NON-PEER-REVIEWED PREPRINT

Submitted to EarthArXiv

High-resolution agent-based modelling of non-exhaust emissions reveals the limits of urban fleet electrification

Marc Sturrock[1], *marcsturrock@rcsi.ie*

https://orcid.org/0000-0002-7435-5256

[1]Department of Physiology and Medical Physics, Royal College of Surgeons in Ireland, Dublin, Ireland

February 2026

High-resolution agent-based modelling of non-exhaust emissions reveals the limits of urban fleet electrification

Marc Sturrock

Department of Physiology and Medical Physics
Royal College of Surgeons in Ireland, Dublin, Ireland

26th February 2026

Abstract

As exhaust emissions decline, non-exhaust emissions (NEEs) from brake and tyre wear are emerging as the dominant source of traffic-related particulate matter. This transition is complicated by the increasing mass of electric vehicles and the push toward high-capacity public transport. We introduce a high-resolution, moving-observer agent-based model calibrated against hyperlocal Google Air View street-level measurements at the Phibsborough junction in Dublin, Ireland, a documented pollution hotspot. Global sensitivity analysis reveals that the system is fundamentally interaction-dominated: all first-order Sobol indices fall below 0.013, while total-order indices for fleet composition (0.27–0.34), mean speed (0.27), and speed variability (0.23) are substantial, confirming that traffic dynamics are as influential as fleet mix. We evaluate four policy domains informed by this sensitivity structure through passenger-preserving scenario simulation. Fleet electrification benefits are entirely contingent on regenerative braking: at 50% penetration, UFP ranges from −19% with effective regenerative braking to +13% without, while $PM_{2.5}$ increases by 1–5% regardless of braking technology. Bus modal shift delivers the largest dual-pollutant benefits of any single intervention, peaking at 50% shift (−30% UFP, −34% $PM_{2.5}$), though a bus paradox emerges at 75% where additional heavy buses partially offset UFP gains. Replacing ICE cars with 1,000 kg microcars yields initially modest reductions (−4% UFP at 25% adoption, rising to −30% at 75%), with diminishing returns as the unchanged SUV fleet dominates brake wear. Speed management is the most robust technology-independent lever: a 30 km/h limit reduces UFP by 16.5% and $PM_{2.5}$ by 9% independently of fleet composition, while "slow + smooth" traffic can reduce UFP by up to 45%. A combined policy applying all five levers delivers −27% UFP and −12% $PM_{2.5}$, demonstrating compounding benefits from coordinated action. These findings establish that speed management provides a guaranteed floor of improvement, that bus modal shift is the most effective single intervention, and that electrification must be paired with effective regenerative braking and mass-conscious vehicle choices to avoid worsening tyre-derived $PM_{2.5}$.

1 Introduction

The environmental and health impacts of road transport are a well-established area of scientific inquiry. While exhaust emissions have historically been the primary focus of regulation and research, there is a growing recognition of the significant contribution of non-exhaust emissions (NEEs) to urban air pollution [Committee on the Medical Effects of Air Pollutants (COMEAP), 2020]. These emissions, which primarily consist of particulate matter (PM) from brake wear, tyre wear, and road surface abrasion, are not mitigated by the transition to electric vehicles (EVs) and, in some cases, may be exacerbated due to the increased weight of these vehicles.

This paper focuses on a specific area of concern: the busy Phibsborough junction in Dublin, Ireland. Data from the *Google Project Air View* initiative (May 2021 - August 2022) have identified this junction as a pollution hotspot within the city [Aclima & Google, 2022]. The data reveal that the mean concentrations of the smallest particulate matter fractions, PMch1_perL and PMch2_perL, are 68.72% and 80.47% higher, respectively, than the Dublin average. These values rank in the 83rd and 88th percentiles of all measurements across the city, indicating a significant local pollution problem. Given the dense commercial frontage in the immediate vicinity of the junction, there is a pressing need to understand the exposure of residents, workers, and patrons to these elevated pollution levels.

Agent-based modelling (ABM) offers a powerful paradigm for investigating such complex urban systems. As a bottom-up approach, ABM allows for the simulation of system-level emergent phenomena, such as traffic congestion and pollution hotspots, from the heterogeneous interactions and behaviours of individual agents. This contrasts with traditional top-down models, which often struggle to capture the granular dynamics that define real-world traffic flow. By modelling individual driver decision-making, vehicle characteristics, and their resulting environmental outputs, we can create a more nuanced and realistic link between micro-scale actions and macro-scale environmental outcomes. Our work builds upon previous research in this area, such as the TRAPSim model developed for Seoul [Shin, 2022, Shin and Bithell, 2023], which demonstrated the utility of ABMs for estimating personal exposure to NEEs.

Recent reviews highlight that the majority of transport ABMs still rely on coarse, speed-based emission factors and sparse roadside monitors, limiting their ability to resolve non-exhaust hotspots [Delcea and Chiriţă, 2023, Tzouras et al., 2023, Lam et al., 2025]. Studies that couple MATSim or SUMO to HBEFA-style inventories, such as the Berlin decarbonisation analyses and Munich's queue-based simulations, can quantify network-level behaviour, yet they necessarily proxy event-driven brake and tyre wear with average traffic states [Kaddoura et al., 2022, Hülsmann et al., 2014, Krajzewicz et al., 2015]. Even when individual exposure is computed, as in Tampa's inequality assessment, spatial fidelity is constrained by the lack of street-level reference data [Gurram et al., 2019]. Moreover, only a handful of implementations explicitly represent bounded rationality or attention lapses, despite evidence that distraction and heuristic routing materially alter stop-go dynamics and therefore NEE generation [Nourzad et al., 2014, Manley and Cheng, 2018]. These gaps motivated us to fuse second-by-second car-following, mechanistic brake and tyre modules, and hyperlocal Google Air View measurements within a single framework.

The scarcity of Dublin-focused modelling compounds these challenges. Ireland's capital depends heavily on buses (the dominant public transport mode) and is simultane-

ously accelerating the electrification of its private fleet [Central Statistics Office, 2022a, Society of the Irish Motor Industry, 2023]. Yet no published agent-based study has represented Irish driver behaviour, validated against the city's emerging street-scale pollution datasets, or stress-tested bus-centric interventions. By targeting the Phibsborough junction, where dense retail frontage keeps pedestrians and bus riders in close proximity to braking traffic, we address a critical evidence gap for managing NEEs in the Irish context.

The impetus for this research stems from the significant public health risks posed by NEEs. Exposure to PM constituents has been robustly linked to adverse cardiopulmonary outcomes, including the triggering of cardiovascular hospital admissions [Rich et al., 2019, Ostro et al., 2011] and acute functional decline in patients with Chronic Obstructive Pulmonary Disease (COPD) [Zhou et al., 2021]. Particles from brake systems, rich in metals like copper, have been shown to induce inflammatory and genotoxic effects in lung alveolar cells [Figliuzzi et al., 2020, Gerlofs-Nijland et al., 2019, Karlsson et al., 2006, Barosova et al., 2018]. A growing body of evidence also connects PM exposure during pregnancy to adverse birth outcomes, such as low birth weight [Sun et al., 2016, Meng et al., 2023], and an increased risk of preterm birth and stillbirth [Smith et al., 2020]. Effects are not limited to physical health, with studies also suggesting links between prenatal air pollution exposure and alterations in neurological development [Lubczyńska et al., 2020, von Ehrenstein et al., 2014]. These health risks are particularly acute in the Irish context. Ireland has the highest COPD hospitalisation rate among all 37 OECD member countries [Department of Health, 2020], and one of the highest global prevalences of alpha-1 antitrypsin deficiency (AATD), a genetic condition carried by 1 in 25 of the population that renders the lungs especially vulnerable to particulate-induced inflammation through the unopposed action of neutrophil elastase [Carroll et al., 2011, Wood et al., 2009]. The dense retail and hospitality frontage at the Phibsborough junction means that patrons and employees, including individuals with these prevalent conditions, are likely to have prolonged exposure to these harmful particulates, underscoring the importance of mitigating these emissions at a local level.

In response to these environmental pressures, urban policy is increasingly promoting a shift towards active transport (walking and cycling) and public transport. However, this transition may prove counter-productive for personal exposure: individuals who walk, cycle, or wait at bus stops spend more time in close proximity to traffic hotspots than those inside private vehicles. This underscores the critical need for micro-scale models that can evaluate the localised exposure risks associated with different transport modes and urban designs.

To address this gap, this study introduces a high-resolution agent-based model of the Phibsborough junction, calibrated against the Air View dataset, and uses it to evaluate passenger-preserving scenarios spanning fleet electrification, bus modal shift, microcar substitution, speed management, and a combined policy. The total number of people transported through the junction is held constant across all scenarios to ensure a fair comparison of modal impacts. For each scenario, we analyse the resulting distributional changes in brake and tyre wear emissions.

Although non-exhaust particles also arise from road-surface abrasion and the resuspension of legacy dust, we restrict the calibrated model to brake and tyre wear for three reasons. First, the Google Air View channels used for calibration (PMch1, PMch2, and $PM_{2.5}$) are dominated by freshly generated brake and tyre signatures, providing empirical leverage that is absent for resuspension at this site. Second, independent European inventories show that brake and tyre processes contribute the majority of primary non-exhaust

mass in congested, signalised streets similar to Phibsborough, whereas road abrasion is typically a minor fraction [Grigoratos and Martini, 2014, Organisation for Economic Co-operation and Development, 2020]. Third, Dublin's frequent rainfall (mean 170 rain-days per year) continually washes deposited material from the road surface, suppressing the dust reservoir that drives resuspension in drier climates; at this site, freshly generated brake and tyre wear is therefore the dominant particulate signal. Finally, resuspension depends on poorly observed factors such as street sweeping, moisture, and historical deposition; including it without site-specific data would add unconstrained parameters and dilute the policy insight [Committee on the Medical Effects of Air Pollutants (COMEAP), 2020]. The resulting scope keeps the model tightly coupled to available evidence while still capturing the dominant controllable sources of particulate matter at the junction.

The remainder of this paper is organised as follows. Section 2 describes the agent-based model, its emission sub-models, and the calibration procedure. Section 3 presents the validation results, the global sensitivity analysis, and the passenger-preserving scenario simulations covering speed management, electrification, bus modal shift, microcar substitution, and a combined policy. Section 4 discusses the implications for urban transport policy and identifies limitations and extensions.

