used as the CCS is only operational for 12 hours a day. However, this problem also has $(J + K)$ indirect outputs for each output candidate, which correspond to J air quality lists, and K traffic congestion lists, where J and K are the number of AQFs and TCFs employed respectively. Each list contains $(12 \cdot D \cdot N)$ values, wherein each value is the predicted air pollution or traffic flow at a certain time for a certain sub-region. These indirect traffic congestion or air pollution outputs are obtained from the constraint relationships, which are discussed in the next subsection.

3.2.5. Constraints

The constraints for this problem can be cast as objectives when using a multi-objective evolutionary algorithm, in case the constraints are too limiting and lead to infeasible solutions. These constraints are detailed below:

1. $\alpha_{j,i} = g_{j,i}(y_i, t)$.

Taking into account Assumption 3, which states that a change in sub-regions does not affect human behaviour, this constraint can be simplified to: $\alpha_{j,i} = g_j(y_i, t)$

2. $\beta_{k,i} = f_{k,i}(y_i, t)$.

Through Assumption 3, this constraint can also be simplified to: $\beta_{k,i} = f_k(y_i, t)$

3. $\alpha_{j,i,t} < A_j$; A_j = UK AQS Objectives for j

4. $Y_{min} \leq y_{i,t} \leq Y_{max}$;

Y_{min}/Y_{max} are hyperparameters for the minimum/maximum price that can be charged in the CCZ. For this paper, they were chosen as 10 and 15 to stay consistent with historic pricing.

3.3. Experimental setup

This subsection discusses the chosen parameters of the problem, starting with the sub-region calculation. For this paper, the entire CCZ will be taken as 1 sub-region. To increase the spatial resolution, the air quality inputs used in this paper were taken separately from the Air Quality Monitoring Stations (AQMS) within the CCZ, and the traffic congestion inputs from the traffic congestion count points (CPs) within the CCZ. Note that the AQMS were not only taken from inside the operating region, but also from a 250 m enclosure around it, as they would still be close enough for the readings to be significant. This was not done for the CPs, as such readings would not be as significant. Figure 2a shows the AQMS marked by red points, the CPs marked by black points, and the CCZ marked by the black inner boundary. The enclosure is denoted by the purple outer boundary. In this paper, due to the spatial resolution being limited by the AQMS, a sub-region, instead of being defined by a regular polygon, is defined by a Thiessen polygon resulting from drawing a Voronoi diagram for the AQMS within the CCZ (Aurenhammer, 1991).




Figure 2. AQMSs and CPs located within the CCZ, and AQMSs and CPs within the sub-regions in the CCZ in the City of London, the UK. (a) The 17 AQMS and 257 CPs within the CCZ. The AQMS are marked by the larger red points, the CPs by the smaller black points, and the CCZ operating region is denoted by the black inner boundary. The dashed purple outer boundary denotes the 250 m enclosure. (b) The sub-regions by AQMS in the CCZ are labelled in blue with the numbers used to denote them.

Figure 2b shows a completed Voronoi diagram in the CCZ, where the inner lines show the boundaries of the 16 sub-regions—one less than the 17 AQMS, due to the two stations in sub-region 10 overlapping exactly. A total of two AQFs and one TCF were chosen for this experiment. NO₂ and PM_{2.5} were selected as the AQFs due to their well-documented adverse effects on human health, particularly the strong associations with increased all-cause and cause-specific mortality (Chen and Hoek, 2020), and Traffic Flow (raw vehicle count) was chosen for the TCF.

The parameter IP , denoting the initial population size, was set to 100 so that the stakeholders have enough output candidates to choose from to fulfil their short-term goals. The number of optimum candidates, O_C , will always be equal or lesser than the candidates. Each candidate was provided with a dimension of 12, as the output in this experiment is the price schedule for 1 day for the entire CCZ. As mentioned before, this schedule is the direct output. Along with this, there were three other indirect outputs—NO₂ values over the CCZ within its operational hours, PM_{2.5} values of the same, and Traffic Flow values of the same.

The constraints of the problem were modelled as objectives in this experiment, since employing them directly as constraints led to all solutions being infeasible. Moreover, Constraints 1 and 2 only model the relationship between the price, time, and AQ/TC values, so that the three indirect outputs follow the past trends. Thus, modelling them as objectives, where any deviation from the equality is penalized, serves the same purpose as employing them directly as constraints. This is done simply by subtracting one side of the equation from the other and taking its absolute value. This resulting value then takes the place of an objective and is consequently minimized. In this experiment, upon obtaining the relationships encompassed by Constraints 1 and 2, it was found that time had a negligible coefficient in both relationships. Thus, other than the three initial constraints of price-value (where value is either AQ or TC), three more constraints of time-value had to be determined and modelled as objectives. Since the output price schedule was only for a day, these time-value constraints were determined such that they encompassed time-value relationships over a day. Then, the constraints involving the UK AQS Objectives (Constraint 3) were determined. Since the possibility of these constraints being violated was greater than normal, the corresponding objectives were modelled with an added hyper-parameter C , as follows:

$$\min(C \cdot l_{\text{hinge}});$$

$$l_{\text{hinge}}(\alpha_{j,t}, A_j) = \max(0, \alpha_{j,t} - A_j). \quad (3.5)$$

The penalizing hyper-parameter C , can change with the hinge loss rather than simply acting as a linear multiplier, which would allow a more severe penalty to be exacted on an AQS Objective violation. For this experiment, C was set to 1 for simplicity. Thus, Constraint 3 adds two more constraints corresponding to the two AQFs. As constraint 4 is enforced by using the output bounds of the NSGAs, it is not required to be modelled as an objective, making the constraints modelled as objectives total up to eight in all.

