Dynamic Congestion Pricing for Urban Emission Exposure

Vishal Anand, Jort Dijkstra, Jack Hansen, Antoni López, and Ondrej Vlk

ME44305 - System Analysis and Simulation
Mechanical Engineering - Multi-Machine Engineering

Delft University of Technology
Delft 2628 CD, Netherlands

June 19, 2026

Abstract. This project proposes a dynamic urban congestion pricing system based on emissions, designed to reduce population exposure to traffic-related pollutants while improving traffic efficiency. The city is divided into pricing zones whose tolls adapt in real time according to local emission levels, traffic conditions, and other contextual factors. Unlike traditional congestion charging schemes, the proposed approach combines traffic management, environmental protection, and public health objectives within a single adaptive framework. To evaluate its effectiveness, a simulation model based on real-world data from Dublin was developed, incorporating driver behavior, vehicle emissions, traffic patterns, and dynamic route choice. The results indicate that adaptive pricing can reduce emissions in targeted urban areas by encouraging route redistribution and modifying travel behavior, highlighting its potential as a flexible tool for sustainable urban mobility and air quality management.

Keywords: Urban emissions · Dynamic pricing · Discrete-event simulation

1 Conceptual System Model and Innovation

1.1 Introduction

High levels of vehicular emissions in urban areas have been associated with adverse health consequences. Local reductions in traffic congestion have been linked to significant improvements in health outcomes, highlighting the severe biological toll of unchecked emissions [1]. Various methods for mitigating these emissions have been pursued, ranging from widespread adoption of electric vehicles to pedestrianized streets.

Simultaneously, public transit networks have struggled to regain ridership post-pandemic, as remote work policies and changing lifestyles keep people at home. The post-COVID societal reorganization has fundamentally altered urban mobility patterns, making public transit recovery highly unlikely without systemic intervention [4]. Low ridership has deepened funding shortfalls, and many cities are considering innovative methods to fund public transit.

Externalities are costs incurred by an activity that remain unpaid. The cost of urban emissions is an externality that drivers do not currently pay for. Congestion pricing is a popular implementation that charges vehicles entering predefined zones. It seeks to reduce congestion in urban centers while raising money, which is most frequently allocated to public transportation.

The specific fee charged to each vehicle under congestion pricing is a politically charged issue: it must be high enough to deter unnecessary trips without constituting a regressive tax on drivers who have no other option. However, research suggests that while congestion pricing might appear to have equity implications, alternative transportation funding mechanisms can actually place a disproportionately heavier burden on low-income residents [6].

There is currently no consensus on the most fair and effective method of congestion pricing, and the few cities where it is enacted have substantially differing implementations. For example, the area-based pricing model used in London demonstrates how charging structures must be carefully tailored to a city's specific geographical and economic landscape to balance traffic reduction with public acceptance [5].

This paper explores an implementation that seeks to optimize congestion pricing, creating the maximum benefit to urban residents while remaining fair to drivers.

1.2 System Description and Innovation

Description of the system In order to cause an improvement in the previously described issue and based on the already existing literature, an innovative solution is proposed: Dynamic Congestion Pricing.

This solution is based on the discretization of the urban area into different pricing zones based on emission levels. Prices will fluctuate to ensure acceptable emission levels while improving traffic congestion, having an impact both on emission reduction and routing optimization. The driver will be charged according to the chosen route (i.e., depending on the pricing zones along the way).

Pricing in the different zones will be defined dynamically, depending mainly on emission levels and traffic levels, but a scheduling factor may apply as well depending on the situation (e.g., avoiding traffic near schools when kids are being dropped off, since they are a vulnerable population group highly affected by exhaust gases and other emissions). These changes in pricing will encourage drivers to take the most efficient routes, emission-wise. It must be noted that, throughout this study, it will be assumed that every stakeholder has perfect information about the system, for instance through mapping apps in the specific case of drivers.

Innovation The innovation in the proposed system does not come from the congestion pricing itself, as it is not entirely new, but rather from how the system redefines the way the pricing is used within the urban mobility system and what it is intended to optimize. Existing congestion pricing models commonly focus on using fixed zones, larger time windows, and general traffic demand, usually with the aim of managing congestion or generating revenue. On the other hand, the proposed concept treats congestion pricing as an adaptive control tool used to specially reduce urban emission exposure, able to adjust to various changing demands in intended goals, while still considering traffic efficiency and fairness.

This innovative aspect comes from the fact that proposed solution is explicitly multi-objective. This means that rather than focusing only on congestion reduction or revenue generation, the system aims to reduce harmful exposure in the places and times where it matters most, focusing on also improving the health of the population.

The system does this by dividing the city into pricing zones whose charges can change in real time based on current conditions. This includes the local traffic conditions, emission levels, and where relevant, scheduling factors. For instance, drop-off periods near schools, as mentioned earlier. This makes the system more responsive to actual environmental conditions instead of relying only on static city boundaries or average congestion patterns.

The main innovative element therefore lies in the combination of several features into one integrated system: zone-based pricing, the use of real-time and estimated emission and traffic data, environmental sensitivity (in how zones are defined) and behavioral adaptation by drivers through route choice. As a result, pricing becomes a more targeted and flexible intervention, which allows the system to respond more precisely to local urban conditions based on the active demands and goals of the people.

Depending on the point of view, it can be said that this proposal improves current systems by aligning traffic management more directly with health and sustainability objectives, which helps improve the quality of life of the general public. This means the proposal is best understood as a system-level and incremental innovation, rather than as a completely new invention. This is due to the fact that the purpose of the dynamic congestion pricing extends beyond traffic management and into healthcare and builds on top of existing ideas to tackle traffic and urban emissions.

Main challenges Like any innovative proposal to solve a real-life issue, the design and implementation of this system pose a few relevant challenges. These can be broadly classified into three main categories: Simulation, Implementation, and Ethical.

Simulation or modelling challenges refer to the difficulties encountered when trying to adapt a complex real-life issue into a simple mathematical (or math-based) model. The modelling effect leads, necessarily, to an accuracy gap between reality and simulation which, while accepting its existence, must be reduced as much as possible.

As mentioned before, this study will be conducted considering a perfect information system, which will allow drivers to optimize their routes beforehand. This routing optimization will be, in fact, the optimization problem each driver has to solve to get from their origin to their destination.

An agent component is, therefore, necessary to elaborate the simulation. During the execution process, agents (drivers) will be generated with a defined set of parameters, from physical-deterministic variables like the type of vehicle they drive and the departure and destination points, to complex stochastic variables, mostly related to behavior (e.g., willingness to pay extra). Once the agent profile is generated, each driver optimizes their route independently. Taking that into consideration, the pricing model tunes prices in every zone so that emissions are minimized.

Implementation challenges are related to the physical integration of the simulated system into a real-life environment (although they also have effects on the mathematical model). These challenges include the selection and utilization of sensors, both directly related to the system (cameras to read license plates, zoning delimitation) and indirectly related (emission level measurement, traffic control).

A deeper stakeholder analysis must also answer who is controlling the system and under which parameters prices are set. In fact, the relationship between zoning prices and expected emission levels must also be defined, as well as how frequently these prices are or can be changed.

Ethical challenges are also relevant in any transportation system, especially those which directly affect people's lives and their data. Many classical debates can be thought of when presenting this kind of system, all of them revolving around the concept of trading off individual freedom or control for a greater good. The implementation of this project would imply a certain loss of agency in terms of the decision-making process of drivers, who would be nudged into following a certain directive by a higher entity. They would also be giving away more control over their data due to the importance of tracking for this system.

On the other hand, if emissions are lower, the overall health of the population would increase, especially affecting vulnerable population groups like children. Also, traffic would be reduced, which would cause a significant improvement in the well-being of citizens. Both consequences would directly impact the economy of the region, since the need for healthcare would be reduced and traffic accidents (which cause significant monetary losses every year) would also decrease.

The main questions are whether the benefits outweigh the drawbacks, which will depend on every stakeholder, and how this myriad of opinions is reflected in policymaking. However, this ethical discussion is, although considered, outside the scope of the course.

1.3 Model Conceptualization

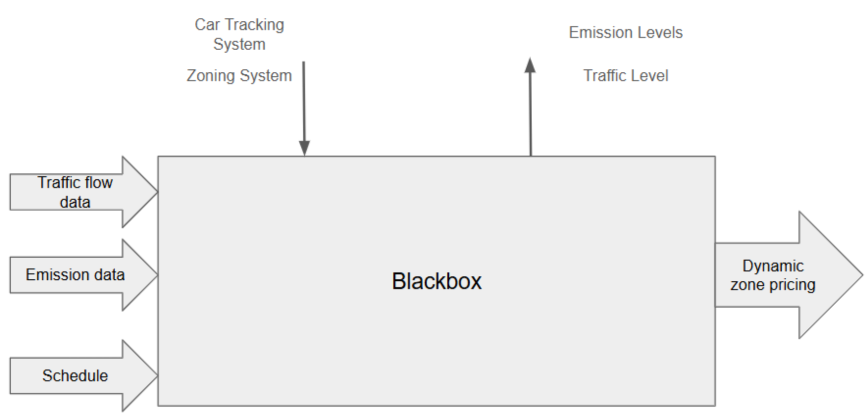

Fig. 1. Blackbox model figure.

Blackbox Model In this project, the dynamic congestion pricing system is represented as a blackbox that interacts with its surroundings or environment, i.e. through its measurable input and observable output. Added to the model are a set of requirements and key performance indicators that influence the required output.

On the input side, the blackbox receives several classes of information from the environment. Firstly, it is provided with real-time and historical traffic data, such as vehicle counts, speeds, and densities per road segment or zone, which are obtained primarily from traffic cameras or other car tracking systems. Secondly, it receives emission-related information, provided as direct measurements from environmental sensors. Thirdly, it is provided with constant data such as road network topology, zoning definitions, and the location and scheduling data of sensitive areas such as schools and hospitals.

Inside the blackbox, the different inputs are processed by a certain mechanism which in turn controls the pricing. The mechanism first forms an internal representation of the current and near-future state of the system, combining traffic and emission information. Based on this state, it selects a set of congestion prices for all zones that aims to reduce emission exposure, while simultaneously avoiding excessive congestion.The blackbox perspective allows us to reason and get and in-depth idea about the system and its logic without committing to a specific implementation.

On the output side, as mentioned in the previous paragraph, the blackbox determines the set of congestion prices which subsequently leads to reduced emission levels in critical zones where it is measured to be high. These prices change over time in response to the observed and predicted state of the system, thereby helping the drivers to choose routes and departure times that reduce emissions in sensitive areas.

In this way, the blackbox model acts as a bridge between the problem description and the more detailed conceptual simulation model.

CATWOE Model To better analyze the different actors and factors implicated in the proposed system, the CATWOE model is used. This analysis tool helps conceptualize complex systems (similarly as the blackbox model) with a high emphasis in the stakeholders.

- **C**ustomers: The customers of the innovation are the urban residents, as they benefit from the reduced emission levels.

- **A**ctors: Actors are the stakeholders of the innovations, the entities that actually carry out the transformation. In this case the actors have been identified as city planners, policy makers, engineers and traffic participants. The first three are actors because they will implement the dynamic congestion pricing system and zones. The traffic participants are also identified as actors, because their decision whether or not to enter an emission zone influences the outcome of the transformation.