2 Methodology

The methodology is centred around a bespoke agent-based model. This section details the study area, the model's structure, the sub-models for emissions, and the behavioural rules governing the agents.

2.1 Study area and data

The model is geographically centred on the Phibsborough junction, Dublin (53.36050, -6.27268). The baseline particulate matter data for model calibration is derived from the *Google Project Air View Dublin City (May 2021 - August 2022)* dataset [Aclima & Google, 2022]. The model environment is a 200x200 grid, with building footprints explicitly defined to create a realistic urban geometry.

2.2 Baseline fleet and modal share

To ensure scenario comparisons are meaningful, the baseline fleet mix and the scenario perturbations are defined in terms of passengers moved rather than just vehicle counts. In practice, we keep the total number of transported passengers constant across all scenarios and adjust fleet composition accordingly (e.g. increasing bus occupancy when reallocating car passengers). This design is anchored to Dublin city data, specifically the National Transport Authority's Canal Cordon counts for inbound city-centre travel and Census 2022 commuting patterns, so that the baseline composition and the magnitudes of the scenario shifts remain consistent with observed modal mixes in Dublin [National Transport Authority, 2023a, Central Statistics Office, 2023a].

2.2.1 Temporal alignment with Air View data collection

A critical consideration for baseline calibration is the temporal alignment between the model setup and the Air View measurements. Analysis of the GPS timestamps in the Google Project Air View dataset reveals that the Phibsborough measurements were collected predominantly during afternoon and inter-peak hours rather than the morning commute peak. Specifically, 43.6% of observations (464 measurements) were recorded between 15:00–17:00 UTC (16:00–18:00 IST), 54.4% (592 measurements) during the midday period 10:00–14:00 UTC (11:00–15:00 IST), and only 2.0% (16 measurements) during the morning peak 07:00–10:00 IST. The collection spanned 40 distinct dates between May 2021 and July 2022, with the Google Street View vehicle repeatedly traversing the junction during afternoon drives.

This temporal pattern has important implications for the baseline modal share assumptions. Afternoon traffic typically exhibits lower bus ridership than the morning peak, more varied trip purposes beyond work commuting, and different vehicle occupancy patterns. Ideally, calibration would use afternoon or inter-peak passenger modal shares that match the 14:00–18:00 window when most Air View data were collected. However, official transport statistics in Dublin are predominantly reported for the morning peak inbound commute (07:00–10:00), reflecting the focus of capacity planning and policy evaluation on this critical period.

2.2.2 Modal share parameters and occupancy assumptions

The NTA 2022 Canal Cordon Report provides the definitive empirical basis for our baseline passenger modal share [National Transport Authority, 2023a]. During the November 2022 AM peak inbound period (07:00–10:00), manual counts recorded 39,219 private cars carrying 49,679 passengers and 1,069 buses carrying 36,047 passengers. When considering only these two motorised road modes, private cars accounted for 57.9% of passengers and buses for 42.1%. This measurement establishes the passenger throughput that our model must reproduce.

Three critical factors require careful consideration when applying AM peak data to our afternoon calibration window. First, the implied average car occupancy from the Canal Cordon count is 1.27 passengers per vehicle, consistent with the NTA M50 Cordon finding that 91% of cars operate with single occupancy across both AM and PM peaks [National Transport Authority, 2023b]. This static, low occupancy throughout the day reflects a structural inefficiency in private vehicle use that is not confined to commuting trips. Second, bus occupancy exhibits pronounced time-of-day variation. The M50 Cordon data show that 49% of buses operate above 50% capacity during the AM peak, declining to only 31% during the PM peak (16:00–19:00), indicating lower afternoon ridership despite similar road congestion [National Transport Authority, 2023b]. Third, the directional asymmetry of commuting means that afternoon inbound travel is less concentrated than the morning peak, further reducing bus passenger loads during the Air View collection window.

Given these dynamics, we adopt the AM peak passenger modal share (57.9% private, 42.1% bus) as our baseline target, acknowledging that this likely represents a minimum private vehicle share for the day. Afternoon traffic almost certainly exhibits higher private vehicle dominance due to lower bus occupancy and ridership, making our baseline conservative in the sense that it maximises bus presence. Any underestimation of private vehicle share would bias our scenarios towards understating the NEE reductions achievable through modal shift or fleet electrification, thereby strengthening rather than weakening our policy conclusions.

Vehicle occupancy assumptions are derived directly from the 2022 NTA data. Private cars and SUVs carry 1.27 passengers per vehicle on average, reflecting the empirically observed single-occupancy dominance. Buses in our baseline scenario operate at 25% of their 90-passenger rated capacity (22.5 passengers per bus), chosen to reflect afternoon inter-peak conditions. This is substantially lower than the AM peak average of 33.7 passengers per bus observed at the Canal Cordon but aligns with the M50 Cordon evidence of declining occupancy outside peak hours. In bus-reallocation scenarios, we increase bus occupancy to 95% (86 passengers) to represent the service frequency and capacity improvements that would necessarily accompany a major modal shift policy. Microcars, when introduced in scenarios, carry 1.2 passengers on average, reflecting their smaller capacity relative to conventional cars.

2.2.3 Private vehicle fleet composition: SUV dominance and electrification trends

Analysis of the Dublin private car fleet in 2022 reveals two critical trends that directly impact non-exhaust emissions modelling. First, Sport Utility Vehicles (SUVs) and crossovers have surpassed traditional saloons and hatchbacks as the dominant vehicle segment. Data from the CSO on new vehicle registrations show that SUVs accounted for approximately

50.4% of combined car and SUV sales in 2022, establishing them as the default choice for Irish car buyers [Central Statistics Office, 2023b]. This shift has profound implications for vehicle mass and, consequently, for tyre-wear emissions.

Second, and more concerning from an emissions perspective, electric vehicle adoption in Ireland exhibits a pronounced bias towards heavier SUV formats. Analysis of the top five best-selling Battery Electric Vehicles (BEVs) in 2022 shows that three were SUVs or crossovers, including the market-leading Volkswagen ID.4 [Society of the Irish Motor Industry, 2023]. When examining the broader EV sales data, approximately 72% of top-selling electric vehicles were SUV-format vehicles rather than cars or hatchbacks. This indicates that Irish consumers are either maintaining their vehicle class during electrification (e.g., ICE SUV to electric SUV) or, more problematically, using the transition to electric powertrains as an opportunity to upsize from traditional cars to heavier electric SUVs.

Vehicle mass data underscore the NEE implications of this trend. The weighted average mass of top-selling EVs in 2022 was 1,877 kg, representing a 23.1% increase (+352 kg) over the average ICE vehicle mass of 1,525 kg. While regenerative braking in EVs reduces brake-wear emissions by recovering up to 80% of braking energy, the mass penalty directly increases tyre-wear particulates, which scale approximately linearly with vehicle weight. For example, replacing an ICE car with a similarly sized electric car (e.g., Nissan Leaf at 1,700 kg) adds 200 kg, whereas replacing it with an electric SUV (e.g., VW ID.4 at 2,049 kg) adds 550 kg; this mass difference translates directly to higher $PM_{2.5}$ generation from tyre wear.

For the Dublin-specific baseline, we synthesise multiple data sources to estimate the 2022 fleet composition [Central Statistics Office, 2022b, Department of Transport, 2023]. Electric vehicle penetration in Dublin exceeded the national average, with the CSO's Sustainable Mobility and Transport survey finding that EV ownership in Dublin (4.9%) was nearly double the national rate (2.5%). Applying this urban correction factor to year-end 2022 national fleet data yields an estimated Dublin private car fleet of approximately 541,400 vehicles, with the following composition: ICE cars 43.1%, electric cars (BEV+PHEV) 2.5%, ICE SUVs 51.9%, and electric SUVs (BEV+PHEV) 2.5%. This distribution reflects both the overall SUV dominance and the nascent but growing electric vehicle presence during the Air View measurement period.

The model therefore distinguishes six vehicle classes by powertrain and body type: ICE cars (1,500 kg), electric cars (1,700 kg), ICE SUVs (2,200 kg), electric SUVs (2,600 kg), buses (12,000 kg base + 75 kg per passenger), and microcars (1,000 kg, scenario only). Vehicle masses are based on manufacturer specifications for representative models in each class, and the 80% regenerative braking efficiency for electric vehicles is consistent with published performance data for contemporary BEVs. Buses are assigned 85% of the sampled free-flow speed to reflect their frequent stopping for passenger service and the lower operating speeds typical of urban bus routes. This granular classification enables the scenario analyses to isolate the mass effect on NEEs independent of powertrain technology.

2.2.4 Scope limited to road-based motorised modes

The model explicitly represents the six private vehicle classes described above plus scheduled bus services. We exclude rail (DART, Luas) because the Phibsborough junction has no proximate light rail or heavy rail stop and these modes do not interact with the signalised carriageway. Cycling and walking are likewise omitted: while both are impor-

tant for sustainable transport, they generate negligible non-exhaust emissions and do not contribute to the brake wear and tyre wear particulates that are the focus of this study. Freight and commercial vehicles (HGVs, vans, pickup trucks, and other light commercial vehicles) are excluded from the baseline fleet. HGVs are rare at this junction during the afternoon measurement window, and lighter commercial vehicles are not directly distinguishable in the Air View data; including them would introduce unconstrained parameters without corresponding observational constraints. This scoping keeps the model tightly coupled to the observed pollutant signature for the private vehicle and bus fleet that dominates traffic at this site.

While our micro-simulation necessarily operates with discrete vehicle counts to reproduce observed traffic flows at the Phibsborough junction, the implied passenger modal shares are aligned with the Dublin datasets cited above, ensuring that scenario comparisons reflect realistic policy-relevant shifts in urban transport mix.

2.3 The agent-based model (ABM)

The agent-based traffic simulator was developed in the Julia programming language. The model was prototyped using the `Agents.jl` package, which provides high-level scheduling and grid abstractions well suited to rapid iteration and debugging. A structurally equivalent production implementation was then built using `Ark.jl`, an entity-component-system (ECS) framework that stores agent state in flat, cache-friendly arrays rather than per-agent structs, achieving a 4–7× speedup. All CMA-ES calibration and global sensitivity analysis runs used the `Ark.jl` implementation. Julia's just-in-time compilation kept these batch workloads tractable even when thousands of parameter combinations were explored, and the pure-Julia stack avoids any dependence on proprietary software.