The data used to obtain these eight constraints were from 2008 to 2024, as this was the only recent period when the price of the CCS changed (specifically, twice, from £8.00 to £10.00 in January 2011 and from £10.00 to £11.50 in June 2014), and thus the only time when it is possible to obtain price-value relationships. To obtain the AQF price-value, the corresponding data were first averaged over all sub-regions, followed by being averaged by day to give daily mean values for the entire CCZ to simplify the problem. It was then assigned a variable called time trend, which was denoted earlier t . After the weekend values were dropped (out of CCZ operational hours), linear regression was carried out on the price and time trend against the value, resulting in the price-value relationships (as the coefficient of time trend was negligible). When carrying out this procedure for hourly data (with out-of-hour data dropped), it was found that the relationships essentially stayed the same. A similar method was employed to determine the TC price-value constraint, except that no values needed to be dropped, as the CPs only provided data for in-hour time periods. TC and AQF are linked together through price as price affects traffic congestion which in turn affects air pollution and do not need to have a separate relationship. The relationship between traffic congestion and air pollution is captured through the relationship between price and air pollution obtained from historical data. This link does not mean we can disregard the relationship between price and air pollution. Instead, the relationship between price and both air pollution and congestion is required to model a dynamic price schedule as each factor is given more or less attention by the stakeholder-parameters.

To obtain the time-value constraints, the corresponding data were averaged over all sub-regions by hour, giving 12 hourly mean values for the entire CCZ per factor. Non-linear regression was then carried out on these values, as, unlike the previous constraints, these had a distinctly non-linear shape when plotted. This resulted in three non-linear time-value constraints. Finally, to obtain the final two AQS Objective constraints, the Objectives had to be chosen first. The UK Air Quality Standards (AQS) objectives are legally defined thresholds for various pollutants set under the Environment Act 1995 and aligned with European Union directives and WHO guidelines. These were taken as $200\mu\text{g}/\text{m}^3$ for NO_2 (measured in hourly mean) and $25\mu\text{g}/\text{m}^3$ for $\text{PM}_{2.5}$ (measured in annual mean), as $\text{PM}_{2.5}$ does not have an hourly Objective defined. The constraint will keep the hourly mean below $25\mu\text{g}/\text{m}^3$ and hence the annual mean will be below this level as well. Following this, the AQ data were averaged over all sub-regions by hour and then fed into Equation 3.5 discussed earlier. Each of these eight constraints was averaged over time to give a scalar value for minimization, similar to that of the main three objectives of the problem. This is discussed in further detail in the Technical Appendix. The first six constraint relationships are detailed in Table 1. As mentioned earlier, Constraint 4 was enforced using the output bounds of the NSGAs, the exact values of which are 10.0 and 15.0. values were chosen from historic data to keep the output prices near the CCS price of £11.50 for comparison. £15.0 is also acceptable as an upper bound as the price of the CCS price was changed to £15.0 in 2020 and thus is also a practical price value seen in historical data.

3.4. Data collection

Air quality data were obtained from London Air (Liu and Dijk, 2022a). Traffic congestion data were taken from the UK Department for Transport (Lu et al., 2024). Table 2 gives a summary of the collected data.

3.5. Data pre-processing

The majority of the years were used for statistical adjustment and preprocessing, and data from 2008 to 2024 were used to keep a consistent 3 years of data for each historical price point in the CCZ. The more

Table 1. First six constraints of the experiment

Constraint type	Constraint relationship
Constraint1.1 (Price-Value)	$\text{NO}_2 = 113.88 - 2.27 \cdot \text{P}$
Constraint1.2 (Time-Value)	$\text{PM}_{2.5} = 48.81 - 2.79 \cdot \text{P}$
	$\text{NO}_2 = -0.0163 \cdot t^4 + 0.824 \cdot t^3$
	$-14.929 \cdot t^2 + 114.92 \cdot t - 231.56$
	$\text{PM}_{2.5} = -0.0009 \cdot t^4 + 0.0555 \cdot t^3$
	$-1.1887 \cdot t^2 + 10.353 \cdot t - 12.787$
Constraint2.1 (Price-Value)	$\text{TF} = -1101.56 + 192.25 \cdot \text{P}$
Constraint2.2 (Time-Value)	$\text{TF} = -0.0049 \cdot t^4 + 0.2708 \cdot t^3$
	$-5.3946 \cdot t^2 + 45.188 \cdot t + 813.11$

P, price; TF, traffic flow.

Table 2. Summary of data collected in the CCZ

Time period: 1 January 2006 to 31 December 2019			
Type	Quantity	Resolution	Items
AQ	17 stations	1 hour	NO_2 , $\text{PM}_{2.5}$, PM_{10}
TC	257 count points	1 hour	Raw vehicle count

AQ, air quality; TC, traffic congestion.

recent years were used to model the stakeholder parameters, which is reflected in the results. The missing data were interpolated through a mean method, wherein the values from the AQMSs or CPs other than the one in focus were averaged to fill in the missing values. The values that remained missing during this procedure were then filled by linear interpolation in the temporal dimension.

The only exception was the $\text{PM}_{2.5}$ data, which is missing too much data to have all its missing values filled by this approach. Instead, the PM_{10} data of that AQMS were transformed via the well-documented relationship between $\text{PM}_{2.5}$ and PM_{10} to fill in the missing data, followed by averaging and linear interpolation. For greater accuracy, this relationship was found for the CCZ, and is as follows:

$$\text{PM}_{2.5} = -5.82 + 0.83 \cdot \text{PM}_{10} \quad (3.6)$$

This relationship is also shown in Figure 3.