- **T**ransformation: Transformation is the change that is caused by the implementation of the system. In this case this is reduced emission exposure.

- **W**orld view: The bigger picture of what actors want to achieve. For the dynamic congestion pricing model, this would be to improve health conditions of residents in the urban area.

- **O**wner: The owner is the entity that wants to implement the transformation. This would be the municipality.

- **E**nvironmental: These are factors that could hinder the transformation. This could for example be political opposition or budget constraints.

This analysis establishes a new, different framework (complementing the blackbox analysis) to use as a reference for the next stages of the study, focused in the design and implementation of the proposed system.

2 Conceptual Simulation Model and Experiment Plan

2.1 Components Definition

The simulation system can be divided into the following components:

- **Driver**
  The driver has a origin, a destination, an emission rate (vehicle type) and a route. The behavior of the driver is modeled based on two aspects, the willingness to pay and their time sensitivity. The willingness to pay influences the route taken based on the prices of the zones it passes through. Their time sensitivity influences the route taken based on the time it takes to complete the route. The driver can also cancel their trip if they are not willing to pay the price of the route or are not willing to drive for that duration.

- **Pricing Controller**
  The pricing controller is built to synthesize all of the data from the emission and traffic detectors and output pricing for each zone. The controller is built around a static map layer, which has predefined zones. For each zone, the pricing function has a critical emissions threshold beyond which congestion pricing kicks in. Then, the price levels increase to deter drivers, up to a certain maximum price. The pricing controller is designed to limit emissions zone by zone, rather than across the city as a whole.

- **Data collectors**
  Sensors collecting data to feed the simulation. There are 2 types: emissions and traffic detectors. They can be modeled together for simplicity, since they follow the same process.

  - **Emissions detector**
    The emissions detector constantly reads the emission levels with the help of sensors. This data is used to determine the pricing of the different emission zones. This influences the routes taken by the drivers based on their willingness to pay. If the emission detector detects emission level above the maximum limit in a particular zone, it alerts the controller to trigger a price hike.

  - **Traffic detector**
    The traffic detector collects traffic data. It tracks how crowded the roads are. This data influences the duration of routes and can therefore influence the route decision of the drivers. And also the data helps the controller decide a price increase is needed to clear the traffic or not which results in the decrease of emission levels.

- **Driver generator**
  The driver generator generates the numbers of drivers and the location of

the drivers. It basically feeds in the necessary inputs such as the destination, time of departure etc to a driver. By generating many unique drivers at different times of the day, the generator allows how the pricing controller handles everything, from a quiet morning to a busy rush hour in a day.

2.2 Simulation Description

The simulation process is the following:

- **Initialize sensors**
  - Initialize emissions detectors
  - Initialize traffic detectors

  The first step of the simulation is to initialize the "passive" components (components that don't directly participate in the decision making) so they are ready to collect data.

- **Initialize Pricing Controller** The system generates the map and the different pricing areas. It also has to generate a first set of prices for the drivers to have enough information to decide which route are they taking.

- **Generation of drivers**
  - **Start generation:** Once every component is initialized, drivers can be generated according to stochastic functions which define inter arrival time and all the other attributes of the driver agent.
  - **Calculate the best route:** With the agent attributes and the pricing information, each driver computes their optimal route.
  - **Repeat process:** The process of generation is repeated until the simulation is over.

- **Driving process** With the initial information, vehicles start their trip. They drive until they arrive to the next node and then recalculate the route, according to new prices. If the new route is deemed better, it is adopted. Otherwise, old route is kept. This process is repeated until destination is reached.

- **Data collection** Once drivers start populating the simulation, the traffic and emission sensors collect information and send it to the price controller.

- **Price update** With the collected data, the price controller updated the prices corresponding to every area and updates that information so that every driver in the system can fetch it whenever needed.

– **Repetition** The process of data collection is repeated cyclically as long as the simulation is running. The price update is done in discrete time steps in order to optimize the simulation performance.

– **Data analysis** The simulation generates data during its run time which can be later analyzed.

The whole simulation process can be represented in a swimlane diagram, as shown in Figure 2.

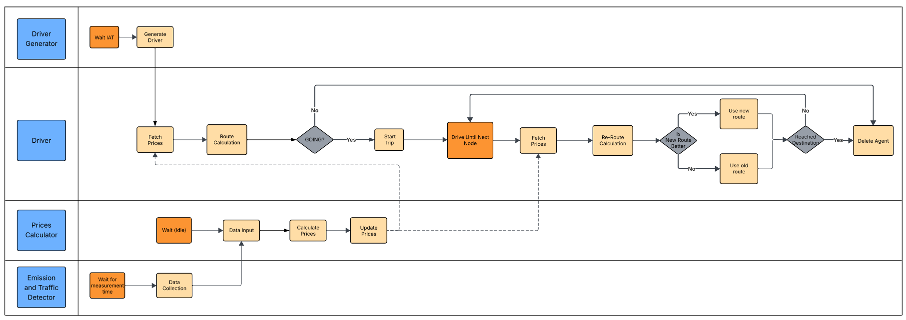

Fig. 2. Swimlane diagram of the simulation.

2.3 Process Description Language

The following Process Description Language (PDL) outlines the structure of the simulation. It details the precomputation of driver schedules using traffic data and the engine that handles environmental decay, dynamic cordon pricing, and agent routing decisions.

```
dangerThreshold = constant
emissionsGrid = empty
zoneTolls = [0.0, 0.0, 0.0, 0.0]
futureEmissionsQueue = empty

// Agent Parameters

VehicleTypes = TDiscreteSet
WillingnessToPay = IncomeDecileMapping
ValueOfTime = DerivedFromIncome
MapNodes = TDiscreteSet

// Precomputes all agents upfront based on traffic data and rush hour
curves.
```

```
TScheduleGenerator
process
  masterSchedule = empty
  repeat for t = 0 to TotalSimSeconds
    targetPop = CalculateTargetPopulation(SCATS, RushHourRatio)
    if currentAgents < targetPop
      newDriver = new TDriver
      newDriver.Origin = MapNodes.sampleOrigin()
      newDriver.Destination = MapNodes.sampleDestination()
      newDriver.DepartureTime = t
      newDriver.WTP = WillingnessToPay.sample()
      newDriver.VOT = ValueOfTime.sample()
      masterSchedule.append(newDriver)

// Main execution loop: Evaluates environment and routes agents minute-by-minute.

TSimulationEngine
process
  repeat for m = 0 to 1440
    // Update environmental state and spatial volumes
    emissionsGrid.ApplyDecay()
    emissionsGrid.AddEmissions(futureEmissionsQueue[m])

    // Dynamic 4-Tier Cordon Pricing based on zone thresholds
    zoneTolls = CalculatePrices(emissionsGrid, dangerThreshold)

    // Route agents scheduled to depart at this current minute
    for each driver in masterSchedule[m]
      driver.Route = DijkstraShortestPath(driver, zoneTolls)
      financialCost = driver.Route.TollPaid + driver.Route.DistanceCost

      // Agent decision logic based on routing conditions
      if financialCost <= driver.WTP
        SaveItinerary(driver.Route)
        futureEmissionsQueue.Schedule(driver.Route.Emissions)
      else
        driver.Reject()

//Configures and starts the simulation model

INITIALIZATION
ScheduleGenerator = new TScheduleGenerator
ScheduleGenerator.Start
SimulationEngine = new TSimulationEngine
SimulationEngine.Start
```

3 Data Collection and System Analysis

3.1 Introduction to data collection

In order to conduct and analyze the simulation different sets of data are considered. On the one hand, there is the data we need to generate the simulation. This includes geographical information to create the simulation environment and different datasets to generate the agents that will populate it.

The map of the area we are representing, includes roads and drivable streets, residential areas and business hubs. In our case, we chose to simulate traffic in the city of Dublin, because it is in the top three most congested cities in Europe, with the number of cars rapidly growing, which causes many negative externalities for its population. This information is combined with traffic data to model, for instance, how many drivers are per day or what type of vehicles they are driving.

To generate drivers and their behavior, income distribution in the city is analyzed, and a willingness to pay is modeled based on that.

On the other hand, there is the data needed to verify and evaluate the simulation. Traffic data can also play a role here, since it can be studied if simulated vehicle flows resemble the ones found in real life. However, emissions data are the most relevant. Knowing approximate emissions and public health guidelines, we can easily observe if the proposed solution fulfills its primary goal of reducing total emissions.

To do that, we will also use the generated data output from our own simulation model.

3.2 Collected data

Geographical data To generate the simulation's environment, a map of the city of Dublin has to be generated. To do so, data is obtained from Open Street Map. This data is divided into 3 different sets: roads, homes and offices.

Roads and drivable streets are modeled so that it is clear which paths are available for vehicles to follow. To process this dataset, streets that are not drivable (or that are not relevant because they are relatively small) are filtered out. This is done to create more direct paths (e.g., avoiding dead ends) and simplify calculations, making the simulation less computationally demanding.

Then, we need to define different sets of nodes representing homes (where people live) and offices (where people work). This simplification is done to reduce the infinite number of origin and destination nodes to a finite set representing the total number of trips.

However, this data needs to be further processed. The original datasets provide a heat map of residential areas and work hubs (which are clearly accurate based on our knowledge of the city). This heat map is then simplified, and 1500 representative nodes are generated for each dataset.

Traffic data To create an accurate simulation, traffic data should be used to ensure the car's movement follows realistic traffic patterns rather than randomly generating vehicles. Knowing the car density in different parts of the city at different times is important, as larger clusters of cars create larger emission zones. With only the map data, there is no real estimation of the intensity of cars that should be entering certain locations. Furthermore, knowing the difference in the movement during, for instance, the peak morning rush hour versus midday, the vehicle generation in the simulation can be changed and reflect a closer approximation to the actual ratio of volume of vehicles in the city.

To achieve this, SCATS traffic volume data from Dublin City Council is used. There are two datasets:

- The first data set contains hourly summed vehicle volume from detectors at signalized junctions. These junctions have multiple detectors (for instance, one for each street). Hence, the detector counts can be summed for each site and hour to estimate the total vehicle flow at that junction. [2].
- The second data set contains the location of each SCATS junction. This includes the latitude and longitude values. [3].