The Phibsborough junction is represented as a 200×200 regular `GridSpace` with periodic boundaries. Each lattice cell corresponds to a $1\text{m} \times 1\text{m}$ patch of carriageway, providing sufficient resolution to encode individual lanes, stop lines, and pedestrian refuges without imposing prohibitive computational costs. Building footprints and impermeable verges are stored as a boolean mask on this grid, and a set of scaling functions maps the canonical lane centre-lines (digitised from the baseline junction layout) to whatever grid resolution is required in sensitivity experiments. Vehicles are spawned stochastically along the outer 40% of each approach while respecting that mask; candidate positions are repeatedly sampled until a vehicle footprint can be placed without overlapping existing traffic, guaranteeing collision-free initial conditions for every run.

Two agent classes are employed. `CarAgent`s represent individual vehicles and carry their current and free-flow speeds, instantaneous heading, vehicle class (`:smart_car`, `:car`, `:suv`, `:e_suv`, `:bus`), and associated mass. The vehicle class determines the number of lattice cells that the vehicle occupies (two cells for microcars, three for cars, four for SUVs, and seven for buses) so that physical length constrains overtaking and saturation effects. Additional state variables store the intended manoeuvre, a stochastic planning distance drawn from a Poisson distribution (mean given by the calibrated `planning_distance`), the progress of any left turn, a cached reference to the governing traffic signal, and a continuous fractional position for sub-cell movement tracking. `TrafficLightAgent`s hold their signal state (`:green`, `:amber`, `:red`), a light-type flag (horizontal versus vertical approaches), and a cycle counter. The controller reproduces the 35 s green, 50 s amber, 50 s red stage lengths recorded during on-site surveys, ensuring that the simulated vehicle queues and release patterns match the observed signal

progression.

2.3.1 Adaptive time-stepping

The model employs an adaptive time-stepping scheme to balance computational efficiency with physical accuracy across the full range of mean speeds explored during calibration (5–15 m s^{-1}). The timestep Δt is computed from the maximum expected vehicle speed rather than the mean, ensuring that even the fastest vehicles in the speed distribution do not exceed the target movement distance per step:

$$\Delta t = \frac{d_{\text{target}}}{v_{\text{mean}} + 2\sigma_v} \tag{1}$$

where $d_{\text{target}} = 1.5$ m is the maximum distance any vehicle may travel per timestep, v_{mean} is the calibrated mean free-flow speed, and σ_v is the speed variability. The factor of two standard deviations captures approximately 95% of the speed distribution under the log-normal sampling used for individual vehicle speeds. This formulation prevents fast-moving vehicles from "hopping" over obstacles or stop lines, a critical failure mode in discrete-grid traffic models that would otherwise compromise collision detection. For the calibrated baseline ($v_{\text{mean}} = 9.13$ m s^{-1}, $\sigma_v = 4.05$ m s^{-1}), Equation 1 yields $\Delta t \approx 0.087$ s. At the lower extreme of the prior range ($v_{\text{mean}} = 5$ m s^{-1}), the timestep expands to $\Delta t \approx 0.25$ s, maintaining numerical stability while reducing computational cost for slower traffic scenarios.

2.3.2 Fractional positioning and two-phase movement

To eliminate discretisation artefacts inherent in integer-grid movement, particularly the spurious "sticking" of slow-moving vehicles when $v\Delta t < 0.5$ would round to zero displacement, the model tracks vehicle positions at continuous sub-cell resolution. Each `CarAgent` maintains a fractional position $(x_f, y_f) \in \mathbb{R}^2$ that evolves continuously according to its instantaneous speed, independent of the discrete lattice used for collision detection and visualisation. The fractional position is updated each timestep as:

$$\mathbf{r}_{t+\Delta t} = \mathbf{r}_t + v_t \Delta t \, \hat{\mathbf{d}} \tag{2}$$

where $\mathbf{r}_t = (x_f, y_f)$ is the continuous position, v_t is the instantaneous speed, and $\hat{\mathbf{d}}$ is the unit direction vector. The discrete grid position used for spatial queries and rendering is obtained by rounding: $(x_{\text{grid}}, y_{\text{grid}}) = (\text{round}(x_f), \text{round}(y_f))$, subject to clamping within the domain boundaries $[1, 200] \times [1, 200]$ to prevent indexing errors.

This continuous representation is coupled with a two-phase movement protocol that decouples intention from execution, enabling collision-free movement in dense traffic. During Phase I (planning), each vehicle computes its intended movement based on local traffic conditions, desired speed, and traffic-signal state, storing the proposed new position in a shared conflict registry. Distance calculations for car-following and stop-line detection operate on the fractional positions of both the subject (ego) vehicle and any leading vehicle, providing sub-metre accuracy:

$$d_{\text{gap}} = \begin{cases} r_{\text{leader,rear}} - r_{\text{ego,front}} & \text{if following} \\ r_{\text{stopline}} - r_{\text{ego,front}} & \text{if approaching signal} \end{cases} \tag{3}$$

where $r_{\text{leader,rear}}$ is computed from the leader's fractional front position and vehicle length. A dead-zone buffer of 1.0 m is applied to the minimum safe distance (1 m for stop lines, 1 m for following), within which speed is clamped to zero to prevent numerical creep. Phase II (conflict resolution) iterates over all intended moves in priority order (determined by a randomised tie-breaking rule seeded per timestep) and validates each move against the updated occupancy grid. If the intended cells are occupied, the vehicle remains stationary and its fractional position is reset to align with its current grid cell; otherwise the move is executed and the fractional position propagates forward. This protocol guarantees collision-free trajectories without requiring explicit inter-agent negotiation or centralised traffic assignment, and the fractional tracking ensures that vehicles approaching stop lines decelerate smoothly to rest exactly at the target position rather than stopping prematurely due to rounding.

Lane choice and turning manoeuvres are resolved during the planning phase. Agents intending to turn left begin seeking the dedicated lane once the distance to the stop line drops below their individual `planning_distance`; a lane change is only executed if every cell in the target lane is unoccupied. Discretionary lane changes for through-traffic are more conservative and are triggered only when the current lane is blocked one cell ahead and the adjacent lane is clear. Turning itself is modelled as an instantaneous perpendicular jump: when the vehicle reaches the turn point (one cell into the intersection), its direction rotates by 90° and its position is translated by the vehicle length in the new direction, ensuring the rear clears the cross-traffic stream. This geometric transformation completes in a single timestep, avoiding the gridlock that would arise from vehicles occupying intersection cells over multiple cycles.

The model tracks two floating-point raster fields that accumulate brake-derived and tyre-derived particles. After each agent update these grids undergo exponential decay governed by the global background rate and pollutant-specific decay multipliers, approximating the short-range dispersion observed in street-canyon studies [Shin, 2022, Shin and Bithell, 2023]. Random seeds are fixed at the start of every simulation to make the calibration and scenario sweeps reproducible while retaining stochastic variation within each replicate.

2.3.3 Model validation and robustness

To verify the collision-avoidance and traffic-flow logic, the model was subjected to extensive stress testing across the full parameter ranges explored during calibration. Collision detection operates by comparing the rounded grid footprints of all vehicle pairs at each timestep; a collision is recorded whenever two vehicles occupy the same lattice cell. Over 500-timestep simulations spanning the calibrated parameter space (mean speed 5–15 m s^{-1}, fleet compositions ranging from 100% ICE to 75% electrified, bus occupancy 22.5–86 passengers), the median collision rate is 0.2% (1 collision per 500 steps), with zero collisions observed in repeated calibration runs under the fitted baseline. The collision events that do occur are transient single-timestep overlaps during simultaneous turns at the intersection, resolving within one cycle without propagating gridlock. This collision rate is negligible compared to the traffic throughput (50+ vehicles transiting the junction over the 180-step analysis window) and does not materially affect the emission statistics used for calibration and scenario comparison. Additionally, a diagnostic suite tracking stopped vehicles at green lights, queue lengths, and vehicle dwell times confirmed that traffic flow remains stable and free of artificial blockages across all tested scenarios.

All experiments were executed with Julia 1.11.7 (channel `1.11`) on a Linux workstation (Intel Core i9-13900, 32 threads). Key packages include `Agents.jl` v6.2.10 for the agent-based simulation, `GlobalSensitivity.jl` v2.8.0 for variance-based screening. Full environment manifests are provided with the repository to ensure exact reproducibility.

2.4 Brake wear emissions

Particulate matter from brake wear is generated during deceleration events. Because a braking manoeuvre converts kinetic energy into heat through friction, we model the mass released per time step as being proportional to the kinetic energy dissipated by the vehicle, following the energy-based formulations used in non-exhaust emission models [Shin and Bithell, 2023]. The governing equation is:

$$E_{\text{brake}} = k_b \cdot W \cdot \Delta v \cdot (1 - f_{\text{regen}}) \tag{4}$$

Here E_{brake} is the instantaneous emission mass (μg), k_b is a dimensionless calibration coefficient, W is the vehicle mass (kg), and $\Delta v = \max(0, v_{t-1} - v_t)$ is the positive change in speed over one time step ($\mathrm{m\,s^{-1}}$). Vehicle masses are assigned by class (1.0 tonnes for microcars, 1.5 tonnes for cars, 2.2 tonnes for SUVs, 2.6 tonnes for electric SUVs, and 12.0 tonnes base + 75 kg per passenger for buses) so that heavier vehicles dissipate more energy when braking. The regenerative braking term f_{regen} captures the efficacy of energy recovery systems, and is fixed at 0 for internal-combustion vehicles and 0.8 for the electric SUV class, reducing the friction-brake demand for electrified fleets.

Whenever a braking event is detected ($\Delta v > 0$), Equation 4 is evaluated for every cell in the vehicle footprint. The resulting mass is projected into a 90° downstream cone with radius five cells, and the kernel weights each target cell by $\exp(-d^2/(r^2/4))$ to mimic the rapid dilution away from the stop line. The dispersed mass is added to the brake-emission raster, which is then subject to the global decay factor at the end of the time step. Observer statistics and calibration losses operate on the decayed grid after applying the fitted scaling factor and exponent that convert raw emissions into UFP counts.