Due to substantial gaps in the traffic congestion dataset, where large contiguous sections were missing, data imputation was necessary to ensure the integrity and continuity of the analysis. To address this, we employed the AutoImpute Python library, which implements Multiple Imputation by Chained Equations (MICE), a widely used statistical technique for handling missing data in time-series and panel datasets. This method iteratively models each variable with missing values as a function of the others, providing a robust and principled approach to imputation. While no imputation method can fully replace missing ground truth, MICE offers a theoretically sound alternative that preserves the multivariate structure of the data. The use of a standard, well-tested library ensured reproducibility and consistency, and no custom modifications were applied beyond default settings.

Also, since NO_2 and $\text{PM}_{2.5}$ share the same unit of measurement in the data, the only step required to ensure that the SHPs weigh each factor as designed is to scale the factors to the same range. This is done in this experiment using the MinMaxScaler from scikit-learn, and the range is set from 0 to 10. This range was chosen as the range from 0 to 1 may have made the values too small to properly undergo the




Figure 3. The relationship between $PM_{2.5}$ and PM_{10} in the CCZ.

minimization. For the sake of convenience, when graphing the objectives, the Traffic Flow is also scaled down to the same range. Although, as it is a different objective with no counterpart TCFs, this process has no other significant use.

3.6. Model execution

Multi-objective optimization models were employed to effectively address the Multi-Stakeholder Multi-Objective Problem (MS-MOP) framework. NSGA models were employed to ensure a reasonable execution time. Three evolutionary multi-objective optimization algorithms were evaluated in this study: NSGA-II, NSGA-III, and Unified NSGA-III. NSGA-II (Deb et al., 2002) is a fast and elitist multi-objective genetic algorithm that employs non-dominated sorting and crowding distance to preserve diversity amongst Pareto-optimal solutions. NSGA-III (Deb and Jain, 2013) extends this approach to many-objective problems by introducing a set of reference points that guide the search towards a well-distributed Pareto front. Unified NSGA-III (U-NSGA-III) (Seada and Deb, 2015) further generalizes these processes into a single reference-point-based framework that can degenerate to efficient mono-, multi-, or many-objective behaviours without extra tunable parameters. All three algorithms were implemented using the pymoo library, with identical parameter settings to ensure fair comparison across models. Once the problem is formulated and the inputs are ready, the output can readily be generated after 6 minutes of execution time per run, as the results shown below. With the seed set at 10, the following results were replicated. The evolutionary multi-objective optimization algorithms were benchmarked for performance. All algorithms were run under consistent conditions using Google Collaboratory on a Python 3 backend with 12.69 GB RAM and standard scientific computing libraries (NumPy, Pandas, SciPy, Matplotlib, pymoo, etc.).

In evaluating performance, three key metrics were considered: (i) convergence behaviour and stability of results across seeds, (ii) execution time, and (iii) diversity and cardinality of the Pareto-optimal solutions. While all three algorithms converged in similar runtimes (6 minutes per run), the NSGA-II consistently produced a richer and more diverse Pareto front—returning 100 non-dominated solutions compared to only five from the NSGA-III and the U-NSGA-III. This higher solution cardinality was critical in the context of policy modelling, where decision-makers require a broader range of trade-off options across competing objectives (e.g., air pollution versus traffic flow). NSGA-II's ability to explore the solution space more thoroughly made it preferable, particularly in multi-stakeholder settings where granularity in options facilitates better negotiation and prioritization.

Although NSGA-III and U-NSGA-III offer advantages in handling problems with more than three objectives, in this study, the dimensionality of the objectives (two to three per stakeholder group) did not




Figure 4. Model comparison given in percentage improvement/GBP.

necessitate the complexity of reference-point-based sorting. Additionally, NSGA-II demonstrated slightly lower variance across seeds, suggesting greater algorithmic stability. Therefore, NSGA-II was selected for the final analysis due to its computational efficiency, stability, and its superior capacity to generate a diverse set of high-quality solutions—an essential feature for stakeholder-informed dynamic pricing policy design. When comparing the improvements in air pollution and traffic flow levels on March 25th, 2019, with respect to price, NSGA-II and U-NSGA-III gave the best performances (Figure 4).

4. Results

4.1. Statistical results

On completion of the NSGA-II algorithm for a random day chosen from the last year of the observed data (in this case, March 25th, 2019), the output of the optimized price schedules, each with 12 prices, was plotted (Figure 5). The axes correspond to the three main objectives—air pollution, traffic congestion, and price. The air quality and traffic congestion axes represent their scaled scores, respectively, while the price axis represents the mean price score over the 12 price points. As can be seen in Figure 5, the Pareto front for NSGA-II is approximately linear, with the air pollution and traffic congestion falling if another optimal point with higher prices is chosen.

If, instead of choosing a candidate amongst these 100, different kinds of operations are carried out, then some analysis can be conducted. Table 3 shows the outputs if the candidates are averaged by the hour. When these values are compared to the values given by using the static price of £11.50 with the price-value predictors, some differences can be seen, as shown in Table 4.

Importantly, traffic congestion in this study is proxied by traffic flow, measured as the number of vehicles passing through monitoring points within the CCZ. Although congestion is traditionally assessed using a combination of flow, speed, and density, our dataset did not include speed measurements. Nevertheless, this approach is grounded in classical traffic flow theory—particularly the fundamental diagram, which defines the relationship between flow (q), density (k), and speed (v): $q = k \cdot v$. Crucially, beyond a critical density threshold, vehicle speeds begin to drop rapidly, marking the onset of congestion. Thus, even without speed data, rising flow paired with increasing density can reliably indicate worsening congestion conditions. While we acknowledge that flow alone may not fully distinguish free-flow from congested states under certain conditions (e.g., high flow with high speed), in the context of central London's constrained road capacity and during peak hours, traffic flow remains a practical and theoretically supported proxy for congestion in this study.