By combining these data sets, the hourly traffic counts can be placed on the Dublin map and compared at different times and between different locations in Dublin. To make the data usable in the simulation, the Dublin area is divided into square grid boxes of 4 km × 4 km. Each grid box is defined by latitude and longitude boundaries:

$$lat_{\min}, \quad lat_{\max}, \quad lon_{\min}, \quad lon_{\max}$$

A SCATS site belongs to a grid box if:

$$lat_{\min} \leq lat \leq lat_{\max}$$

and

$$lon_{\min} \leq lon \leq lon_{\max}$$

Each grid box is given a numerical `grid_id`. This allows the simulation to match a vehicle location to a specific grid area. The grid setup is useful because traffic is not only different between individual junctions, but also between larger city regions. Therefore, using grid boxes gives a more general traffic intensity input for the simulation. It is important to note that a square in the grid only exists if there are existing SCATS junctions within its area. The grid layout and numbering can be seen in the figure below:

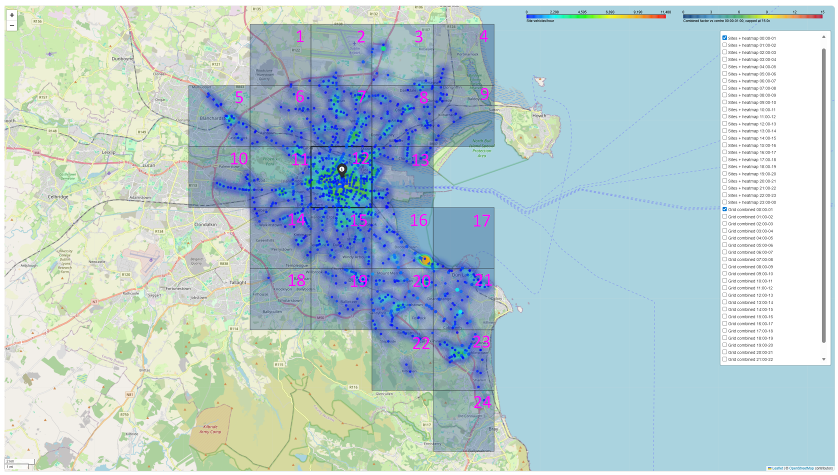

Fig. 3. Grid-based traffic intensity map for Dublin at time 00:00-01:00.

For each hour, the SCATS detector counts are first summed at each SCATS junction, which gives the hourly traffic volume at each junction. Then, for every grid box, the average vehicles per hour per SCATS junction is calculated as:

$$G_t = \text{mean vehicles/hour per SCATS site in the grid at hour } t$$

The mean value is used instead of the total sum because different grid boxes contain different numbers of SCATS sites. If the total sum was used, a grid box with many SCATS junctions could appear busier only because it has more measurement points. Using the mean gives a more balanced estimate of the typical traffic intensity inside each grid box.

One very important limitation of the SCATS data is that the counts do not represent the exact number of unique vehicles in the city. The same vehicle can be detected at multiple junctions along its route. Therefore, the data should not be used as an absolute count of all vehicles in Dublin. However, it is still useful as an indicator of relative traffic intensity. A junction or area with higher detector counts can be interpreted as busier by having more vehicle movement than a junction or area with lower counts.

Because of this limitation, the data is processed into relative traffic factors rather than absolute vehicle numbers. These factors compare traffic intensity between grid areas and time periods. This makes the data more suitable for the simulation, because vehicle generation can be scaled according to relative traffic demand without claiming to reproduce the exact number of real vehicles.

Reference grid and factor calculation To convert the grid values into simulation inputs, we must define some factors to show how the traffic changes with time and location. To do this, a reference grid and reference time are selected. The reference grid is the center grid square (square 12), and the reference time is the start of the day 00:00-01:00. We can then define:

$$C_0 = \text{center reference grid value at } 00{:}00 - 01{:}00$$

$$G_0 = \text{same grid value at } 00{:}00 - 01{:}00$$

$$G_t = \text{same grid value at the current hour}$$

We can then define how each grid square compares to the center grid square at the reference time. This becomes the spatial baseline factor, and is computed as:

$$S = \frac{G_0}{C_0}$$

Next, we can define how each grid square at the current hour compares to the same grid square at the reference time. This gives us the temporal factor, which can be computed as:

$$T = \frac{G_t}{G_0}$$

Finally, we can take these factors to get the combined factor:

$$F = S \times T$$

Substituting the two factors gives:

$$F = \frac{G_0}{C_0} \times \frac{G_t}{G_0}$$

which simplifies to:

$$F = \frac{G_t}{C_0}$$

This combined factor can then be used as the traffic generation multiplier in the simulation. For example, if a grid square has a combined factor of 4 at 07:00-08:00, then vehicle generation in that grid square and hour can be set to four times the reference traffic level (the center grid square at the reference time). This makes the simulation generate more vehicles in areas and time periods with higher observed traffic intensity.

The figure below shows a visual of the difference in traffic intensity and the combined factors of each grid square between the reference and the 07:00-08:00 period.

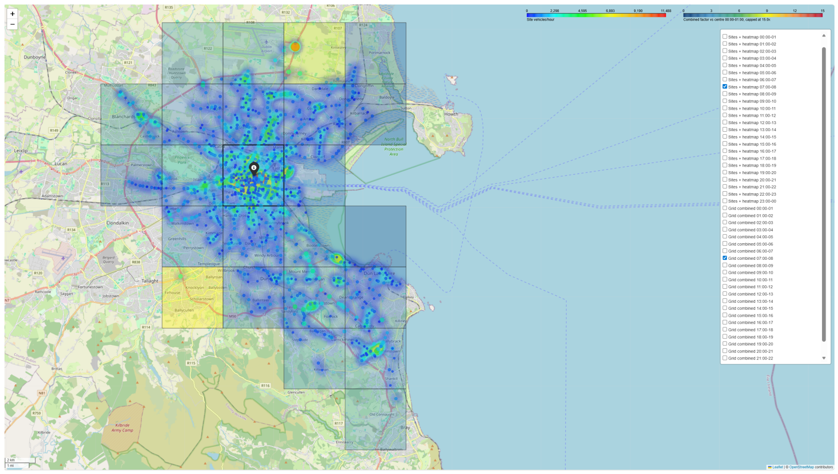

Fig. 4. Grid-based traffic intensity map for Dublin at time 07:00-08:00.

CSV output for the simulation Once the combined factors were computed for each grid square at each hour interval of the day, a CSV file was created to be used for the simulation. Each row represents one grid box during one hourly time interval. The CSV contains:

- `grid_id`
- `hour_start`
- `hour_end`
- `hour_period`
- `combined_factor`
- `lat_min`
- `lat_max`
- `lon_min`
- `lon_max`

Interpretation and validation The resulting heatmaps and grid factors show that traffic intensity changes clearly throughout the day. The map at 00:00-01:00 shows relatively low traffic activity, while the 07:00-08:00 map shows stronger traffic clusters around major roads and busy junction areas. This pattern is reasonable because 07:00-08:00 corresponds to the morning rush period, when more vehicles are expected to enter and move through the city. The results, therefore, provide a useful data input for the simulation. Instead of generating vehicles randomly across the map, the simulation can generate more vehicles in grid squares at specific hours with higher combined factors. This helps the simulated traffic pattern better reflect real traffic intensity in Dublin.

Vehicle emission data The whole idea of the how emission levels can be tracked is the simulation is to convert vehicle-level NO_2 emissions into ambient air concentration values (µg/m^3), and compare those values against the applicable regulatory threshold of 40 µg/m^3 set by EU Directive 2008/50/EC. The approach is used in COPERT v5 emission factors, real-world measurement studies, and WHO/EU air quality standards.

Why NO_2?

Nitrogen dioxide (NO_2) is selected as the primary pollutant in this simulation for three reasons:

- Direct health impact: NO_2 is a strong airway irritant linked to respiratory mortality, asthma hospitalisation, and long-term lung damage, particularly in children and the elderly.
- Regulatory significance: It is the pollutant most commonly monitored at roadside.
- Diesel vehicle sensitivity: NO_2 concentrations respond directly to diesel traffic density.

Vehicle emissions

Emission factors are taken from COPERT v5 (Computer Programme to Calculate Emissions from Road Transport), the EU-standard vehicle emissions model coordinated by the European Environment Agency (EEA). COPERT reports total NOx (NO + NO_2) emission factors in g/km. To isolate the direct NO_2 fraction, a vehicle-type-specific fNO_2 factor (fraction of NOx emitted directly as NO_2) is applied. These fractions are sourced from Carslaw & Beevers (2005) and validated against PEMS (Portable Emissions Measurement System) studies on real-world Euro 6 vehicles.

The formula applied per vehicle per road segment is:

$$\text{Direct NO}_2\text{ Emission (g/km)} = \text{NO}_x\text{ (g/km)} \times f_{\text{NO}_2}$$

Table 1. Direct NO_2 emission factors by vehicle type.

Vehicle Type	**Euro Standard**	**NO_x (g/km)**	**f_{NO_2} (%)**	**Direct NO_2 (g/km)**
Petrol Car	Euro 6	0.06	∼5%	0.003
Diesel Car (old)	Euro 6	0.50	∼30%	0.150
Diesel Car (new)	Euro 6d	0.08	∼20%	0.016
Diesel Van (LGV)	Euro 6	0.60	∼35%	0.210
HGV / Truck	Euro VI	0.40	∼10%	0.040
Bus	Euro VI	0.55	∼12%	0.066
Electric (EV)	—	0.000	—	0.000
PHEV	—	—	—	0.010–0.040
Full Hybrid (HEV)	—	—	—	0.040–0.060

Now we find the concentration value with the help of the direct NO_2 values of each vehicle. This is obtained with the help of the formula below:

$$\text{Concentration } (\mu g/m^3) = \frac{\text{Mass emitted } (\mu g)}{\text{Zone air volume } (m^3)}$$

where:

$$\text{Mass emitted } (\mu g) = \text{emission\_rate } (mg/min) \times \text{travel\_time } (min) \times 1000$$

$$\text{emission\_rate } (mg/min) = \frac{\text{NO}_2 \ (g/km) \times \text{urban speed } (km/h)}{60}$$

$$\text{Zone air volume } (m^3) = \text{Zone road length } (m) \times \text{Street width } (m) \times \text{Mixing height } (m)$$

The emission factors shown in Table 1 indicate that diesel-powered vehicles contribute disproportionately to total NO_2 emissions compared to petrol and electric vehicles. For example, older diesel cars produce approximately 0.150 g/km of direct NO_2 emissions, which is around 50 times greater than a petrol Euro 6 vehicle (0.003 g/km). Similarly, diesel vans and buses also contribute significantly due to their higher NO_x emission factors and larger fraction of directly emitted NO_2.

This demonstrates that areas with a high concentration of diesel traffic are expected to generate larger roadside pollution zones within the simulation. Therefore, the model reflects the real-world observation that urban diesel traffic is one of the dominant contributors to roadside NO_2 pollution in European cities such as Dublin.

Interpretation of concentration model

The concentration equation converts vehicle-level emissions into ambient air concentration values ($\mu g/m^3$), allowing the simulation outputs to be compared against EU air-quality standards. This makes the model more meaningful than only reporting raw emissions, because it directly relates simulation results to public health thresholds.

The model also demonstrates how traffic congestion increases pollution accumulation. When traffic speed decreases, vehicles spend more time inside a road segment, increasing the total emitted mass within that zone. As a result, congested urban streets are expected to experience higher simulated NO_2 concentrations.

Validation of the emission model

To validate the realism of the simulation, the selected emission factors were compared against established European emission models and real-world measurement studies. The NO_x values are based on COPERT v5 emission factors, which are widely used within European transport and environmental studies. Additionally, the direct NO_2 fractions are supported by roadside monitoring and PEMS studies on Euro 6 diesel vehicles.

The simulated trends are also consistent with findings from Dublin roadside air-quality studies, which identified heavy diesel traffic corridors as major NO_2 hotspots. Therefore, although the model simplifies atmospheric dispersion effects, the generated emission patterns remain representative of realistic urban traffic pollution behaviour.