2.5 Tyre wear emissions

Tyre wear emissions arise both from the constant rolling contact between rubber and tarmac and from the additional shear that accompanies speed changes. To capture these processes we decompose the source term into a baseline abrasion component that scales with instantaneous speed and a transient slip term that scales with the magnitude of acceleration or braking:

$$E_{\text{tyre}} = k_t \cdot W \cdot (c_1 v + c_2 |\delta v|) \tag{5}$$

In Equation 5, E_{tyre} denotes the instantaneous tyre-emission mass (μg), k_t is a dimensionless calibration coefficient, and W is the vehicle mass (kg). v is the current vehicle speed and $|\delta v| = |v_t - v_{t-1}|$ captures the magnitude of the speed change since the previous step. The weighting constants c_1 and c_2 control the relative influence of steady-state rolling versus manoeuvre-induced shear; in practice they are absorbed into the calibrated factors `tyre_to_pm25_factor` and `tyre_exponent` (Table 2), allowing the optimiser to match the asymmetric $\mathrm{PM}_{2.5}$ distribution observed in the Air View dataset.

Tyre emissions are evaluated at every time step, regardless of whether the vehicle is accelerating or braking, and are distributed onto the grid using the same conal kernel as

for brake wear. This produces a broader plume that follows the vehicle trajectories along each arm of the junction. The subsequent exponential decay step avoids unbounded accumulation under stationary vehicles, yielding the spatial gradients that were later validated against the moving observer.

2.6 Spatial dispersion and temporal decay

The treatment of particulate dispersion and decay in this model is tailored to the spatial scale (200 m × 200 m junction), temporal resolution (1-second timesteps), and measurement characteristics (street-level sensors at ∼1.5 m height) of the Google Project Air View dataset. Rather than implementing full three-dimensional atmospheric dispersion models appropriate for city-wide or regional scales, we adopt a parsimonious near-field approach that captures the dominant physics governing pollutant concentrations at breathing height within a signalised urban junction.

2.6.1 Conal dispersion kernel

Particulate emissions from both brake and tyre wear are dispersed spatially using an anisotropic conal kernel aligned with each vehicle's direction of travel:

$$C(x, y) = E \cdot \exp\left(-\frac{r^2}{R^2/4}\right) \cdot \mathbb{1}_{|\phi-\theta|<45°} \tag{6}$$

where r is the Euclidean distance from the source, $R = 5$ m is the effective dispersion radius, ϕ is the angle to the target cell, and $\mathbb{1}$ restricts deposition to a 90° cone. The 5-meter radius was selected to match the spatial resolution of the junction geometry and the observation that NEE concentrations exhibit steep gradients over distances of a few vehicle lengths in congested traffic [Shin, 2022, Shin and Bithell, 2023].

2.6.2 Temporal decay and effective removal

The concentration field for each pollutant evolves according to a first-order exponential decay process:

$$\frac{\partial C}{\partial t} = -\lambda C \tag{7}$$

where $\lambda = \lambda_0 \cdot \lambda_p$ is the effective decay rate, composed of a global baseline rate λ_0 and a pollutant-specific multiplier λ_p. This effective decay rate lumps together multiple physical removal mechanisms that operate at the junction scale: horizontal advection, turbulent diffusion, gravitational settling, and dry deposition.

2.7 Model calibration strategy

To ensure the model's outputs are quantitatively consistent with real-world measurements, a rigorous calibration process is undertaken. The model is fitted to two distinct empirical datasets from the Google Air View study: PM2.5 mass concentrations (linked to tyre wear) and ultrafine particle (UFP) number concentrations (from PMch1 and PMch2, linked to brake wear). The calibration uses a bespoke implementation of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), a derivative-free, population-based optimiser well suited to noisy, non-convex landscapes. At each generation, the algorithm

samples a population of candidate parameter vectors from a multivariate Gaussian, evaluates each candidate by running the full ABM simulation, and updates the mean, step size, and covariance matrix based on the ranked fitness values. The mean vector is injected into the first candidate at generation one to ensure the search begins from the best known point.

The objective is to minimise the distributional distance between simulated and observed UFP and $PM_{2.5}$. Each candidate parameter vector is evaluated on three independent random seeds; the per-seed losses are combined via the geometric mean, so that a single poorly fitting seed cannot be masked by the others. For each seed and pollutant, the loss combines quantile-based distance measures (RMSE and R^2 on the 1–99% quantile–quantile correspondence) with kernel density overlap in both linear and log space, with additional weight on the upper tail to ensure the model captures the heavy right-skewed distributions characteristic of junction-level NEE. The total loss is the equally weighted mean of the UFP and $PM_{2.5}$ components.

2.8 Evaluation strategy

2.8.1 Agent behaviour and moving observer

The behaviour of each `CarAgent` is governed by car-following logic (speed adjusted to the vehicle ahead), randomised driver reaction times, traffic-light response, and rules for turning and discretionary lane changes. To quantify localised exposure, the model attaches a moving observer to a specific vehicle, recording the cumulative pollutant concentration experienced by its occupant as it traverses the junction. The observer time-series provides the one-dimensional distributional data used for calibration and scenario comparison.

2.8.2 Stochastic replication

Each scenario is evaluated using three independent random seeds, yielding approximately 3,000 post-warmup samples (1,000 per seed at 1 Hz over 1,000 s of simulated traffic after a 100 s burn-in). The three realisations are pooled to compute scenario means and distributional statistics, providing robustness against seed-specific fluctuations in vehicle arrival patterns and signal-cycle alignment. The model-validation density plots (Figure 2b, c) show all three individual seeds alongside the pooled distribution, illustrating the consistency across realisations. Spatial fields and all scenario comparisons use the full three-seed pool.

2.8.3 Spatial evaluation via radial profiles

The physical geometry of the Phibsborough junction is an offset, staggered intersection rather than a symmetric perpendicular cross. Direct two-dimensional grid-cell-to-grid-cell spatial correlations between simulated and observed Air View fields would therefore incur systematic misalignment penalties unrelated to the emission physics. To evaluate spatial decay and hotspot placement independent of this geometric distortion, we compute one-dimensional radial profiles. Both simulated and observed pollutant fields are spatially binned into 5 m concentric annuli from the junction centre, and Spearman rank and Pearson linear correlations are reported on the resulting radial means. This approach isolates the physical decay of brake and tyre wear plumes from the discrete footprint of

the road network, providing a geometry-independent test of whether the model places emission mass at the correct distances from the stop lines.

2.9 Global sensitivity analysis

To identify which behavioural and fleet-composition parameters most strongly control NEEs, we performed a variance-based global sensitivity analysis (GSA) using a polynomial-chaos surrogate [Sudret, 2008, Crestaux et al., 2009]. Global methods explore the full prior ranges, revealing non-linearities and interactions that local perturbations can miss.

2.9.1 Parameter space and sampling strategy

The uncertain parameter vector comprises eleven controls: independent vehicle counts for cars (0–15), e-cars (0–5), SUVs (0–15), e-SUVs (0–5), buses (0–3), and microcars (0–5), each sampled independently over its full feasible range (constrained only by carriageway capacity), plus the mean free-flow speed (30–50 km/h), speed variability (0.25–3.0 m/s), green signal duration (15–45 s), effective sight distance (5–17 cells), and planning distance (±20% of calibrated value). Vehicle counts are treated as independent rather than passenger-preserving to allow the GSA to explore the full combinatorial fleet space without imposing an occupancy constraint that would couple fleet parameters *a priori*. A parallel analysis using passenger-preserving fleet shares (in which vehicle counts are renormalised to maintain constant total passengers) yielded remarkably consistent ST rankings and S2 interaction structure, confirming that the sensitivity results are robust to the choice of fleet parameterisation.

2.9.2 Polynomial-chaos surrogate

To compute Sobol indices efficiently, we fit a degree-3 polynomial-chaos expansion (PCE) to approximate the input-output map. The surrogate represents the response $Y(\boldsymbol{\theta})$ as a weighted sum of orthogonal polynomials:

$$\hat{Y}(\boldsymbol{\theta}) = \sum_{k=0}^{M} c_k \Psi_k(\boldsymbol{\theta}) \tag{8}$$

where $\{\Psi_k\}_{k=0}^{M}$ is an orthonormal Legendre basis on independent uniforms, c_k are expansion coefficients estimated by least-squares, and $M+1$ is set by the degree-3 truncation.

2.9.3 Analytical computation of Sobol indices

A key advantage of PCE is that Sobol sensitivity indices can be computed analytically from the expansion coefficients, avoiding the sampling noise inherent in traditional Monte Carlo estimators. The total-order Sobol index S_{T_j} for parameter θ_j captures its direct effect plus all interaction effects involving θ_j:

$$S_{T_j} = \frac{\sum_{k:j\in\mathcal{I}_k} c_k^2}{\sum_{k=1}^{M} c_k^2} \tag{9}$$

Parameters with $S_{T_j} \approx S_j$ exert mainly independent effects, while $S_{T_j} \gg S_j$ indicates strong interactions. The second-order index S_{ij} captures the pure interaction variance between parameters θ_i and θ_j:

$$S_{ij} = \frac{\sum_{k:\mathcal{I}_k=\{i,j\}} c_k^2}{\sum_{k=1}^{M} c_k^2} \tag{10}$$

3 Results

3.1 Google Air View data reveal Phibsborough as a Dublin non-exhaust emission hotspot

Analysis of the Google Air View dataset confirms that the Phibsborough junction is a significant hotspot for traffic-related air pollution within Dublin. As shown in Table 1, the mean concentrations of the particle fractions used for model calibration are substantially elevated compared to the city-wide average.

Specifically, the UFP-range channels `PMch1_perL` and `PMch2_perL` are 68.7% and 80.5% higher, respectively. The significance of this is underscored by their percentile ranks; the average `PMch1_perL` concentration at Phibsborough falls in the 83rd percentile of all measurements across Dublin, while `PMch2_perL` is in the 89th percentile. This indicates that the junction is among the most polluted locations in the city for these particle sizes, which are of high toxicological concern. Similarly, the overall `PM25_ugm3` mass is also elevated, ranking in the 68th percentile. These data provide a clear empirical justification for focusing detailed modelling efforts on this specific location to understand the dynamics driving these high emissions.