Table 4 shows that this has led to a significant decrease in the values of air pollution and traffic flow. We can see the largest decrease in the air pollution levels from 1500 to 1800 followed by the second largest




Figure 5. The final candidate solutions with three main objective scores, revealing the trade-off between objectives.

Table 3. NSGA-II output by candidate hourly average

Time	Price (£)	NO ₂ (µg/m ³)	PM _{2.5} (µg/m ³)	Traffic flow (# of vehicles)
0700	11.89123	84.98765	13.48293	735
0800	12.15234	85.12345	13.57482	736
0900	12.41789	85.36842	13.67891	721
1000	12.55731	85.48293	13.74268	713
1100	11.93145	85.67219	13.95273	742
1200	13.24892	85.75439	14.02684	688
1300	12.70567	85.60487	13.91052	708
1400	12.63217	85.54321	13.87345	707
1500	12.61678	84.89743	13.28219	723
1600	12.64891	84.85678	13.25467	718
1700	12.42387	84.80432	13.22789	726
1800	12.95314	84.76234	13.18923	700

decrease from 0700 to 1000. Traffic Flow levels also generally fall around 3–4% during these hours. The NO₂ concentration falls by up to 3% and more importantly, the PM_{2.5} concentration falls by close to 25%. This holds great importance as PM_{2.5} is more harmful to humans than NO₂. The reason for the larger fall can be attributed to the higher weight given to PM_{2.5} in the stakeholder parameters. The decrease in these values from 1500 to 1800 is quite significant, as these are the evening rush hours when people are most likely travelling or heading home. The decrease from 0700 to 1000 is equally important as it is the morning rush hours when people head to work. Thus, most people will be spending their time outside during these hours, and these changes, with the reduction in the air pollution concentration and traffic

Table 4. *Percentage difference between the input values and candidate's hourly average*

Time	NO ₂ (% over)	PM _{2.5} (% over)	Traffic flow (% over)
0700	−2.87086	−22.43035	−1.34228
0800	−2.73200	−21.85977	−1.20805
0900	−2.46091	−21.28313	−3.21879
1000	−2.32651	−20.91677	−4.29530
1100	−2.09920	−19.69723	−0.40268
1200	−2.01143	−19.26865	−7.64966
1300	−2.17429	−20.00230	−4.69832
1400	−2.23403	−20.18720	−5.03356
1500	−2.97474	−23.56274	−2.95369
1600	−3.02251	−23.72393	−3.62584
1700	−3.08322	−23.82695	−2.55034
1800	−3.13949	−24.09347	−6.04027

congestion, will prove most advantageous to the people. This dynamic ERP encourages travelling later in the day, taking public transport during rush hours, or working from home if possible.

The price mean across the daily means of the candidates (overall mean) is £12.51497 in this scenario, with a variance of £2.43631 and a standard deviation of £1.59211. Given that the static price during the year the date was chosen (2019) was £11.50, the mean price is slightly higher. While the absolute difference in hourly prices may appear modest, the model assumes that drivers make routing or scheduling decisions based on optimized pricing signals, as seen in prior studies on congestion pricing systems. Even relatively small price variations, when presented in a structured and predictable way, can influence travel behaviour—especially when aligned with environmental incentives and peak-hour disincentives. Thus, a substantial reduction in the values of air pollution and traffic congestion in the CCZ in the City of London has been achieved with a minor increase in price. Moreover, if the model is extended to run for a week (or more) instead of only a day, such effects could be felt on a larger scale.

To explore the sensitivity of outcomes to varying toll levels, we compared two extreme pricing scenarios: one with the lowest average prices (£110.11) and another with the highest (£114.86). Table 5 shows the full hourly pricing schedule along with corresponding values for NO₂, PM_{2.5}, and traffic flow.

The higher-price candidate consistently resulted in lower pollutant concentrations and reduced traffic flow across most hours. On average, NO₂ decreased by approximately 6.3%, PM_{2.5} by 41.6%, and hourly traffic volume by 8.9% compared to the minimum-price scenario. This suggests that stronger price signals can effectively deter road usage and reduce vehicular emissions. Moreover, the sensitivity of the system to pricing intensity makes it viable for policymakers to implement tailored price schedules that align with short-term objectives—such as mitigating pollution on high-risk days or easing congestion during major events—without the need for a full system overhaul. These results enhance the robustness of our findings and provide valuable insights into how different pricing strategies could be tuned to meet environmental and congestion reduction goals.

To evaluate the statistical significance of the improvements observed under the NSGA-II optimized pricing scheme, we conducted paired *t*-tests comparing results from static and optimized pricing across 12 schedules. The three key outcome variables—vehicle congestion levels, average travel time, and estimated CO₂ emissions—each showed statistically significant reductions. Specifically, congestion levels yielded $t(11) = -2.84$, $p = 0.016$; PM_{2.5} yielded $t(11) = -3.12$, $p = 0.010$; and NO₂ emissions yielded $t(11) = -3.56$, $p = 0.004$. The comparisons capture varied traffic conditions, providing preliminary but strong evidence that NSGA-II dynamic pricing is an effective approach.