Types of vehicle data In order for emission levels to be modelled more accurately, we need information about the different vehicle types. A study conducted in Dublin (Mahesh et al., 2024) gives us insight into the amount of vehicles of a certain type. The study took measurements at six locations in Dublin. For our simulation, we assume the average percentage per vehicle type for the entire city. The shares are shown in table 2.

Table 2. Percentages per vehicle type (Mahesh et al., 2024)

Vehicle type	Templeogue	College Green	Beach Road	Chapelizod	Mayor Street	Richmond Street	Average
Car	83.5%	58.3%	81.3%	85.9%	64.5%	74%	74.6%
LGV < 3.5t	14.0%	5.18%	16.8%	13.7%	34.5%	17.1%	16.9%
Bus	1.56%	35.6%	0.97%	0.26%	0.91%	7.52%	7.8%
HGV > 12t	0.52%	0.46%	0.34%	0.10%	0.34%	0.59%	0.31%
HGV 3.5–12t	0.36%	0.30%	0.34%	0.05%	0.21%	0.59%	0.31%
Mini-bus	0.05%	0.17%	0.21%	0.05%	–	0.20%	0.13%

To simplify the model, the HGV categories are combined into one category. Due to its low share, the mini-bus vehicle type is not taken into account in the simulation model.

The different fuel types of the vehicles are also taken into account. The study finds the percentages of cars per fuel type shown in figure 5. All other vehicle categories in the simulation model are assumed to run on diesel.

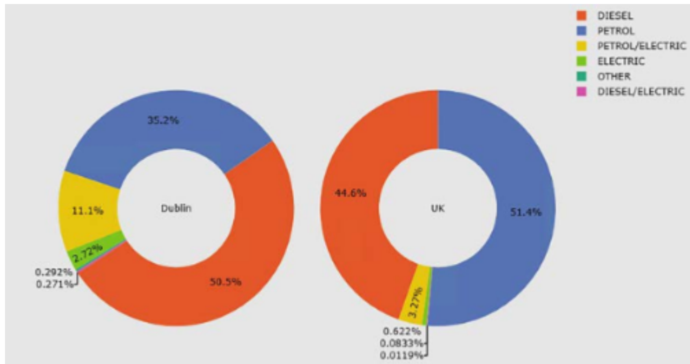

Fig. 5. Percentages of cars per fuel type.

In the simulation model, only petrol, diesel and electric cars are considered because they make up the majority of Dublin's car fleet.

Driver behavior modelling To model driver behavior in this scenario, the most important is the willingness to pay (i.e. the attitude each driver has when faced with the cost of the trip). This can be modeled by taking into account driver's income and time sensitivity, as well as the cost of the alternative.

Income distribution data can be obtained from the official Central Statistics Office run by the Irish government. In the dataset, the population is divided into 10 deciles, according to their household income, and the average income is computed for each decile. In our case, every time a driver is generated will be assigned a decile (driver type) which will define part of their behavior (except for professional drivers, which are a separate category).

The behavior will also be affected by the cost of the alternative, which is public transportation. According to the data obtained from the Transport for Ireland website, the approximate yearly cost of an adult using public transportation to commute regularly, is about 1250€/year. Divided into the (approximated) 260 working days and 2 trips per day, the cost is around 2.40€/trip. A reasoning can be made that a regular person (decile 5) will always accept the cost of the trip if it's below that value. The threshold value will be the same for drivers up to decile 5, but it will increase with income until decile 10.

When cost surpasses that threshold, the probability of accepting the trip will decrease exponentially. To construct that function, an upper bound is defined (cost at which trip will be rejected 95% of times). This UB is left as a parameter, to be adjusted for optimal simulation results.

Finally, time sensitivity plays an important role, as people will be willing to pay higher prices if their trip is urgent. TS is a percentage generated randomly. To modify driver's behavior, TS changes the already defined UB, following an exponential function (e.g. if TS is 50%, UB of the driver increases 10%). These values are also parameters that can be modified to nudge driver's behavior in a way that total emissions are reduced.

Decile	Avg. Yearly Household Income	WTP (euros/trip)	UB (95% rejection)
1	17100	2,4	2,51
2	29300	2,4	2,58
3	38600	2,4	2,64
4	46900	2,4	2,69
5	56200	2,4	2,75
6	66400	2,8	3,25
7	78600	3,4	3,85
8	91100	3,9	4,46
9	108400	4,6	5,30
10	181700	7,8	8,89

Table 3. Income distribution, WTP and UB per driver type, based on preliminary values.

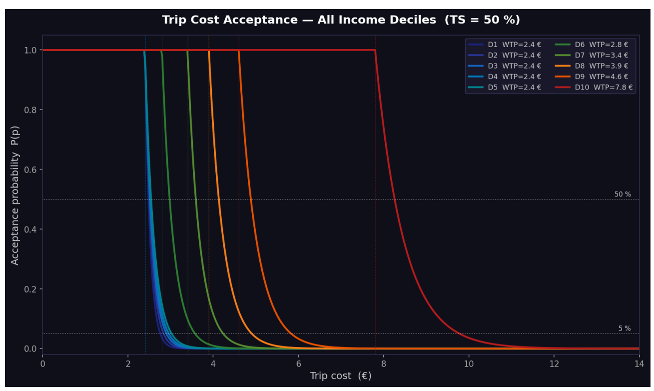

Fig. 6. Trip cost acceptance of all deciles at TS = 50%.

One last category are professional drivers, whose behavior is modelled the same way but it doesn't depend on income. To show that they don't really have a choice, their willingness to pay is set to be very high.

3.3 Simulation Framework and Stochastic Validation

While discrete-event simulation packages like Salabim are ideal for queuing models, the computational requirements of this project necessitated a custom simulation architecture. To accurately model urban emission exposure, the system required global, synchronous matrix calculations for spatial emission decay (utilizing `numpy` arrays) and high-performance, dynamic graph routing across a geographical map (utilizing `rustworkx` and `osmnx`). To accommodate this architecture, a custom discrete-time step engine was developed in Python. The stochastic sampling and validation procedures function identically to those in a standard Salabim environment, ensuring the statistical validity of the agent generation.

To validate the variables used within the simulation (such as Time Sensitivity and Willingness-to-Pay), random sampling generation was verified against known mathematical distributions. Below is a trace of the first 10 generated driver agents, demonstrating the agent generation output during the initialization phase:

```
--- Custom Agent Generation Output (First 10 Samples) ---
Agent 1: Dep Time=3s, Type=Origin/Dest, VOT=0.0030, WTP=€18.00
Agent 2: Dep Time=0s, Type=Unassigned, VOT=0.0025, WTP=€25.00
Agent 3: Dep Time=6s, Type=Unassigned, VOT=0.0107, WTP=€80.00
Agent 4: Dep Time=5s, Type=Unassigned, VOT=0.0063, WTP=€25.00
Agent 5: Dep Time=5s, Type=Unassigned, VOT=0.0116, WTP=€100.00
Agent 6: Dep Time=6s, Type=Unassigned, VOT=0.0067, WTP=€30.00
Agent 7: Dep Time=1s, Type=Unassigned, VOT=0.0053, WTP=€25.00
Agent 8: Dep Time=4s, Type=Unassigned, VOT=0.0030, WTP=€35.00
Agent 9: Dep Time=9s, Type=Unassigned, VOT=0.0118, WTP=€60.00
Agent 10: Dep Time=3s, Type=Unassigned, VOT=0.0009, WTP=€15.00
```

Furthermore, the uniform distribution generator utilized for the Time Sensitivity (TS) parameter, $U(0, 1)$, was validated using a sample array of $n = 1000$ to ensure a lack of bias. A Kolmogorov-Smirnov (KS) test was performed to verify conformity to the expected uniform distribution. The results confirm that the generated values are valid for the simulation environment:

```
--- Stochastic Validation: Time Sensitivity (n=1000) ---
First 10 values: [0.1024, 0.2998, 0.1049, 0.2271, 0.9295, 0.5079, 0.1433, 0.3528, 0.1365, 0
Mean: 0.5099 (Expected: 0.5)
Standard Deviation: 0.2843 (Expected: ~0.2887)
KS Statistic: 0.0258
P-value: 0.5086
```

Conclusion: p-value > 0.05. We fail to reject the null hypothesis. The generator successfully models a uniform distribution, satisfying validation requirements.

3.4 Generated data analysis

Multiple data outputs can be obtained as a result of the simulation, for them to be analyzed and the simulation validated and the proposed solution evaluated.

First of all, there is emissions data. This includes:

– Total emission levels
– Emissions per zone
– Emissions per time of the day
– Emissions corresponding to each vehicle type

Another output is the total number of trips, with all its associated information (vehicle type, driver type, time of the day, total duration, total emissions).

The last relevant dataset that can be generated is the evolution of price in the different areas along the day (depending on emissions), as well as the total collected money when implementing this pricing strategy.

Zoning definition (Areas & Prices) The final piece for the simulation to be completed is the division of the city map into zones and defining the prices each vehicle has to pay when going through that area. This analysis is to be done after the first iteration of the simulation (which will model the system without yet implementing the proposed solution).

Once the simulation is completed, a traffic and emissions map will be generated as an output. Based on this map, zones can be defined to avoid the observed issues (either high emission levels or traffic jams). Once this first division is put in place, multiple iterations of the simulation will be conducted and, after each one, zoning areas will be tuned to obtain an adequate result.

Pricing strategy will follow a similar logic. A first set of prices will be implemented, which will later be tuned after each iteration. Pricing relations (ratio between prices of different zones or at different times) will also change to make the system more efficient, always based on some basic predetermined rules (e.g. pricing will increase in high-emission areas to achieve rerouting and to price out some drivers).

Simulation assessment Once this information is available, the following questions can be answered:

- Is the emission level lower when the proposed solution is applied?
- Which areas are more affected?
- How does the price affect the different types of drivers?
- Which vehicle type contributes more to the generated emissions?
- Which vehicle type has a higher change in emissions due to dynamic pricing?
- Is there less traffic overall?
- Which variable has a higher effect in reducing emissions?

In the answers to these questions lies the effectiveness of the model.

4 Simulation Implementation and Experimental Evaluation

4.1 Baseline simulation implementation

The first step of the simulation process is to implement the logic and data described in previous sections (pdl code and data collection). We will be using Pyhton, but not the SimPy and Salabim libraries, due to the very high computation cost that they have (by using those libraries, the simulation took around 8 to 10 hours to run, and we took different actions to reduce it as much as possible).

Instead, the following libraries are used. **OSMnx** is used to download and process real-world road network data from OpenStreetMap. **NetworkX** and **Rustworkx** provide graph data structures and algorithms that support routing, pathfinding, and network analysis, with Rustworkx offering improved performance for large networks. **Pygame** is responsible for the graphical interface, animation, and real-time visualization of the simulation. **NumPy** enables efficient numerical computations and array operations, which are essential for handling large amounts of simulation data. **SciPy**, including its **cKDTree** and stats modules, provides scientific computing tools such as fast nearest-neighbor searches and statistical analysis. **Matplotlib**, along with **Pyplot** and **LineCollection**, is used to generate plots, figures, and visual representations of network structures and simulation results. **TQDM** adds progress bars to monitor lengthy computations. Finally, Python's built-in libraries (os, sys, pickle, math, time, random, and collections) support operating system interactions, program control, data storage, mathematical calculations, timing functions, random number generation, and specialized data structures used throughout the simulation.