Table 1: Comparison of mean particulate matter concentrations at the Phibsborough junction versus the Dublin city-wide average for the metrics used in model calibration. The percentile rank indicates the position of the Phibsborough mean within the distribution of all Dublin measurements.

Metric	Phibsborough Mean	Dublin Mean	Difference (%)	Percentile Rank
`PMch1_perL`	39543.8	23436.9	+68.7%	83.0%
`PMch2_perL`	5399.7	2992.0	+80.5%	88.7%
`PM25_ugm3`	8.75	7.98	+9.7%	68.4%

Figure 1 presents the study area (a) and the spatial distribution of measured pollutant concentrations within the 200 m × 200 m box. The $PM_{2.5}$ (b) and UFP (c) heatmaps show elevated concentrations along the junction's approach arms, with the highest values occurring near the centre of the junction where vehicles decelerate and queue, broadly consistent with the expected spatial signature of brake and tyre wear emissions.

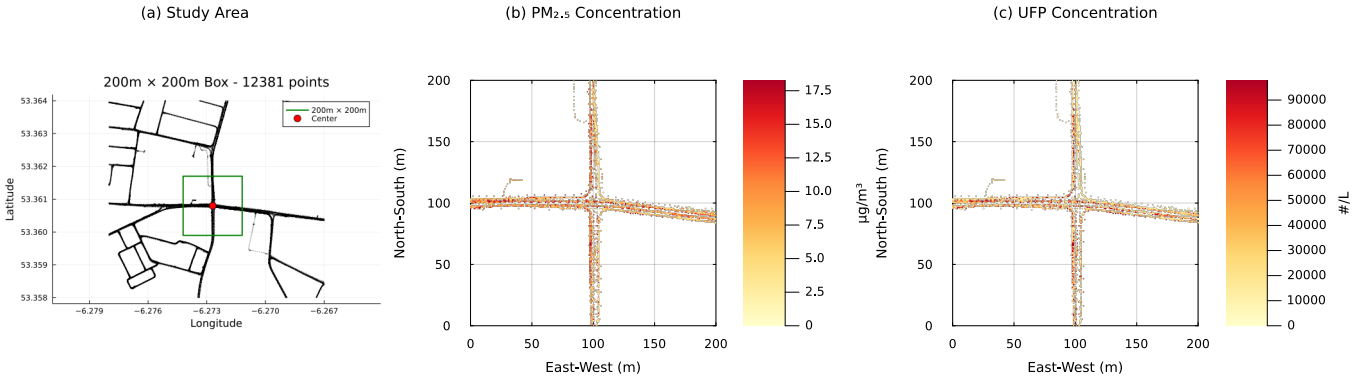

Figure 1: Study area and experimental data. (a) Dublin map showing the 200 m × 200 m study area centred on the Phibsborough junction. (b) Spatially binned mean $PM_{2.5}$ concentration from Google Air View measurements. (c) Spatially binned mean UFP concentration. Both pollutant fields peak at the junction's stop lines, consistent with non-exhaust emission hotspots.

3.2 Calibrated model quantitatively captures non-exhaust emissions

Having established the area's hotspot status, the next critical step was to ensure our model could accurately replicate these specific real-world conditions. The calibration, performed over 2,000 evaluations using the CMA-ES optimiser described in Section 2.7, successfully minimised the quantile-based distributional loss between the simulated and observational data.

Figure 2 presents the calibration and validation of the model. Panel (a) shows a representative model snapshot with vehicle positions, the moving observer, and the brake-emission overlay. The calibration targets are purely distributional: the CMA-ES optimiser minimises the quantile-based loss between simulated and observed UFP and $PM_{2.5}$ distributions, with no spatial information entering the objective function. The model reproduces the strongly right-skewed UFP distribution (brake wear) and the heavy-tailed $PM_{2.5}$ distribution (tyre wear), with simulated and observed modes closely aligned. To quantify shape agreement across the full distribution, we report the squared correlation between the 1–99% quantiles (R^2_q) annotated in each panel. The density plots (b, c) show all three individual seed realisations alongside the pooled distribution: all three seeds closely track the experimental data, with per-seed $R^2_q \geq 0.95$ for both UFP and $PM_{2.5}$. Pooling across the three seeds yields $R^2_q = 0.95$ for both $PM_{2.5}$ and UFP, confirming that distributional agreement is robust to stochastic variation [Montgomery et al., 2012].

As an independent validation, we compared the model's time-averaged spatial emission fields (panels d–f) against the spatially binned Air View measurements. Because the CMA-ES objective function operates entirely on one-dimensional, non-spatial quantile distributions, the accurate reproduction of two-dimensional spatial hotspot patterns constitutes a genuine out-of-sample test of the underlying mechanistic vehicle and dispersion models. Radial profiles of both $PM_{2.5}$ and UFP, binned at 5 m intervals from the junction centre (Section 2.8.3), show strong rank correlation between simulated and observed fields (Spearman $\rho = 0.69$ for $PM_{2.5}$, Pearson $r = 0.51$ for $PM_{2.5}$; Spearman $\rho = 0.47$ for UFP, Pearson $r = 0.56$ for UFP). The tyre-wear field (d) forms broad plumes along the approach arms, while the brake-wear field (e) concentrates at stop lines and left-turn approaches; the combined normalised field (f) confirms the overall hotspot placement. Together, the distributional calibration and the emergent spatial agreement provide a high-confidence baseline for scenario analysis, and the calibrated parameters are listed in Table 2.

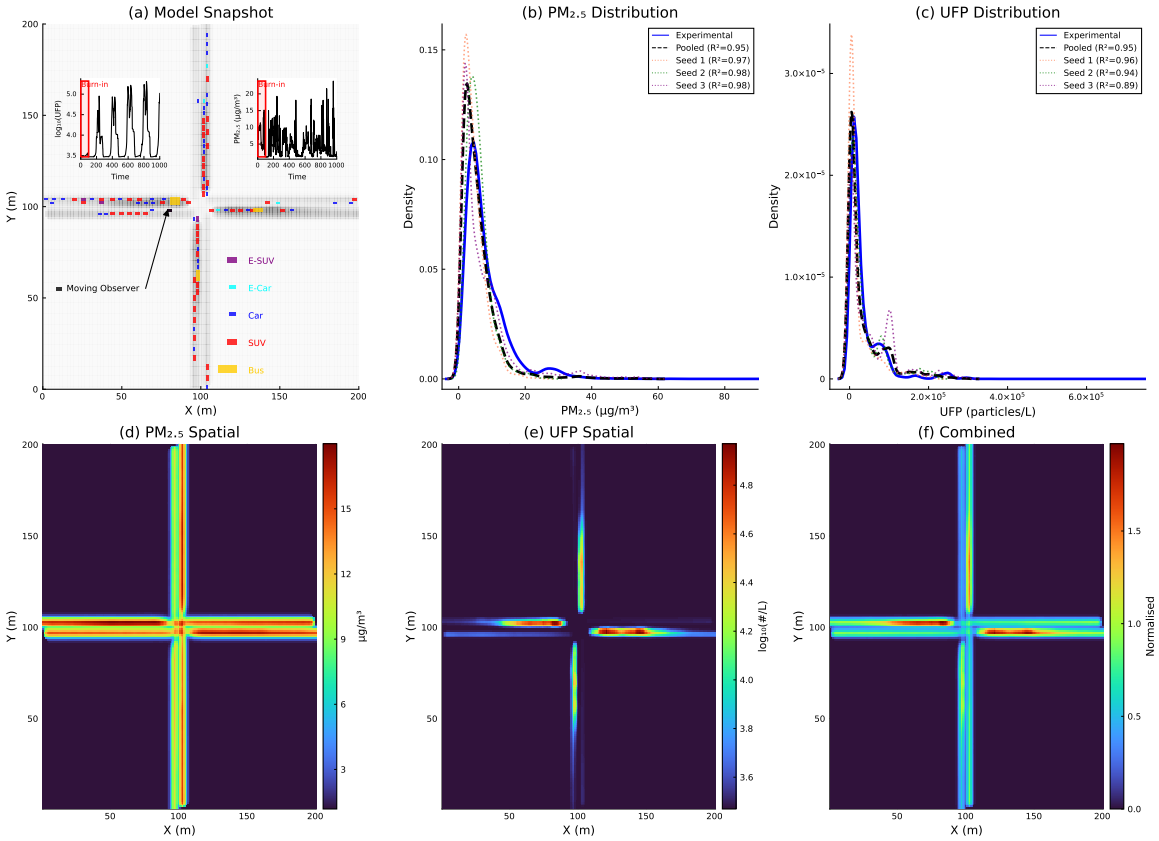

Figure 2: Calibrated model validation and spatial emission fields. Top row: (a) model snapshot showing vehicle positions, moving observer with emission time-series insets, and brake-emission overlay; (b) $PM_{2.5}$ distribution comparing three individual seed realisations (dotted), pooled model (dashed black), and Google Air View data (solid blue); (c) UFP distribution with the same format. All three seeds closely track the experimental distributions, with pooled $R_q^2 = 0.95$ for both $PM_{2.5}$ and UFP. Bottom row: time-averaged spatial fields averaged across three independent seeds for (d) $PM_{2.5}$ (tyre wear), (e) UFP (brake wear), and (f) combined normalised emissions. Brake wear concentrates at stop lines and left-turn approaches; tyre wear forms broader plumes along the approach and exit arms.

Table 2: Final calibrated model parameters after optimisation.