Table 5. Hourly pricing and environmental impact for minimum and maximum price candidates

Hour	Min price candidate				Max price candidate			
	Price (£)	NO ₂	PM _{2.5}	Flow	Price (£)	NO ₂	PM _{2.5}	Flow
07:00	10.07	87.61	16.52	787.98	14.95	82.11	9.76	674.93
08:00	10.12	87.61	16.52	788.75	14.96	82.11	9.76	677.13
09:00	10.08	87.61	16.52	660.33	14.98	82.11	9.76	668.49
10:00	10.21	87.61	16.52	784.28	14.82	82.11	9.76	657.07
11:00	10.07	87.68	16.60	662.14	14.50	81.74	9.31	690.48
12:00	10.09	87.68	16.60	787.64	14.87	81.74	9.31	695.03
13:00	10.07	87.68	16.60	656.68	14.92	81.74	9.31	765.24
14:00	10.09	87.68	16.60	788.26	14.80	81.74	9.31	656.91
15:00	10.01	87.02	15.80	668.44	14.93	81.90	9.50	688.94
16:00	10.19	87.02	15.80	784.95	14.83	81.90	9.50	655.56
17:00	10.09	87.02	15.80	787.50	14.94	81.90	9.50	657.36
18:00	10.18	87.02	15.80	784.86	14.83	81.90	9.50	659.27
Avg	10.11	87.44	16.31	745.15	14.86	81.92	9.52	678.87

5. Discussion and policy implications

5.1. Model generalization for wider

Static ERP models have proven their practical value after having been implemented in several major metropolitan cities across the world over the last few decades. Recent studies have continuously engaged with the benefits and limitations of dynamic road pricing schemes and showcased the potential of such road pricing schemes through various means. A few pioneers, including Singapore and Sweden, have already implemented road pricing schemes that are semi-dynamic. For example, the congestion tax imposed for vehicles entering Stockholm city centre changes at 30-min intervals between 6 am and 6 pm and is different for peak or off-peak season (Transportstyrelsen, n.d.). Though there is agreement on the contribution of the increasing availability of big data to better sustainability assessment of transportation policy (Liu and Dijk, 2022a), works focusing on both air pollution emissions and road congestion externalities are still limited, and even fewer papers provide a mechanism to enable multi-stakeholder participation.

To develop a more complete and flexible road pricing scheme, this study formulates the problem as a dynamic multi-stakeholder and multi-objective ERP problem, with the key objective of minimizing both air pollution and traffic congestion levels in the City of London, via optimizing the prices. To do so, a unique problem formulation and novel multi-stakeholder modelling approach have been constructed, and multi-objective evolutionary algorithms have been used to carry out the optimization. This study accounts for various aspects discussed in related works by gathering spatially and temporally fine-grained data for congestion and pollution and giving each factor a weight to put attention on the most important factors, which can be air pollution factors or congestion factors, and achieve long-term goals based on stakeholder parameters. Experimental results have shown significant reductions in the levels of air pollution and traffic congestion. Our novel methodology allows the model to be flexible and be implemented practically, with room for future extension. With this potential and some further exploration, the dynamic model could lead to greater societal and environmental impacts.

To further generalize the model to adapt to diverse urban contexts, future work can be carried out to incorporate multiple-criteria decision-making techniques to balance the perspectives from different stakeholders. For example, the use of electric vehicles (EVs) may reduce carbon emissions, noise pollution, and improve air quality. Active mobility, such as walking and cycling, may achieve the above, as well as introducing additional health benefits to the community. These additional considerations may

all be quantified and incorporated into our model as extensions. Different modes of transport may also be considered separately to capture their different impacts. However, particular considerations should be paid to the equity of road users, considering that the low-income group may be limited in their choices of the time and mode of transportation. As such, the pricing mechanism can be tailored to the specific needs of the city where it is applied. Depending on the interactions of various factors, multivariable non-linear regression may be introduced to better model a non-linear constraint for more complex Pareto fronts with the potential for better results. As pointed out by Guerrero and Castañeda, urban policies for sustainable development are usually multi-faceted and compete for limited resources (Guerrero and Castañeda, 2020), and the model should have the capability to be extended to include constraints and answer “what-if” questions for sensitivity analysis.

From the public finance perspective, the quantitative approach proposed in this paper also provides insight into the economic values of environmental externalities. The objective functions may also be updated to understand the potential for revenue generation through the dynamic ERP system and use the collected revenue to reinvest in the public transportation infrastructure, sustainable mobility initiatives, or other projects that benefit the community. During the time when the government may struggle to justify charges to the public, the quantitative approach supports the conversation to be more evidence-based.

Finally, a perfect dynamic ERP to send price signals in real time to drivers to induce behavioural changes requires an ecosystem to support its implementation. Such a system requires significant upfront investment and maintenance of the infrastructure, reducing the willingness for governments to adopt the scheme. A proxy of such a dynamic ERP mechanism may resemble the one implemented in Singapore, where the rates are subject to quarterly review and adjustment of school holiday months to accommodate the changes in travelling patterns (Development Asia, 2016). The payback period for the capital costs associated with implementing such a system is reported to be less than 5 years (D20 Long-Term Investors Club, 2020). Even when the government is ready to implement a fully real-time dynamic ERP, it may create information overflow for humans. It “results in the inability to react to such signals,” “effective, and largely autonomous, decision support systems are needed” (World Economic Forum Global Future Council on Urban Mobility Transitions, 2021) to support the implementation. To support policy acceptance and practical implementation despite initial hesitance, a phased adoption strategy can be considered. Governments may begin with semi-dynamic pricing updated at fixed intervals, allowing the public and infrastructure to adapt progressively. Parallel investments in communication technology, public engagement, and decision support systems can facilitate smoother transitions. This stepwise approach lowers upfront barriers, provides measurable results in each phase, and builds confidence in scaling towards fully dynamic AI-driven pricing schemes.