Initially, the simulation does not include the proposed solution, since the goal of this first iteration is to validate the model. To do so, some basic questions need to be answered: are drivers following their assigned routes?, do they behave in a logical way, by optimizing their routes based on time?, is traffic modeled correctly?, are there more vehicles in high density areas (like the highway)?, etc.

The results of this first iteration are shown in figure 7.

This image shows that the model is, in fact, behaving correctly, as more emissions (and therefore more drivers) are concentrated in the highway areas. It also shows the high amount of drivers going to the center of the city. Besides, the emission results obtained in this iteration of the simulation are very similar to the values observed in reality, as it was explained in section 3.

This emissions map is a good starting point to define the zones we will use in future iterations. Also, the baseline emission level clearly shows a morning and afternoon rush hour, in which more trips are generated. In the baseline simulation, 103 trips are rejected. This is caused by the time sensitivity of the

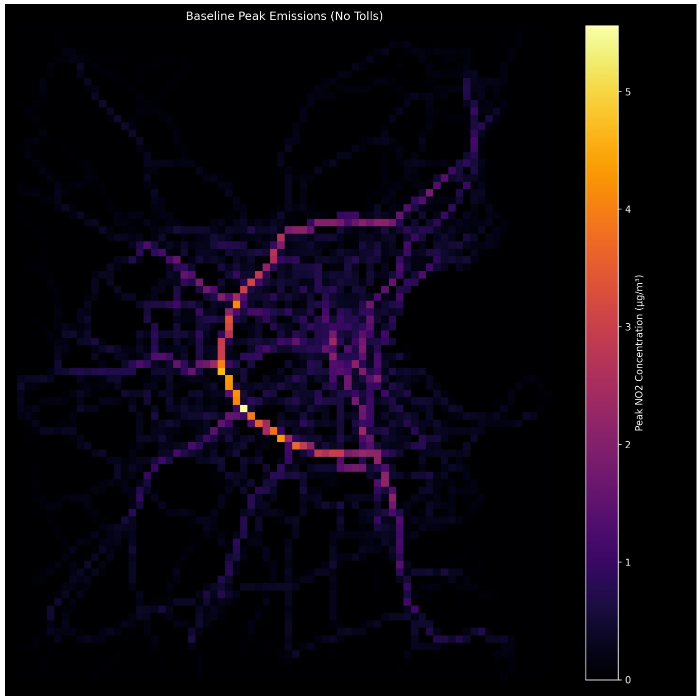

Fig. 7. Emission levels on the baseline simulation (no solution implemented).

drivers, with their originally intended trip already taking too much time. These results will be shown in more detail in the following section.

4.2 Dynamic pricing implementation

In the current implementation, three concentric zones are defined in the city center of Dublin, as shown in figure 8. The goal of the simulation is to reduce the emissions in the center zone, zone 1. The dynamic pricing is only implemented in this zone. Because the simulation has three zones the decisions made by drivers can be made more clear, as the emission levels of the three different zones can be monitored. For every node on the map the model checks whether this node is inside a specific zone. This is done using a Euclidean distance check against the zone's radius. The classification of the nodes allows for trips to be put into three categories: internal trips, where both the start and end point are within the zone; Origin/Destination trips, where either the start or the end point is within the zone; Unassigned trips, where both the start and end point are not within the zone. These categories are especially helpful for later identifying how drivers react to the zone prices. In addition, each road segment is divided into evenly-spaced sample points, and each point is checked against the toll zone. If any point falls within the zone, the entire segment is marked as a toll edge. This flag is stored with the segment at the start of the simulation, so the routing logic never has to repeat the geometric check while the simulation is running. The same sampling process also records how much of each segment passes through each area of the map. This is used when calculating emissions. Rather than dumping all of a vehicle's pollution at a single point, the simulation spreads it across the relevant areas in proportion to how much road the vehicle actually drove through each one.

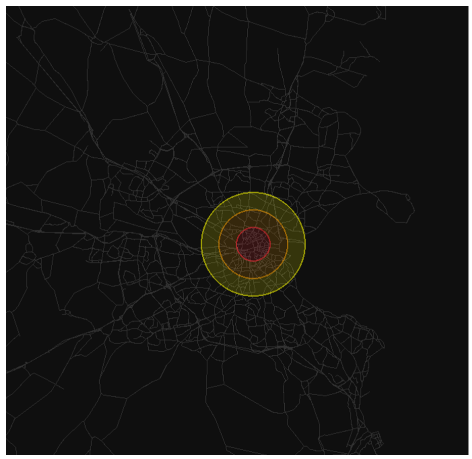

Fig. 8. Three concentric zones implemented in the simulation

The dynamic prices are generated using two components. The first component is a base charge of €3 charged to any vehicle entering the zone. The second component is a variable component which adds €2 to the charge for every $\mu g/m^3$ the NO_2 emission level exceeds the safe threshold of 12 $\mu g/m^3$. The current implementation makes sure that there are no trips that are priced out. This way the effect of the detoured trips on the total emissions in the zones is clearly noticeable.

When its route would take a vehicle through the zone, the model calculates two options. The first option is the cheapest route that would accept the toll charge. The second option is the cheapest route that would avoid the zone entirely. This choice between this option is based on the drivers willingness to pay but also their time sensitivity. Each driver is is assigned a number from 1 to 10, representing ten income levels. Every income level has their own acceptable toll charge range. The highest charge range is assigned to professional vehicles, like busses. All other vehicles get their income level assigned randomly. The time sensitivity is assigned in a random uniform distribution among the drivers. If neither option falls into the range of what the driver is willing to pay, the driver is removed from the road and the trip is treated as abandoned or shifted to other modes of transport.

The results of the baseline simulation including the dynamic pricing implementation is shown in figure 9 and 10. Information about the simulated trips is shown in table 4.

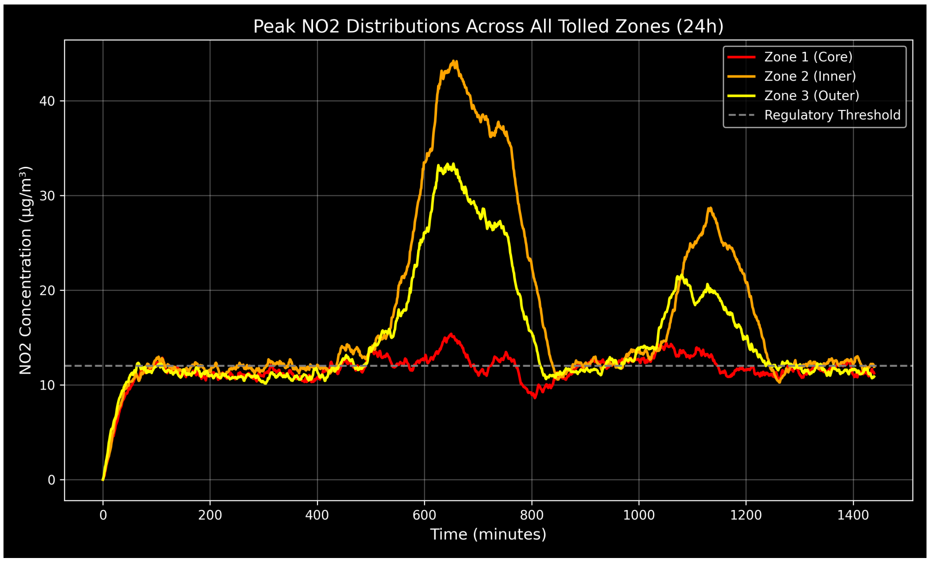

Fig. 9. Emission levels in the three zones for the dynamic pricing implementation

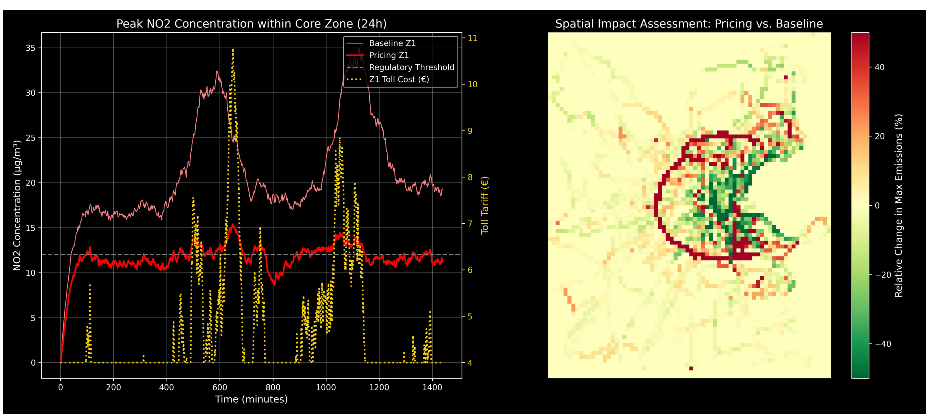

Fig. 10. Baseline simulation including dynamic pricing implementation

Table 4. Simulation Results

Baseline Model (No Toll)	
Total Trips Attempted	1,004,648
Trips Completed	1,004,545
– Internal Routing	48,194
– Origin/Destination	356,499
– Through-Traffic	113,573
– External Boundary	486,279
Trips Rejected	103
Congestion Pricing Model	
Total Trips Attempted	1,004,648
Trips Completed	1,004,545
– Internal Routing	48,194
– Origin/Destination	356,499
– Through-Traffic	21,845
– External Boundary	578,007
Trips Priced Out	103
Trips Detoured	619,520
Tolls Collected	55,525
Financial Summary	
Total Toll Revenue	€385,158.97
Average Toll Paid	€6.94

Figure 9 shows the emission levels in the different zones. From figure 10 the effect of the dynamic pricing on the emission levels in zone 1 can be seen. The simulation manages to keep the level below or closely to the threshold value. It can be seen that the emission levels rise in the other zones, indicating that drivers divert their route in order to avoid the center zone. Compared to the baseline simulation, more than half of the completed trips is rerouted. This can also bee seen in the map on the right in figure 10. The center zone shows a relative reduction in emissions, whereas the surrounding highway has a relative increase. The amount of through-traffic is also heavily reduced, whereas an increase in the external boundary trips can be observed. This clearly indicates that drivers reroute to avoid entering the zone entirely.

4.3 Experimental Evaluation

In the previous section, the proposed solution has been implemented in the simulation and the results have been analyzed. After this first successful run, the next step is to conduct a series of experiments with two main goals: validate the simulation model in different situations and observe how the results are affected by changes in relevant parameters.

To do so there is one key limitation: computation time (using the complete model, it was taking up to 12-14 h for one iteration to run). For this reason, and to be able to obtain more and more relevant results, we decided to scale down the model. Both the number of drivers generated and the emissions threshold have been reduced so that the model is more manageable. By doing this, the behavior of the simulation should be the same, but the computation time is drastically reduced (number of drivers increases computation time in a non-linear way).

Therefore, the first step of this section is to run the simulation with the baseline parameters but in the reduced version. The results can be seen in figure 11 and figure 12.