Parameter	**Description**	**Calibrated Value**
mean_speed	Mean free-flow target speed between braking events (m/s)	9.13
speed_variability	Random variation in individual vehicle speed (m/s)	4.05
brake_to_ufp_factor	Scaling factor for brake wear to UFP count	0.00187
tyre_to_pm25_factor	Scaling factor for tyre wear to PM2.5 mass	7.32
brake_decay_rate	Grid decay rate for brake particles	0.509
tyre_decay_rate	Grid decay rate for tyre particles	183
brake_exponent	Non-linearity of brake emissions	2.50
tyre_exponent	Non-linearity of tyre emissions	0.646
max_sight_distance	Effective headway distance a driver reacts to (cells)	16
planning_distance	Mean look-ahead distance for lane selection (cells)	5.59
green_duration	Calibrated green-phase duration (s)	25.0
brake_disp_radius	Dispersion radius for brake particles (cells)	5
tyre_disp_radius	Dispersion radius for tyre particles (cells)	5
ufp_offset	Additive UFP background (counts)	1387
pm25_offset	Additive $PM_{2.5}$ background ($\mu g/m^3$)	0.749

3.3 Global sensitivity analysis identifies speed and fleet interactions as dominant levers

We conducted a global sensitivity analysis on moving-observer NEE exposure to identify which parameters, and which parameter interactions most strongly control emissions at the junction. Fleet vehicle counts were varied independently (Section 3.3) so that the GSA explores the full combinatorial fleet space. A degree-3 polynomial-chaos surrogate was fitted to 2,000 Latin-hypercube samples $\times$ 5 independent repeats, and the Sobol indices were averaged across repeats to suppress surrogate-fitting noise. Figure 3 presents the results as first-order Sobol indices (S1), total-order indices (ST), and a second-order (S2) interaction heatmap.

The ST ranking (Figure 3b) reveals that car count ($ST = 0.34$) and e-car count ($ST = 0.34$) are the most influential individual parameters, followed by e-SUV count ($ST = 0.30$), bus count ($ST = 0.29$), SUV count ($ST = 0.27$), and mean speed ($ST = 0.27$). Microcar count ($ST = 0.26$) and speed variability ($ST = 0.23$) rank next, with behavioural parameters trailing: sight distance ($ST = 0.15$), green duration ($ST = 0.12$), and planning distance ($ST = 0.10$).

First-order effects (S1) are uniformly small, with all parameters below 0.013; the largest direct effect is bus count ($S1 = 0.013$). Total effects (ST) reach 0.34 (Figure 3a,b). The contrast between these two sets of indices demonstrates that the majority of output variance arises through parameter interactions rather than through isolated direct effects. This reflects the high dimensionality of the fleet-composition space: six vehicle-count parameters collectively span a large prior volume, and because NEE physics couples mass and velocity multiplicatively (Equations 4 and 5), the effect of any single parameter depends on the values of the others. The low S1 values should not be read as evidence that individual parameters lack influence; the scenario analyses that follow demonstrate that targeted changes to speed or fleet composition at the calibrated baseline produce large, policy-relevant emission shifts. Rather, the GSA reveals that the system is interaction-dominated, so interventions should be evaluated in the context of a specific fleet rather than in isolation.

The S2 heatmap (Figure 3c) confirms that pairwise interactions are diffuse: the largest S2 values are SUV$\times$e-SUV (0.035) and car$\times$e-car (0.030), reflecting the electrification mass trade-off, while speed$\times$variability (0.013) and bus$\times$microcar (0.010) round out the top pairs. All pairwise indices fall below 0.04, indicating that the degree-3 surrogate distributes variance across higher-order interaction terms rather than concentrating it in pairs, consistent with the strongly non-linear emission physics at this junction. The dominant policy-relevant insight comes from the ST ranking: traffic parameters (mean speed $ST = 0.27$, speed variability $ST = 0.23$) rival fleet-composition parameters in importance, and bus count ($ST = 0.29$) ranks above SUV count ($ST = 0.27$), confirming that public transport interventions carry substantial leverage.

These findings motivate the scenario analyses that follow: we evaluate electrification, bus and microcar modal shift, and speed management as the policy domains whose underlying parameters dominate the sensitivity structure, using 3-seed repeats (3,000 samples per scenario) to ensure distributional robustness.

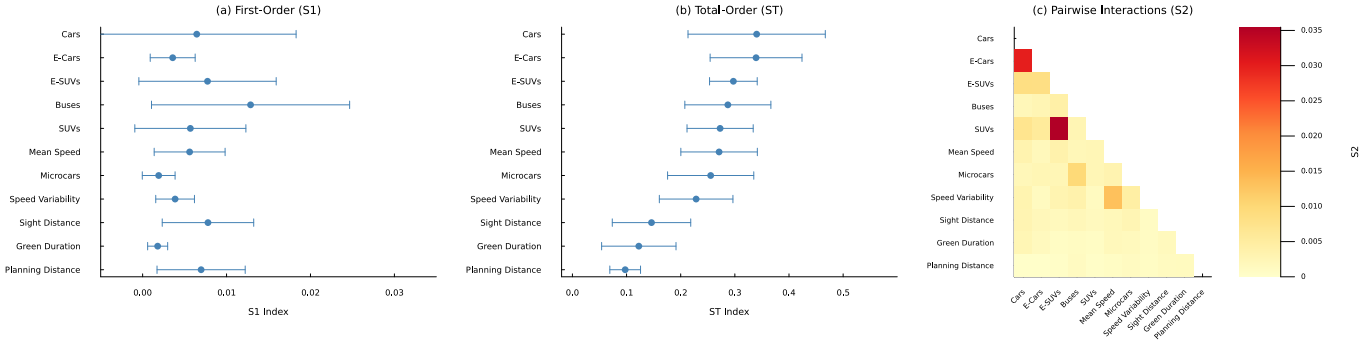

Figure 3: Global sensitivity analysis of moving-observer NEE exposure (degree-3 PCE, 2,000 samples $\times$ 5 repeats). (a) First-order Sobol indices (S1) are uniformly small (all below 0.013), confirming that direct effects are negligible relative to interactions. (b) Total-order indices (ST) reveal that car and e-car counts are the most influential parameters when interactions are included (ST $\approx$ 0.34), followed by e-SUV count (0.30), bus count (0.29), SUV count and mean speed (both 0.27). (c) The S2 heatmap identifies the dominant pairwise interactions: SUV$\times$e-SUV (0.035), car$\times$e-car (0.030), speed$\times$variability (0.013), and bus$\times$microcar (0.010). The contrast between small S1 and substantial ST values demonstrates that the system is fundamentally interaction-dominated.

3.4 Electrification benefits depend critically on regenerative braking assumptions

The high total-order indices for e-car (ST = 0.34) and e-SUV (ST = 0.30) counts, combined with their pairwise couplings to their ICE counterparts, reflect the strong sensitivity of junction-level emissions to electrification assumptions. To expose this interaction across a range of policy-relevant penetration levels, we simulated 25%, 50%, and 75% fleet electrification, each under two regenerative braking conditions: 80% efficiency (optimistic: warm battery, moderate state-of-charge, smooth driving) and 0% efficiency (pessimistic: cold battery, high SoC, or aggressive driving). Each electrification scenario replaces ICE cars with e-cars and ICE SUVs with e-SUVs at the same spawn positions, preserving the car/SUV ratio and initial spatial configuration of traffic; a 100 s burn-in period (Section 2.8) allows the traffic to reach a realistic flow state before data collection begins. Table 3 reports the fleet composition and mass accounting; Figure 4 presents the quantile-ratio distributions for UFP (a) and $PM_{2.5}$ (b) across all six scenarios relative to the baseline, where values below 1.0 indicate improvement.

Table 3: Fleet composition and emission changes for electrification scenarios (per direction). ΔUFP and $\Delta PM_{2.5}$ are mean quantile-ratio changes relative to the baseline.

Scenario	**Vehicles**	**Mass (t)**	**kg/pass.**	**ΔUFP (regen)**	**ΔUFP (no regen)**	**$\Delta PM_{2.5}$**
Baseline	25	59.4	1121	—	—	—
25% Electric	25	61.0	1151	−11.0%	+10.4%	+1.0%
50% Electric	25	62.8	1185	−19.0%	+12.7%	+2.6%
75% Electric	25	64.4	1215	−30.2%	+17.1%	+5.2%

The results reveal that regenerative braking is the decisive factor determining whether electrification helps or harms local UFP exposure, and the gap widens with penetration level. At 25% electrification, UFP falls by 11% with effective regenerative braking but increases by 10% without it. At 50% electrification the regenerative braking assumption becomes decisive: with 80% regenerative braking, UFP falls by 19%, driven by the recovery of braking energy across the electrified fleet; without it, the same electrification level produces a 13% UFP increase because the heavier electric vehicles (2,600 kg e-SUVs replacing 2,200 kg ICE SUVs) brake harder without energy recovery. At 75% electrification, the spread is larger still: −30% with regen versus +17% without. Whether electrification delivers a substantial benefit or a disbenefit therefore hinges entirely on regenerative braking performance, a parameter that remains poorly constrained by field data.

$PM_{2.5}$ increases slightly across all electrification levels (+1.0% at 25%, +2.6% at 50%, +5.2% at 75%), reflecting the heavier rolling mass of the electrified fleet (up to 64.4 t from 59.4 t at 75%). This steady increase is consistent with the direct mass-tyre-wear relationship: heavier electric vehicles produce more tyre-derived $PM_{2.5}$ regardless of braking technology. The near-identical $PM_{2.5}$ distributions within each regen/no-regen pair confirm that regenerative braking is irrelevant for tyre-derived particulates, as expected.

This finding reframes the policy debate. Electrification delivers meaningful UFP reductions only when regenerative braking operates effectively, a condition that varies with temperature, battery state-of-charge, driving style, and vehicle design. Meanwhile, the mass penalty of heavier electric vehicles increases tyre-derived $PM_{2.5}$ regardless of

braking technology. Speed management and modal shift, by contrast, deliver consistent reductions irrespective of vehicle technology (Sections 3.7–3.8).