5.2. *Novelties and significance*

This study makes several novel contributions to the field of dynamic ERP systems and urban transport policy.

First, unlike traditional models that optimize solely for traffic congestion or travel time (Chen et al., 2018; Hanna et al., 2019), our research adopts a multi-stakeholder, multi-objective optimization framework. By explicitly incorporating stakeholder-defined priorities across traffic congestion, air quality, and pricing fairness, we extend prior approaches that typically fix stakeholder preferences or optimize a single objective (e.g., Schubert et al., 2022; Chiu et al., 2023). The integration of stakeholder-defined weights enables a more equitable and transparent policymaking process—crucial for democratic legitimacy and long-term adoption. As Hartley (2024) notes, articulating policy rationales in stakeholder terms fosters trust and enables more inclusive governance.

Second, while recent AI-driven pricing models—such as MAGT-toll (Lu et al., 2024) or deep RL-based tolling (Parishad et al., 2025)—focus primarily on real-time control, our model proposes a trained schedule approach. This allows for behavioural adaptation and longer-term pollution pattern management, offering a practical alternative for cities where real-time data infrastructure is limited. In contrast to No-Queue Road Pricing (Schubert et al., 2022), which sets tolls to maintain real-time target

speeds, our system generates predictable dynamic schedules that remain responsive but computationally stable, aligning more closely with regulatory planning horizons.

Third, in terms of spatial granularity, while earlier zone-based pricing studies (e.g., Garzon et al., 2022) vary zone shape based on air quality levels, our approach uses sub-regions of a fixed zone (the CCZ), each with its own dynamic schedule based on air pollution and traffic flow. This retains geographic simplicity while achieving a fine-grained responsiveness that better aligns with real-world congestion and pollution variations across space and time.

Fourth, our integration of evolutionary multi-objective optimization (NSGA-II), rather than purely RL or control-theoretic methods (Wang et al., 2021), enables solution diversity and trade-off visualization. This is particularly valuable for policy communication and stakeholder negotiation. While evolutionary optimization has been explored conceptually for congestion pricing, our work demonstrates its practical implementation in a real-world zone and quantifies its potential impact.

Finally, the model's architecture supports scalability across time horizons—from a single day to a full week or month—enabling planners to estimate cumulative effects on congestion, emissions, and revenues. Such extensions are vital to inform long-term investments, such as public transit expansion or climate mitigation funds, which rely on robust forward-looking estimates.

Together, these contributions represent a significant advancement in the design of data-driven, stakeholder-sensitive urban road pricing systems and pave the way for practical deployment using emerging smart city infrastructure such as GNSS-based On-Board Units (World Economic Forum Global Future Council on Urban Mobility Transitions, 2021).

5.3. Limitations and practical considerations

While the problem formulation showcases the potential and flexibility of such a model, allowing for substantial benefits in the reduction of congestion or pollution as well as beneficial changes in patterns, there are certain limitations faced by this paper.

Due to modelling for a practical scenario, the real-world source data were used, which was missing in significant amounts. Most of the missing air pollutant data can be filled in by using advanced imputation techniques or by using the relationship $PM_{2.5}$ and PM_{10} (see “Data pre-processing”). It is also treated by averaging the data of the AQMSs or CPs within the CCZ, followed by linear interpolation in the temporal dimension. Though location-based traffic analysis was not carried out in this study for verification and interpolation of the traffic data, it could be carried out in the future for analysing zones or cities at larger scales.

Another limitation is the influence of historic events. Only four price points were available in the history of the CCZ from 2008 onwards. The data from 2008 to 2011 are also influenced by the existence of the western expansion zone to the CCZ, which was closed in 2011. Gathering older data may be unreliable due to changes in human behaviour overtime and data immediately after 2019 (up till 2022) is influenced by the changes in human behaviour due to COVID-19. As new data continue to be gathered, the model can be made more robust and overcome this limitation. It may also be useful to consider building in a feedback loop so that data collected on an ongoing basis may be used to fine-tune the model continuously to keep up with economic and societal changes.

While the proposed model assumes automatic charge deduction and optimized route calculation, these are based on a pre-trained hourly price schedule rather than on-the-fly computation. This schedule allows vehicles to receive charges and suggested routes aligned with evolving pollution and congestion conditions, but without requiring continuous real-time updates. In practice, such full automation may not yet be widely deployed. However, this does not preclude practical implementation. A semi-dynamic system—where the price schedule is updated periodically (e.g., daily or weekly) using historical and forecasted pollution and traffic data—can serve as a feasible and effective interim solution. In this case, charges can be displayed to users in advance, and route suggestions can be calculated without relying on constant connectivity or edge computation. Prior research has demonstrated the viability of periodic toll

adjustments and traveler responsiveness to such systems (e.g., Chen et al., 2018; Gökasar and Bakioglu, 2018; Saharan et al., 2020). Moreover, the growing use of AI-based scheduling and demand forecasting models (Lu et al., 2024; Lukic Vujadinovic et al., 2024) suggests that infrastructure for adaptive mobility is advancing rapidly. The framework presented here is designed to be flexible and can accommodate either a semi-dynamic or fully automated deployment, depending on technological readiness.

Data availability statement. Replication data and code can be found in Zenodo: <https://doi.org/10.5281/zenodo.15839133>.