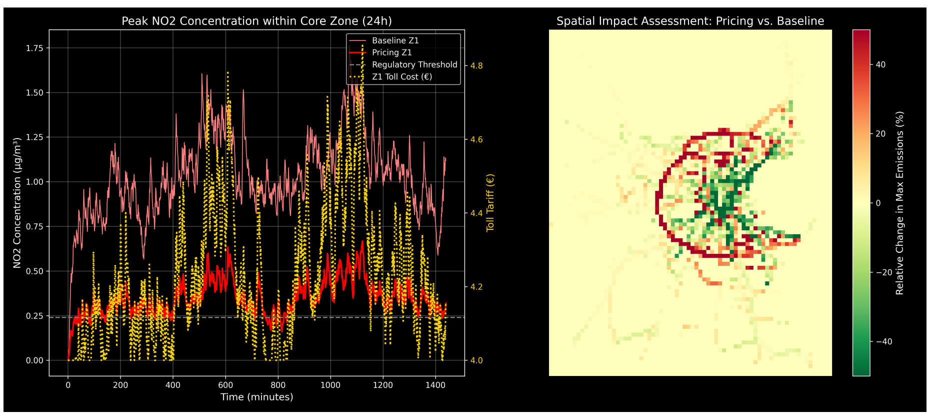

Fig. 11. Results of the reduced baseline simulation.

As it can be observed in both figures, the output datasets follow the same pattern as in the full scale simulation. The most relevant result is that, in this situation (same as in the previous section), emissions of the center area are significantly reduced by implementing the proposed solution. There are certain moments of the day, though, when emissions surpass the established threshold. This is likely due to the nonlinearities introduced in the system by the number of drivers, and is something to be aware of when comparing different situations. Another important change is the introduction of larger noise effect in the graph, due to the reduction in scale.

There are clear limitation to this approach, though, being the main one that traffic won't be the same, so emissions might be lower (proportionally) because of that. This limitation should be tackled in future iterations of the project.

The experiments are chosen to conduct a sensitivity analysis on relevant parameters of the simulation.

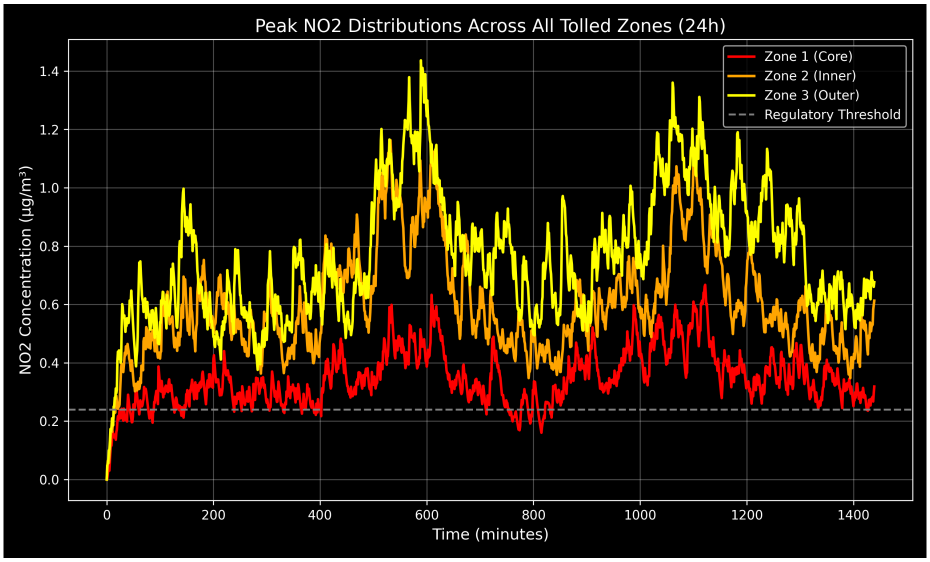

Fig. 12. Results of the reduced baseline simulation (cont.)

Experiment 1: Pricing Sensitivity Analysis The effectiveness of the system relies heavily on the pricing level applied, so various pricing configurations are tested to analyze how this parameter is related to the behavior of drivers.

To do so, two different configurations are run, scaling the base prices per zone at 50% of their original values and at 150%, respectively. The hypothesis is that, as prices increase, more drivers will cancel their trips and more reroutes will be conducted, so emission levels will go down. The inverse is expected to be true when prices are lowered.

It can be observed in figures 13 and 14 that there is a clear difference in emissions when the prices change, being a low price less effective in keeping emissions under the threshold value. That being said, the increase in prices doesn't seem to affect emissions in the core area by much, since the obtained results are very similar to those of the baseline reduced model. Also, the number of rejected trips (priced out) stays the same in the 3 different scenarios.

This suggests that the current model of driver behavior is prone to have a high acceptance ratio. Also, it means that prices affect the system nonlinearly (since reducing the price by 50% has a higher effect in one direction than increasing it by 50% has in the other direction).

Another important result is that the 150% experiment has a higher noise in toll tariffs than the other experiments. A possible explanation for this is that,

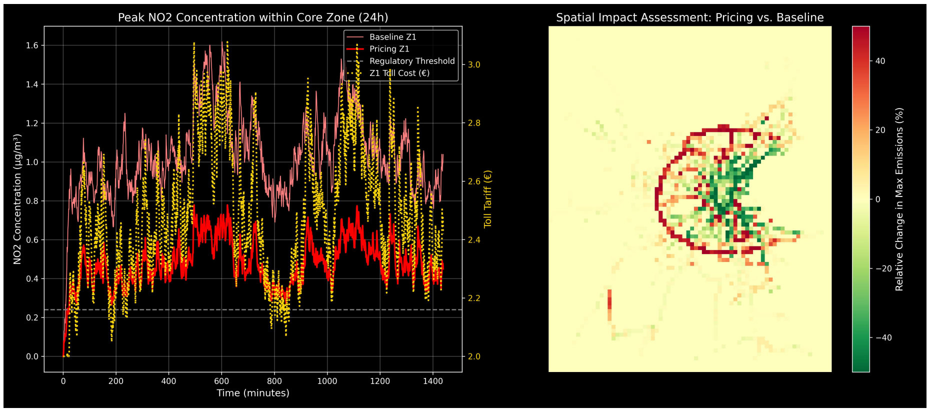

Fig. 13. Zone 1 results of the experiment at **50%** of the original prices.

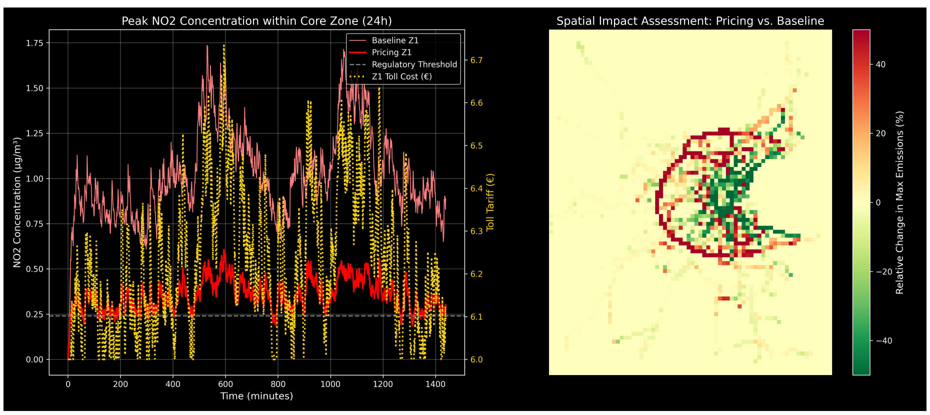

Fig. 14. Zone 1 results of the experiment at **150%** of the original prices.

when toll tariffs kick in, their effect on emissions is faster, since the starting toll price is higher. This higher response causes the dynamic prices to go down again, thus restoring the higher levels of the emissions.

A possible solution for this would be to include in the pricing function a term that provides some "inertia" to the system. With more stable prices would also come more stable and predictable emission levels, that could be better adjusted by our solution.

Despite these limitations, these experiments also help validate the simulation, since emissions in the center area are reduced in both scenarios. It can also be observed in figures 15 and 16 that the higher the price, the greater the difference between the internal zone and the others. However, this difference might materialize in some cases as a trade-off between zones, rather than an overall emission reduction, since trips are rerouted. This situation is definitely worth checking out in future iterations of the simulation.

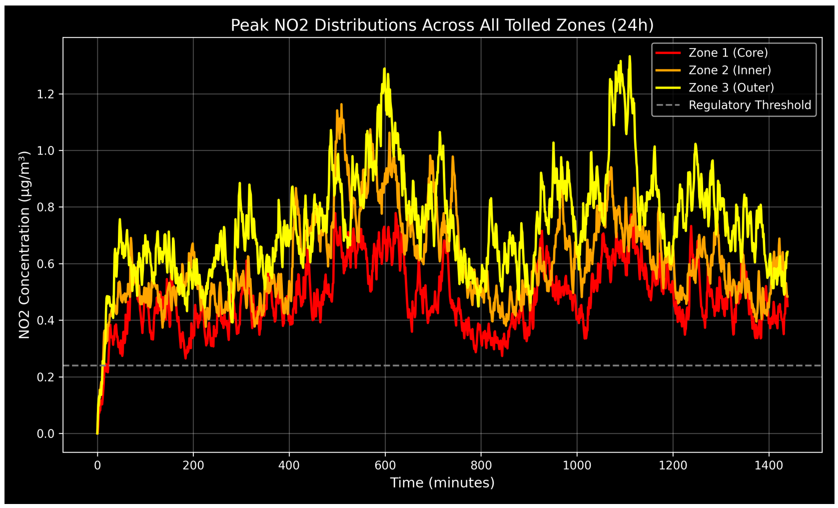

Fig. 15. All zones emissions of the experiment at **50%** of the original prices.

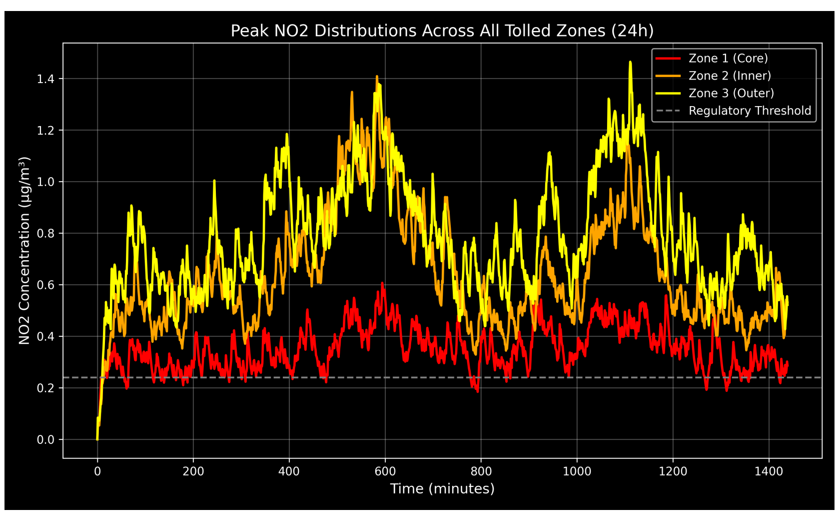

Fig. 16. All zones emissions of the experiment at **150%** of the original prices.

Experiment 2: Emission Threshold Sensitivity In this experiment, the impact of changing the value of the activation emission threshold of the pricing controller is considered. The threshold determines when the dynamic part of the pricing controller starts increasing the toll in response to high emissions. A lower threshold corresponds to a stricter policy, because the system reacts earlier. A higher threshold corresponds to a more relaxed policy, because higher concentrations are tolerated before additional pricing is applied.