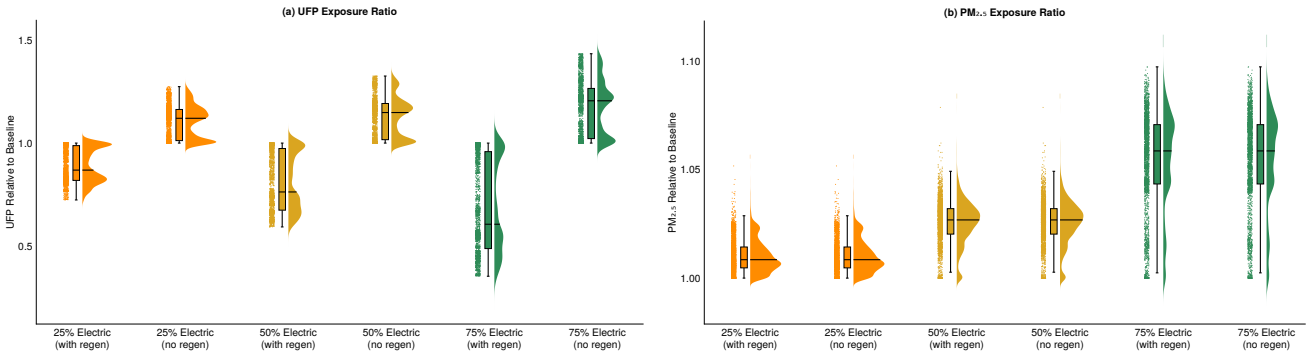

(a) UFP (brake-dominated)

(b) $PM_{2.5}$ (tyre-dominated)

Figure 4: Impact of regenerative braking assumptions across three electrification levels (25%, 50%, 75%), shown as quantile-ratios to baseline (3 seeds, sorted distributions divided element-wise; values below 1.0 indicate improvement). The UFP panel (a) shows that the regen/no-regen gap widens progressively with electrification level: a 21-percentage-point spread at 25%, a 32-percentage-point spread at 50%, and a 47-percentage-point spread at 75%. The $PM_{2.5}$ panel (b) shows a small, consistent increase across all electrification levels (+1% to +5%), reflecting the heavier mass of the electrified fleet.

3.5 Bus modal shift reduces emissions through mass consolidation

A genuinely policy-relevant question is whether shifting private vehicle passengers onto buses, while preserving total passenger throughput, reduces or increases non-exhaust emissions at the junction. The answer is not obvious: each bus weighs approximately 14 tonnes (12 t base + passengers), far more than the private vehicles it replaces, but a single bus consolidates the braking events of many private vehicles into one trajectory.

To test this, we ran a passenger-preserving bus modal shift experiment in which 25%, 50%, and 75% of private vehicle passengers were reallocated to buses, with the private fleet reduced proportionally to maintain constant total passenger throughput (Table 4). Bus occupancy increased from the baseline 22.5 passengers (25% capacity) towards 86 passengers (95% capacity) as the shift fraction grew, consistent with the service improvements that would accompany a major modal shift policy.

Table 4: Fleet composition and mass accounting for passenger-preserving bus modal shift scenarios (per direction).

Scenario	**Vehicles**	**Mass (t)**	**kg/pass.**	**ΔUFP**	**$\Delta PM_{2.5}$**
Baseline	25	59.4	1121	—	—
25% to bus	21	62.4	1148	−18.0%	−21.6%
50% to bus	13	47.5	919	−30.0%	−33.9%
75% to bus	8	49.0	949	−22.5%	−40.6%

The results demonstrate that bus modal shift delivers substantial non-exhaust emission reductions even at moderate shift fractions, but with a non-monotonic UFP response that reveals a "bus paradox" at high shift levels. At 25% modal shift, four private vehicles are removed and one bus is added (21 vehicles per direction, of which 2 are buses), and both pollutants improve substantially: UFP falls by 18% and $PM_{2.5}$ by 22%. At 50% shift, 12 private vehicles are removed and a third bus is added (13 vehicles per direction, of which 3 are buses), and the benefits peak for UFP (−30%) and $PM_{2.5}$ (−34%).

However, UFP reductions do not scale monotonically. At 75% shift, a further 5 private vehicles are removed but no additional bus is needed (8 vehicles per direction, of which 3 are buses at near-full capacity), yet UFP improvement drops back to −22.5%, less than the 50% scenario. This bus paradox arises because the three heavy buses per direction (each approximately 13 tonnes) now dominate the fleet, and the total fleet mass (49.0 t) is actually higher than at 50% shift (47.5 t). Each heavy bus generates intense braking events at the stop lines, partially offsetting the benefit of removing private vehicles. $PM_{2.5}$ continues to improve steadily (−41% at 75%) because tyre wear scales with total rolling mass and the fewer remaining vehicles reduce cumulative tyre contact.

The distributional character of the quantile-ratio shifts qualitatively across the modal shift gradient (Figure 5a for UFP, b for $PM_{2.5}$). At 25% shift, the bulk of the distribution falls below 1.0, indicating broad-based improvement. At 50% shift, the distribution compresses further below 1.0, with the strongest reductions across the full quantile range. At 75% shift, while the median remains below 1.0, a long right tail reflects the intense bus braking events that limit further UFP improvement. In other words, while the time-averaged concentration falls substantially at high bus shift fractions, the intense braking events of 13-tonne buses at stop lines produce localised spikes that may pose risks to

pedestrians and cyclists waiting at the junction.

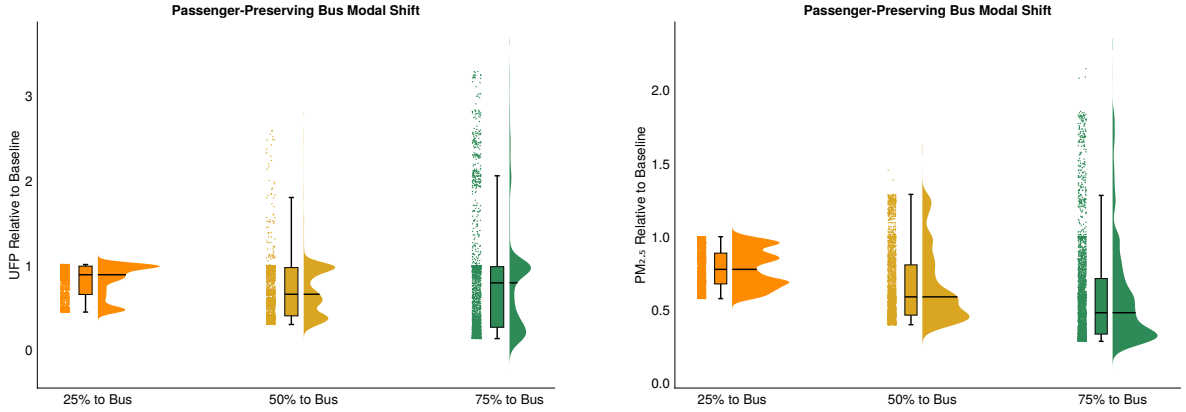

(a) UFP (brake-dominated)

(b) $PM_{2.5}$ (tyre-dominated)

Figure 5: Passenger-preserving bus modal shift: percentage of private vehicle passengers reallocated to buses while maintaining constant total throughput. Raincloud plots show quantile-ratios to baseline (3 seeds; values below 1.0 indicate improvement). All shift levels deliver substantial reductions, but UFP exhibits a non-monotonic "bus paradox": 50% shift is optimal (−30% UFP, −34% $PM_{2.5}$), while 75% shift reduces UFP by only 22.5% as additional heavy buses generate intense braking events that partially offset the benefit of removing private vehicles. $PM_{2.5}$ continues to improve progressively (−22% to −41%).

3.6 Microcar modal shift: mass reduction versus braking frequency

Unlike buses, which consolidate passengers into fewer, heavier vehicles, microcars take the opposite approach: replacing larger vehicles with lighter ones carrying fewer passengers (1,000 kg, 1.2 passengers per vehicle). In this scenario only ICE cars are replaced with microcars; the SUV fleet is held constant, reflecting the distinct consumer segment that purchases SUVs.

Table 5 presents the fleet configurations. Because only ICE cars are replaced, total mass per direction drops modestly from 59.4 to 55.4 tonnes (−7%) at 75% shift, and the per-direction vehicle count remains exactly 25 across all shift levels. Microcars are spawned at the same initial positions as the cars they replace, ensuring that the only experimental change is the vehicle mass, not the spatial configuration of traffic.

Table 5: Fleet composition for passenger-preserving microcar modal shift scenarios (per direction). Only ICE cars are replaced; SUVs, e-cars, e-SUVs, and buses remain unchanged.

Scenario	**Vehicles**	**Mass (t)**	**kg/pass.**	**ΔUFP**	**Δ$PM_{2.5}$**
Baseline	25	59.4	1121	—	—
25% to microcar	25	58.4	1105	−4.2%	−10.0%
50% to microcar	25	56.9	1081	−25.5%	−24.6%
75% to microcar	25	55.4	1057	−29.6%	−17.0%

The results reveal that microcar substitution delivers progressively increasing emission reductions with adoption level (Figure 6a, b). At 25% shift, replacing two ICE cars with two microcars yields a modest −4% UFP and −10% $PM_{2.5}$. At 50% shift, benefits grow substantially (−26% UFP, −25% $PM_{2.5}$), and at 75% shift reductions reach −30% UFP and −17% $PM_{2.5}$.

The consistent but sub-linear scaling reflects the dominance of the unchanged SUV fleet in total brake wear. At 75% adoption, eight microcars coexist with twelve unchanged SUVs (2,200 kg each), one e-SUV (2,600 kg), and one bus (∼13 tonnes); the total fleet mass drops only modestly (59.4 to 55.4 t) because SUVs remain untouched. Since brake-wear emissions scale with vehicle mass through the weight factor (m/1500 kg), the per-event emission reduction from replacing a 1,500 kg car with a 1,000 kg microcar is 33%, but this applies only to the replaced cars, which constitute a diminishing fraction of total brake wear as the SUVs continue to dominate.

$PM_{2.5}$ reductions are largest at 50% (−25%) and smaller at 75% (−17%), reflecting the interplay between reduced tyre-road contact mass and the persistent SUV fleet. The diminishing marginal returns at high adoption reflect saturation of the ICE-car replacement pool: only 10 of 25 vehicles per direction are eligible for replacement, and the remaining SUVs and bus dominate tyre wear regardless of the car segment composition.

The contrasting results of the bus and microcar strategies highlight complementary mechanisms. Bus consolidation reduces the *number* of independent braking trajectories, producing monotonically improving $PM_{2.5}$ reductions but encountering the bus paradox for UFP at high shift fractions. Microcar substitution reduces the *mass* per trajectory, delivering monotonically increasing UFP reductions with adoption but with diminishing marginal returns as SUV-dominated brake wear becomes the binding constraint. At

moderate adoption (25%), bus modal shift delivers a larger UFP benefit (−18%) than microcar substitution (−4%), because consolidating passengers into fewer vehicles removes entire braking trajectories from the junction, whereas lightweighting a small fraction of the fleet produces only a proportionate per-event emission reduction.