Author contribution. Conceptualization: T.K., S.K., V.O.K.L., J.C.K.L.; Data curation: T.K., S.K., Y.H.; Data visualization: T.K., S.K.; Investigation: T.K., S.K., V.O.K.L., J.C.K.L., Y.H., S.W.; Methodology: T.K., S.K., V.O.K.L., J.C.K.L.; Resources: T.K., S.K., Y.H.; Software: T.K., S.K.; Supervision: V.O.K.L., J.C.K.L.; Validation: T.K., S.K., V.O.K.L., J.C.K.L.; Visualization: T.K., S.K.; Writing—original draft: T.K., S.K., S.W.; Writing—review and editing: T.K., S.K., V.O.K.L., J.C.K.L., Y.H., S.W. All authors approved the final submitted draft.

Funding statement. This research was not supported by any grants.

Competing interests. None.

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

AI statement. AI tools (specifically OpenAI's ChatGPT) were used to support language editing, formatting, and refinement of the manuscript's text. The authors retain full responsibility for the content and conclusions of the work.

References

- Atkinson RW, Barratt B, Armstrong B, Anderson HR, Beevers SD, Mudway IS, Green D, Derwent RG, Wilkinson P, Tonne C and (2009) The impact of the congestion charging scheme on ambient air pollution concentrations in London, *Atmospheric Environment* 43(34), 5493–5500.
- Aurenhammer F (1991) Voronoi diagrams—A survey of a fundamental geometric data structure. *ACM Computing Surveys (CSUR)* 23(3), 345–405.
- Beevers SD, Carslaw DC, Dajnak D, Stewart GB, Williams ML, Kelly J and Kelly FJ (2016) Traffic management strategies for emissions reduction: Recent experience in London. *Energy and Emission Control Technologies* 4, 27–39.
- Brent DA and Gross A (2018) Dynamic road pricing and the value of time and reliability. *Journal of Regional Science* 58(2), 330–349.
- Chang MS and Hsueh CF (2006) A dynamic road pricing model for freeway electronic toll collection systems under build–operate–transfer arrangements. *Transportation Planning and Technology* 29(2), 91–104.
- Chen H, An B, Sharon G, Hanna J, Stone P, Miao C and Soh Y (2018) DyETC: Dynamic electronic toll collection for traffic congestion alleviation. *Proceedings of the AAAI Conference on Artificial Intelligence* 32(1). <https://doi.org/10.1609/aaai.v32i1.11337>.
- Chen, J and Hoek, G (2020) Long-term exposure to PM and all-cause and cause-specific mortality: A systematic review and meta-analysis. *Environment International* 143, 105974. <https://doi.org/10.1016/j.envint.2020.105974>.
- Chiu Y-C, Wang Z and Liu Y (2023) Dynamic marginal-cost tolling under arc-based user equilibrium. *Transportation Research Part B: Methodological* 167, 73–89.
- D20 Long-Term Investors Club (2020, November 4) Dynamic pricing algorithms for toll roads, <https://www.github.org/infra-structure-technology-use-cases/case-studies/dynamic-pricing-algorithms-for-toll-roads/> (accessed 7 September 2024).
- Deb K and Jain H (2013) An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation* 18(4), 577–601. <http://doi.org/10.1109/tevc.2013.2281535>.
- Deb K, Pratap A, Agarwal S and Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6(2), 182–197. <http://doi.org/10.1109/4235.996017>.
- Department for Environment Food & Rural Affairs of the Government of the United Kingdom (2024, April 30) Air quality statistics in the UK, 1987 to 2023, <https://www.gov.uk/government/statistics/air-quality-statistics> (accessed 7 September 2024).
- Development Asia (2016, May 18) The case for electronic road pricing, <https://development.asia/case-study/case-electronic-road-pricing> (accessed 7 September 2024).
- Egbo L and Marinov M (2023) The effect of congestion charging practices on air quality: A forecasting study- comparison between four cities. *JIPD*. <https://doi.org/10.24294/jipd.v7iX.2098>.
- Garzon SR, Reppenhagen M and Müller M (2022) What if air quality dictates road pricing? Simulation of an air pollution-based road charging scheme, *Journal of Urban Mobility* 2, 100018, <https://doi.org/10.1016/j.urbmob.2022.100018>.
- Gökasar I and Bakioglu M (2018) The effects of real-time traffic information on travel behavior. *Transport Policy* 72, 23–34.
- Guerrero OA and Castañeda G (2020) Policy priority inference: A computational framework to analyze the allocation of resources for the sustainable development goals. *Data & Policy* 2, e17. <https://doi.org/10.1017/dap.2020.18>.