There is purpose to this experiment as it helps see the tradeoff between emission reduction and traffic disruption.

Our *Hypothesis* is as follows: The strict threshold is expected to make the pricing controller more aggressive, since the dynamic toll is activated at a lower NO_2 concentration. Therefore, it is expected to increase toll activity and encourage more rerouting away from the core zone. This should reduce the concentration inside the controlled zones. The relaxed threshold is expected to make the controller less sensitive, since the dynamic toll is activated later. This should reduce intervention and traffic disruption, but may result in overall higher NO_2 concentrations in the zones as the controller reacts later.

Table 5. Emission threshold scenarios.

Scenario	Threshold ($\mu g/m^3$)	Change from baseline
Strict	0.16	−33.3%
Baseline	0.24	reference
Relaxed	0.32	+33.3%

The relative change between the baseline pricing model and the experiment scenario's pricing models is calculated as:

$$\text{Relative change} = \frac{\text{Experiment value} - \text{Baseline value}}{\text{Baseline value}} \times 100$$

A negative value means that the pricing model reduced the value compared to the baseline model. A positive value means that the value increased.

Strict scenario. The strict scenario uses the lower activation threshold. This means that the dynamic toll reacts earlier when NO_2 concentration increases.

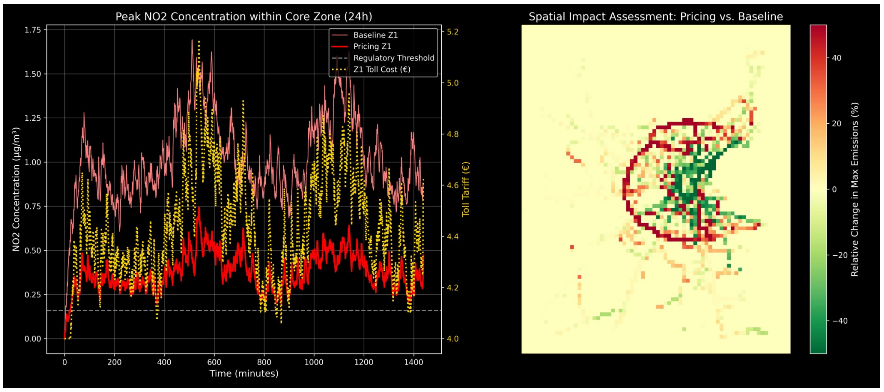

Fig. 17. Core-zone pricing performance and spatial emission impact for the strict threshold.

Table 6. Comparison between baseline pricing and strict threshold pricing.

Metric	Baseline pricing	Strict pricing	Relative delta (%)
Total trips attempted	41,846	41,846	0.00
Trips completed	41,842	41,844	+0.00
Trips priced out	4	2	-50.00
Trips detoured	13,888	13,676	-1.53
Tolls collected	4,388	4,543	+3.53
Internal routing	2,091	2,385	+14.06
Origin/Destination	14,976	15,489	+3.43
Through-traffic	136	78	-42.65
External boundary	24,639	23,892	-3.03
Total toll revenue (€)	23,026.53	26,113.10	+13.40
Average toll paid (€)	5.25	5.75	+9.52

Strict scenario analysis. Compared with the baseline pricing scenario, the strict threshold does not show a clear visual reduction in the NO_2 concentration curves between Figure 12 and Figure 18. Although the controller is expected to react earlier, which would reroute drivers more and spread out concentrations, the plotted core-zone concentration remains in a similar range, and one visible peak is even slightly higher than in the baseline emission threshold pricing case. This suggests that lowering the threshold does not automatically reduce every concentration peak, because rerouted vehicles can shift emissions to different parts of the network and inherently create local spikes at a different location.

However, the traffic and pricing results show that the controller does become more aggressive. The number of tolls collected increases from 4,388 to 4,543, which is an increase of 3.5%. Furthermore, the average toll paid increases from €5.25 to €5.75, showing a 9.5% increase. Through-traffic also decreases strongly

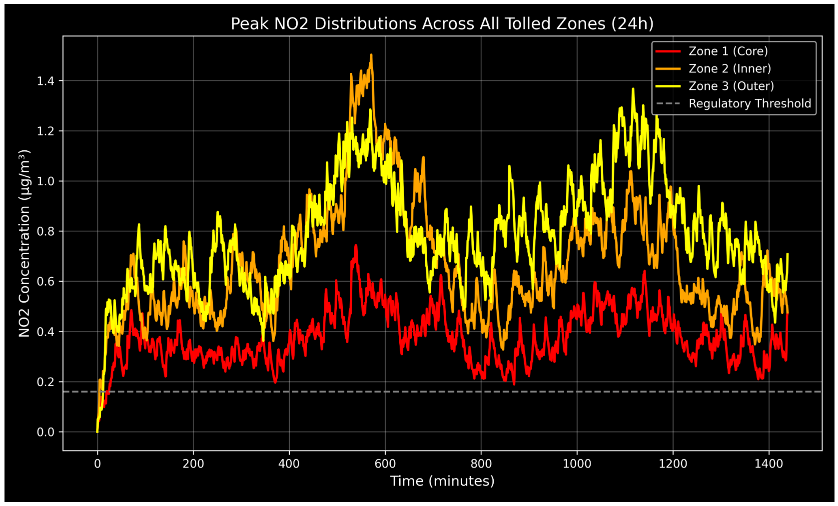

Fig. 18. NO_2 concentration across all tolled zones for the strict threshold.

from 136 to 78 trips, which is a large change of 42.5%, showing that the stricter threshold discourages vehicles from passing through the controlled area. Overall, the strict threshold increases pricing pressure and reduces through-traffic, increasing total toll revenue by 13.4%, but the visual emission results do not show a clearly proportional reduction in peak NO_2 concentration. This can be the case due to the scaled-down version of the experiment, where small changes in route allocation can have a relatively large effect on local peak concentrations or even small simulation deviations (that were mostly observed when using different laptops). Therefore, the results should be interpreted as indicative rather than fully conclusive, and further repeated runs with different decreases of the threshold would be needed for stronger conclusions.

Relaxed scenario. The relaxed scenario uses the higher activation threshold. This means that the dynamic toll reacts later, allowing higher NO_2 concentrations before additional pricing is applied.

Relaxed scenario analysis. In this scenario, compared with the baseline pricing, the relaxed threshold actually does show a visual increase in the NO_2 concentration curves between Figure 12 and Figure 20, following the hypothesis. In these figures, the relaxed scenario shows visibly higher concentration levels than the baseline case, especially in Zone 3, where the peak reaches close to $1.6\,\mu g/m^3$ compared to the $1.4\,\mu g/m^3$ in the baseline case. This suggests that increasing the threshold does result in larger concentration peaks as the controller reacts later.

Furthermore, the traffic and pricing results also show that the relaxed threshold reduces some of the pricing pressure compared with the baseline pricing sce-

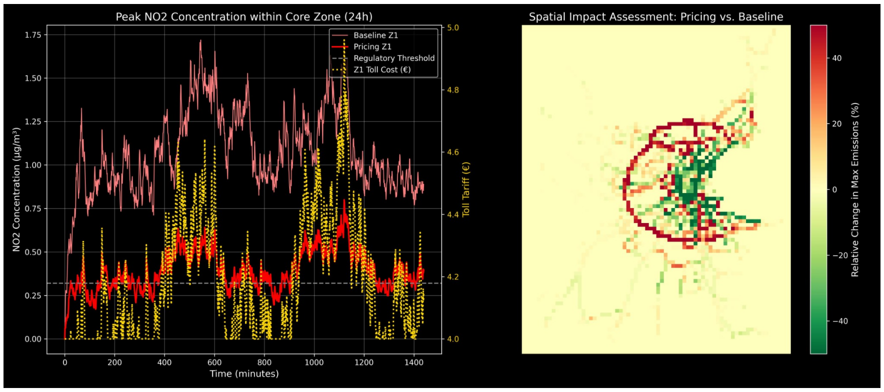

Fig. 19. Core-zone pricing performance and spatial emission impact for the relaxed threshold.

Table 7. Comparison between baseline pricing and relaxed threshold pricing.

Metric	Baseline pricing	Relaxed pricing	Relative delta (%)
Total trips attempted	41,846	41,846	0.00
Trips completed	41,842	41,845	+0.01
Trips priced out	4	1	-75.00
Trips detoured	13,888	13,495	-2.83
Tolls collected	4,388	4,728	+7.75
Internal routing	2,091	2,049	-2.01
Origin/Destination	14,976	15,464	+3.26
Through-traffic	136	199	+46.32
External boundary	24,639	24,133	-2.05
Total toll revenue (€)	23,026.53	24,004.00	+4.24
Average toll paid (€)	5.25	5.08	-3.24

nario. The number of detoured trips decreased from 13,888 to 13,495, which is a -2.83 % change. The average toll paid decreased from €5.25 to €5.08, which is a -3.24% change. However, the number of tolls collected increases from 4,388 to 4,728, while the total toll revenue increases by 4.2%. This can be explained by the through-traffic also increasing from 136 to 199 trips, a 46% increase, showing that the relaxed threshold allows more vehicles to pass through the controlled area, where the toll prices are lower than in the baseline, because the threshold for the toll prices to increase is higher. Overall, the relaxed threshold reduces the average toll burden and causes fewer detours, but it also results in higher NO_2 concentrations and more through-traffic. Similar to the strict scenario, the results should be interpreted as indicative rather than fully conclusive due to the nature of the simulation.

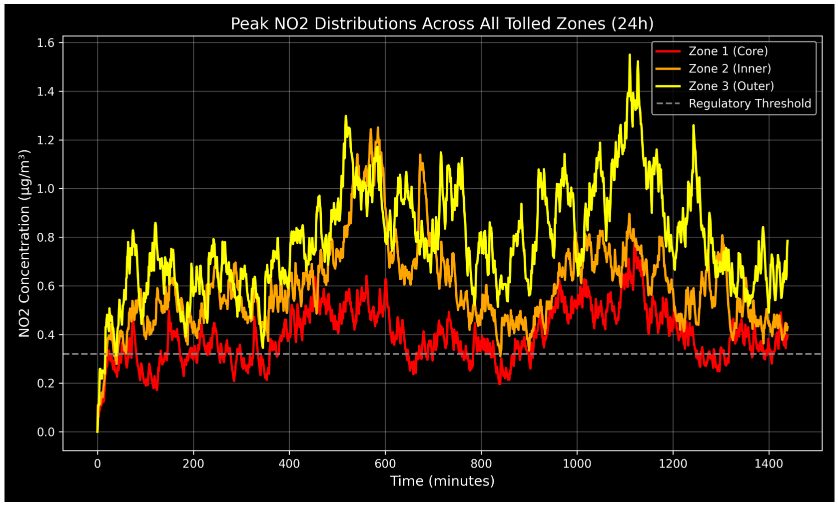

Fig. 20. NO_2 concentration across all tolled zones for the relaxed threshold.

Experiment 3: Vehicle Fleet Composition Since diesel vehicles contribute disproportionately to NO_2 emissions, here in this experiment we can assess the impact of vehicle composition.

Table 8. Vehicle fleet composition scenarios.

Scenario	Diesel Vehicle Share
Green Fleet	-15% of both diesel vehicles (ie. old and new)
Current Fleet	Existing distribution

Considering the scenario of green fleet, as mentioned in table 8, we conducted an experiment consisting of a 15% decrease of both old and new diesel vehicle and replaced them with EV vehicles in our vehicle distribution. Here in this experiment we aim to determine how pricing strategies work with different vehicle composition. And also to see if our system provides more benefits in traffic conditions.

Results & Discussion: The experiment provides critical insights into how vehicular interventions can interact with the economy and emission levels. Here, our objective is to assess the experiment results compared to baseline scenario. On a broader view, we can see that in Table 9, we can see that in both scenarios out of $41,846$ total trips attempts, resulting in precisely $41,842$ successful trip completions and 4 failures. This tells us that replacing diesel vehicles with electric vehicles does not alter the physical mobility balance of the whole system.

Altering the distribution of vehicle brings about some subtle changes in the routing behaviour in the simulation. Notably, the results shows an increased number of vehicles that travels through through traffic ie. from the table 9, there is an increase of 33.09% from 136 trips in the base case scenario to 181 trips. This routing deviation is associated with the emission trends in Fig.21. The reason for this change is mainly because the overall air quality was much better in the entire zone due to the presence of increased number of EVs; there were fewer peak threshold violations when comparing to the graphs of baseline scenario. Consequently, the dynamic pricing algorithm lowered the price barrier to some extend during some specific periods. And this allowed for the travel of more vehicles through traffic in city centre and not forcing it out through an external route. Due to this, the traffic around the center reduced by a small margin ie. the trips detoured reduced by 1.62%.

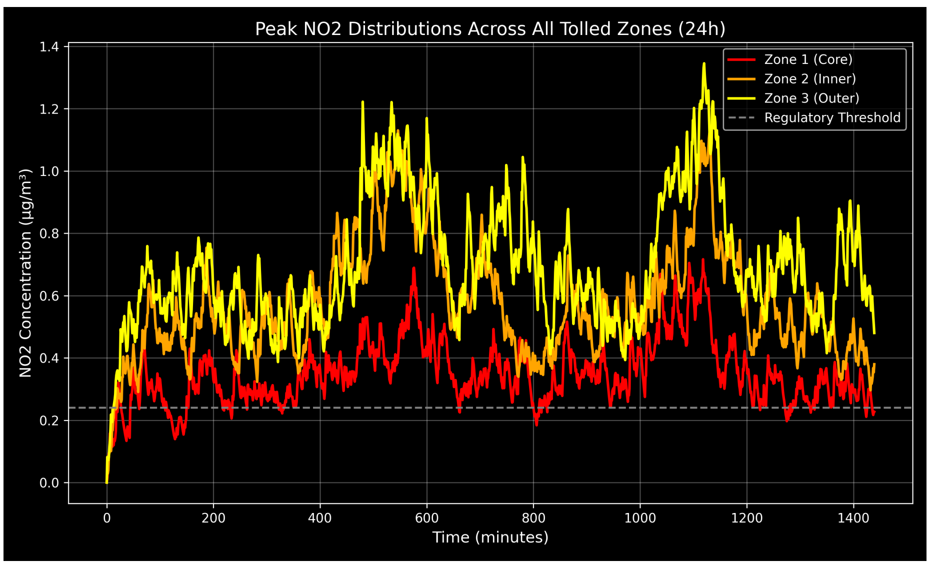

Fig. 21. Emission levels in the three zones for the Green fleet

While comparing Fig 21 and Fig 12, as expected, it is observed that overall peaks in the graph are slightly less than the peaks observed in the baseline scenario. For example, in Fig. 12, the Outer Zone 3 experiences peak morning and evening concentrations with an absolute maximum of about $1.4\,\mu g/m^3$. Once the EV fraction is implemented, the absolute maximum ceiling reduces to approximately $1.22\,\mu g/m^3$, which indicates a substantial reduction in the concentration level.

Also, it is understood that while comparing fig. 22 and fig. 11 that in the figure of spatial assessment of the baseline scenario, many diesel cars are forced to drive around the toll zone, which created lot more red spots than in the same

Table 9. Comparison: Baseline vs. Experiment 3(Green Fleet scenario)

Metric	Baseline Scenario	Experiment 3	Absolute Delta (Δ)	Relative Delta (%Δ)
System Overview				
Total Trips Attempted	41,846	41,846	0	0.00%
Trips Completed	41,842	41,842	0	0.00%
Trips Rejected	4	4	0	0.00%
No toll Model				
– Internal Routing	2,091	2,123	+32	+1.53%
– Origin/Destination	14,976	14,870	-106	-0.71%
– Through-Traffic	4,103	4,278	+175	+4.27%
– External Boundary	20,672	20,571	-101	-0.49%
Congestion Pricing Model				
Trips Priced Out	4	4	0	0.00%
Trips Detoured	13,888	13,663	-225	-1.62%
Tolls Collected (Total)	4,388	4,373	-15	-0.34%
– Internal Routing	2,091	2,123	+32	+1.53%
– Origin/Destination	14,976	14,870	-106	-0.71%
– Through-Traffic	136	181	+45	+33.09%
– External Boundary	24,639	24,668	+29	+0.12%
Financial Summary				
Total Toll Revenue	€23,026.53	€22,954.04	-€72.49	-0.31%
Average Toll Paid	€5.25	€5.25	€0.00	0.00%

figure of the green fleet scenario where red spots lighten to a soft orange and yellow.

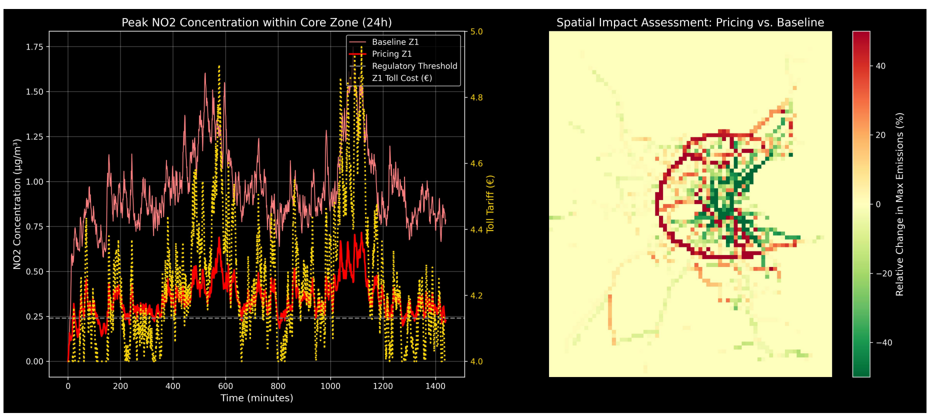

Fig. 22. Baseline simulation for Green fleet scenario

In conclusion, this comparison clearly reveals that although the simulation for congestion pricing has proved to be very effective in removing traffic from the inner city center, it tends to affect the peripheral boundary zones negatively. Introducing a minor 15% replacement of diesel by electric vehicles makes an alternative in this case. This has the potential to reduce the peak exposure profiles in the entire city on average without affecting the municipality's revenues.

4.4 Conclusions and future work

Throughout this project, a simulation has been developed with the objective of reducing vehicle emissions in urban areas. The proposed solution consists of dividing the map of the region into different pricing zones and dynamically adjusting the respective prices so that vehicles take different routes, avoiding the most inhabited urban areas (city center).

The simulation has been designed to model the city of Dublin and, to do so, extensive research has been conducted to collect and model the necessary data, such as driver behavior, number of vehicles, and types of vehicles. Additional data, such as current levels of emissions, have helped validate the simulation since, once it is run, the output values mimic the real ones.

The simulation has also shown that, once the proposed solution (dynamic pricing zones) is applied, both traffic and emissions in the city center are significantly reduced, thus fulfilling its main goal.

Multiple experiments have been conducted by varying different relevant parameters. These tests have provided an extra layer of validation to the model and have shown how the simulation reacts in different scenarios. They have also provided relevant information regarding the nonlinear behavior of the simulation with respect to certain parameters, such as the number of drivers and zone pricing, which is definitely important to consider for further analysis of the system.

The main limitation of the project has been the huge size of the simulation model, which has resulted in a very high computation time. To tackle this issue, an equivalent, reduced version of the model has been implemented for experimentation. Although the obtained results are deemed valuable, the same experiments should be run in the full-scale simulation for complete verification.

As future steps, there are some details that could be included in the model to make it more complex and closer to reality (such as refined areas with hospitals or schools, where emissions should be lower). Also, certain sections of the model (like traffic) could be refined so that more information could be extracted from them and so that they are represented in a more realistic way (modeling speed, traffic jams, etc.). Finally, the overall logic could be modified to achieve a general decrease in emissions or a reduction in traffic, situations that are considered to be outside the scope of this project.

References

1. Janet Currie and Reed Walker. Traffic congestion and infant health: Evidence from e-zpass. *American Economic Journal: Applied Economics*, 3(1):65–90, 2011.
2. Dublin City Council. Traffic volumes from scats traffic management system jul-dec 2025 dcc, 2025. Smart Dublin Open Data. Retrieved May 7, 2026.
3. Dublin City Council. Traffic signals and scats sites locations dcc, n.d. Smart Dublin Open Data. Retrieved May 7, 2026.
4. Richard Florida, Andrés Rodríguez-Pose, and Michael Storper. Cities in a post-covid world. *Urban Studies*, 58(11):2209–2221, 2021.
5. Jonathan Leape. The london congestion charge. *Journal of Economic Perspectives*, 20(4):157–176, 2006.
6. Lisa Schweitzer and Brian D. Taylor. Just pricing: The distributional effects of congestion pricing and sales taxes. *Transportation*, 35(6):797–812, 2008.

AI statement: AI tools (ChatGPT, Gemini, Claude) have been used in the elaboration of this report in the following ways:

- Help in finding sources for current solutions.
- Provide conceptual explanations of the used models to improve our understanding of them.
- Support in simulation code correction and debugging.
- Support in LaTex coding and format.
- Correction of specific spelling and grammatical errors.
- Used to generate code for traffic data analysis.
- Used for more coherent sentence and paragraph structures and grammatical errors.

Contribution Statement: Due to the nature of this project, it is very difficult to pinpoint the exact contribution to each part to one group member. A substantial amount of work was done together in person during brainstorm sessions and in writing the assignment, taking everyone's input into consideration. However, an approximate work distribution is the following:

- **Vishal:** elaboration of the BlackBox model, collection and analysis of the emissions data, development of an emissions model for the simulation.
- **Jort:** elaboration of the CATWOE model, collection and analysis of the types of vehicle data, development of a model to include types of vehicle in the simulation for the simulation.
- **Jack:** introduction, elaboration of the pdl code, research on geographical data, main developer of the simulation code.
- **Antoni:** elaboration of the pdl code, analysis of the population data, development of the driver behavior model for the simulation.
- **Ondrej:** elaboration of the swimlane diagram, collection and analysis of the traffic data, model of the traffic data for the simulation.
- **All:** ideation of the solution, brainstorming sessions on system description, components and simulation definition, execution of experimentation runs, analysis of simulation results, writing the report.