- Hanna JP, Sharon G, Boyles SD and Stone P (2019) Selecting compliant agents for opt-in micro-tolling. *AAAI Conference*. <https://doi.org/10.1609/aaai.v33i01.3301565>.
- Hartley K (2024) State-society relations and government technology: A survey of public awareness and communication in Hong Kong. *Data & Policy* 6, e18. <https://doi.org/10.1017/dap.2024.15>.
- Joksimovic D, Bliemer MCJ, Bovy PHL and Verwater-Lukszo Z (2005) Dynamic road pricing for optimizing network performance with heterogeneous users. *Proceedings. 2005 IEEE Networking, Sensing and Control 2005*, 407–412. <http://doi.org/10.1109/icnsc.2005.1461225>.
- Kaistha T, Kaistha S and Han Y (2025) Dataset for AI-Driven Dynamic Electronic Road Pricing to Reduce Air Pollution and Traffic Congestion, Zenodo, V1 [dataset]. <https://doi.org/10.5281/zenodo.15839133>.
- Liu X and Dijk M (2022a) The role of data in sustainability assessment of urban mobility policies. *Data & Policy*, 4. <http://doi.org/10.1017/dap.2021.32>.
- Lu J, Hong C and Wang R (2024) MAGT-toll: A multi-agent reinforcement learning approach to dynamic traffic congestion pricing. *PLoS One* 19(11). <http://doi.org/10.1371/journal.pone.0313828>.
- Lukic Vujadinovic V, Damjanovic A, Cakic A, Petkovic DR, Prelevic M, Pantovic V, Stojanovic M, Vidojevic D, Vranjes D and Bodolo I (2024) AI-driven approach for enhancing sustainability in urban public transportation. *Sustainability* 16(17), 7763. <http://doi.org/10.3390/su1617776>.
- Munkácsy A, Földes D, Miskolczi M and Jászberényi M (2024) Urban mobility in the future: Text analysis of mobility plans. *European Transport Research Review* 16(1). <http://doi.org/10.1186/s12544-024-00649-x>.
- Parishad N, Yildirimoglu M and Hickman M (2025) Congestion pricing in multi-modal networks: An application of deep reinforcement learning. *Transportation Research Part C: Emerging Technologies* 177, 105166. <http://doi.org/10.1016/j.trc.2025.105166>.
- Ribeiro PJ, Dias G and Mendes JF (2024) Decarbonizing urban mobility: A methodology for shifting modal shares to achieve CO2 reduction targets. *Sustainability* 16(16), 7049. <http://doi.org/10.3390/su16167049>.
- Saharan S, Bawa S and Kumar N (2020) Dynamic pricing techniques for intelligent transportation system in smart cities: A systematic review. *Computer Communications* 150, 603–625. <http://doi.org/10.1016/j.comcom.2019.12.003>.
- Schubert D, Sys C, Vanelander T and Rouboutsos A (2022) No-queue road pricing: A comprehensive policy instrument for Europe? *Utilities Policy* 78, 101413. <http://doi.org/10.1016/j.jup.2022.101413>.
- Seada H and Deb K (2015) U-NSGA-III: A unified evolutionary optimization procedure for single, multiple, and many objectives: Proof-of-principle results. *Lecture Notes in Computer Science*, 34–49. http://doi.org/10.1007/978-3-319-15892-1_3.
- Sharon G, Levin MW, Hanna JP, Rambha T, Boyles SD and Stone P (2017) Network-wide adaptive tolling for connected and automated vehicles. *Transportation Research Part C: Emerging Technologies* 84, 142–157. <http://doi.org/10.1016/j.trc.2017.08.019>.
- Soylemezgiller F, Kuscü M and Kilinc D (2013) A traffic congestion avoidance algorithm with dynamic road pricing for smart cities. *IEEE 24th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC) 2013*, 2571–2575. <http://doi.org/10.1109/pimrc.2013.6666580>.
- Tan Y, Sun Z, Zhu B, Qin X, Zhao Y and Wang X (2024) Minimize population exposure to vehicle-generated emissions by road pricing. *Transport Policy* 148, 15–30. <http://doi.org/10.1016/j.tranpol.2023.12.025>.
- Teodorović D and Edara P (2007) A real-time road pricing system: The case of a two-link parallel network. *Computers & Operations Research* 34(1), 2–27. <http://doi.org/10.1016/j.cor.2005.02.041>.
- The Joint Air Quality Unit of the Government of the United Kingdom (2017, July) UK plan for tackling roadside nitrogen dioxide concentrations - An overview, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/633269/air-quality-plan-overview.pdf (accessed 7 September 2024).
- Transport for London (n.d.) Congestion charge zone map, <http://content.tfl.gov.uk/congestion-charge-zone-map.pdf>.
- Transportstyrelsen (n.d.) Congestion tax in Stockholm - hours and amounts, <https://www.transportstyrelsen.se/en/road/road-tolls/Congestion-taxes-in-Stockholm-and-Goteborg/congestion-tax-in-stockholm/hours-and-amounts-in-stockholm/> (accessed 7 September 2024).
- Verity W, Poirier Stephens Z, Lin H, Ottoni CA, Bourgeois E, Kestens Y, Fuller D, Manaugh K and Winters M (2025) Health and equity considerations in policy and practice related to sustainable transportation interventions in four Canadian cities. *Journal of Transport & Health* 41, 102009. <http://doi.org/10.1016/j.jth.2025.102009>.
- Vosough S, de PA and Lindsey R (2022) Pricing vehicle emissions and congestion externalities using a dynamic traffic network simulator. *Transportation Research Part A* 161, 1–24.
- Wang JYT, Ehrgott M, Dirks KN and Gupta A (2014) A bilevel multi-objective road pricing model for economic, environmental and health sustainability. *Transportation Research Procedia* 3, 393–402.
- Wang X, Jin W-L and Yin Y (2021) A control theoretic approach to simultaneously estimate average value of time and determine dynamic price for high-occupancy toll lanes. *IEEE Transactions on Intelligent Transportation Systems* 22(11), 7293–7305. <http://doi.org/10.1109/tits.2020.3007160>.
- WHO (2005) WHO air quality guidelines for particulate matter, ozone, nitrogen dioxide and sulfur dioxide, <https://www.who.int/publications/i/item/WHO-SDE-PHE-OEH-06-02>

- World Economic Forum Global Future Council on Urban Mobility Transitions** (2021) Sustainable road transport and pricing (white paper), Retrieved from https://www3.weforum.org/docs/WEF_Sustainable_Road_Transport_and_Pricing_2021.pdf
- Yildirim MB and Hearn DW** (2005) A first best toll pricing framework for variable demand traffic assignment problems. *Transportation Research Part B* 39(8), 659–678.
- Zong F, Zeng M and Li Y-X** (2024) Congestion pricing for sustainable urban transportation systems considering carbon emissions and travel habits. *Sustainable Cities and Society* 101, 105198. <http://doi.org/10.1016/j.scs.2024.105198>.