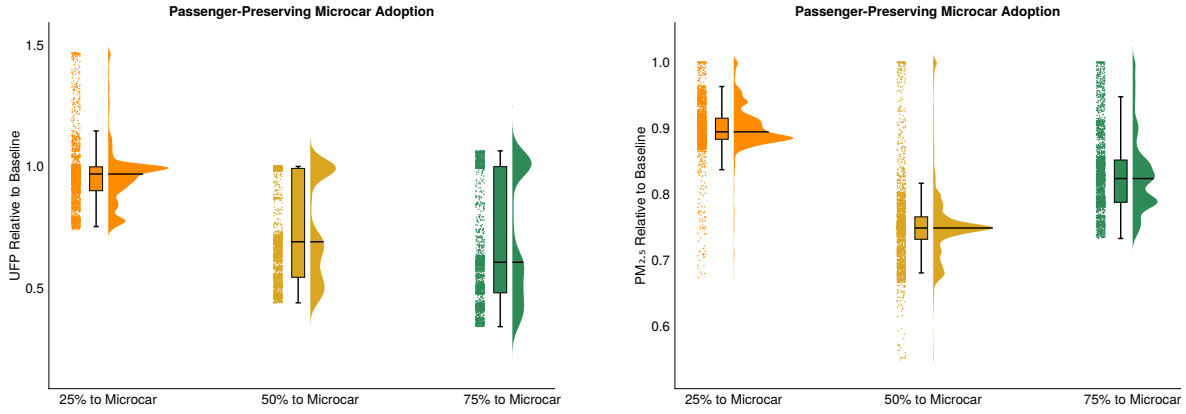

(a) UFP (brake-dominated)

(b) $PM_{2.5}$ (tyre-dominated)

Figure 6: Passenger-preserving microcar modal shift: ICE car passengers reallocated to 1,000 kg microcars while SUVs remain unchanged (quantile-ratios to baseline; 3 seeds; values below 1.0 indicate improvement). Reductions increase consistently with adoption: −4% UFP at 25%, −26% at 50%, and −30% at 75%, with diminishing marginal returns as the unchanged SUV fleet increasingly dominates brake wear.

3.7 Speed × variability: traffic calming yields compounding reductions

Mean speed (ST = 0.27) and speed variability (ST = 0.23) are among the highest-ranked parameters in the GSA, rivalling fleet composition. Their pairwise interaction (S2 = 0.013) is also one of the few non-negligible S2 values in the system. The physical mechanism is clear: mean speed determines how hard vehicles brake (kinetic energy scales as v^2), while speed variability determines how often they brake (more stop-go cycles per unit time). These two effects compound non-linearly.

To expose this interaction, we present a 7×7 parametric grid as annotated heatmaps showing percentage change in UFP (a) and $PM_{2.5}$ (b) from the calibrated baseline (Figure 7). Each cell represents the mean of 3 stochastic repeats (3,000 post-warmup samples), providing robust estimates across the entire speed-variability design space.

The heatmaps show a clear diagonal gradient. The "slow + smooth" corner (25 km/h, 0.10 m/s variability) reduces UFP by 45% and $PM_{2.5}$ by 25% relative to the calibrated baseline (33 km/h, 4.05 m/s), the largest reduction observed across all 49 scenarios. The "fast + erratic" corner (50 km/h, 2.0 m/s variability) increases UFP by 83% and $PM_{2.5}$ by 38%. The two levers are approximately additive for UFP: slowing traffic from 33 to 25 km/h alone (holding variability at baseline) reduces UFP by 42%, while the combined slow+smooth scenario delivers a 45% reduction. For $PM_{2.5}$, slowing to 25 km/h alone reduces exposure by 23%, and combining with smooth flow delivers a 25% reduction, with a modest super-additive bonus of 2 percentage points. This compounding arises because slower vehicles produce gentler stop-go transitions even at the same variability, so the two levers reinforce each other.

For $PM_{2.5}$, the gradient is pronounced across the speed axis. The tyre-wear source term (Equation 5) includes both a speed-proportional rolling component and a variability-proportional shear component, and the mass dependence means that high-variability scenarios push heavier vehicles through more severe acceleration-deceleration cycles. The "fast + erratic" corner produces mean $PM_{2.5}$ of 9.6 $\mu g/m^3$, nearly twice the WHO annual guideline of 5 $\mu g/m^3$, while "slow + smooth" traffic brings values to 4.4 $\mu g/m^3$, below this threshold.

The practical implication is that speed reduction is a powerful lever: slowing from 33 to 25–29 km/h delivers 33–42% UFP reductions and 16–23% $PM_{2.5}$ reductions at any variability level. Urban design interventions such as lane narrowing, raised tables at junctions, chicanes, and planted build-outs have been shown to reduce mean speeds while simultaneously discouraging aggressive acceleration and braking, thereby compressing the speed distribution [Air Quality Expert Group, 2019]. Adaptive signal timing that creates "green waves" for vehicles travelling at the target speed further smooths the approach profile. At Phibsborough specifically, the four-arm signalised junction with its short inter-stop distances already forces frequent deceleration; redesigning the approaches to encourage a steady 25–30 km/h arrival speed rather than a fast approach followed by a hard stop would move traffic from the centre of the heatmap toward its green corner, yielding reductions of up to 25% in $PM_{2.5}$.

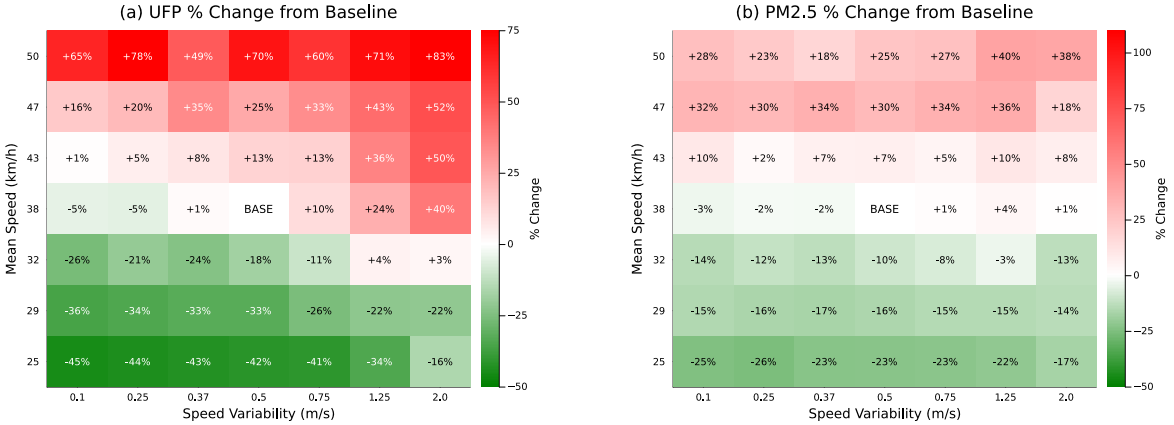

(a) UFP (brake-dominated)

(b) $PM_{2.5}$ (tyre-dominated)

Figure 7: Speed × Variability interaction heatmap: percentage change computed via quantile-ratios (sort both baseline and scenario distributions, divide element-wise, report mean ratio) across the full 7×7 parametric grid (3 repeats per cell). Green cells indicate reductions; red cells indicate increases. Speed reduction is the dominant lever: moving from 33 to 25 km/h delivers 41–45% UFP reductions across all variability levels. "BASE" marks the calibrated baseline (33 km/h, 4.05 m/s).

3.8 Policy comparison: speed management is the most robust lever

The GSA identifies fleet composition as the dominant contributor to total model variance (ST = 0.27–0.34 for vehicle counts), but variance attribution and policy effectiveness are distinct questions. Fleet purchasing trends and regenerative braking efficiencies are difficult for local transport authorities to regulate, and as the electrification scenarios demonstrate, whether their net impact on UFP is positive or negative depends on uncertain assumptions about regenerative braking. Yet the GSA also reveals that mean speed (ST = 0.27) and speed variability (ST = 0.23) are nearly as influential as fleet composition, confirming that traffic calming is a major policy lever. We therefore evaluate speed reduction and modal shift as actionable, infrastructure-level interventions alongside electrification.

Figure 8 compares the percentage change in mean moving-observer NEE exposure (UFP from brake wear, $PM_{2.5}$ from tyre wear) across all major policy options, computed via quantile-ratios relative to the calibrated baseline. The comparison includes a 30 km/h speed limit, smooth traffic flow (reduced speed variability), bus modal shift (25%), microcar modal shift (25%), and fleet electrification (50% with and without regenerative braking, the central case from Table 3). The 25% modal shift level was chosen as a policy-realistic target consistent with incremental improvements to bus service and micromobility infrastructure.

The comparison reveals that most single interventions cluster in a similar range of UFP benefit. Fleet electrification at 50% with regen (−19% UFP), bus modal shift at 25% (−18% UFP, −22% $PM_{2.5}$), and speed reduction to 30 km/h (−16.5% UFP) all deliver substantial reductions. By contrast, microcar substitution at 25% yields only −4% UFP, reflecting the modest mass change from replacing just two cars per direction, though its −10% $PM_{2.5}$ benefit is noteworthy.

Fleet electrification at 50% with regenerative braking delivers a meaningful UFP reduction (−19%), but this benefit is entirely contingent on effective regen: without it, the same electrification produces a 13% UFP increase. In both cases, $PM_{2.5}$ increases by 3% due to the heavier rolling mass of the electrified fleet. A policy whose UFP benefit depends on uncertain regenerative braking performance, and whose $PM_{2.5}$ impact is consistently adverse, is difficult to recommend as a standalone intervention.

Speed reduction is the most robust technology-independent lever: a 30 km/h limit delivers −16.5% UFP and −9.0% $PM_{2.5}$, while smooth traffic flow (reduced speed variability) delivers −15.6% UFP with near-neutral $PM_{2.5}$ (−0.9%). These benefits are independent of vehicle technology, fleet turnover timelines, and regenerative braking assumptions; they require only infrastructure design and enforcement rather than consumer behaviour change.

Finally, we tested a combined policy that applies all five interventions simultaneously: a 30 km/h speed limit with smooth traffic flow, 25% bus modal shift, 25% ICE car replacement with microcars, and 50% electrification of remaining ICE vehicles with regenerative braking (fleet per direction: 3 cars, 4 e-cars, 5 SUVs, 5 e-SUVs, 2 buses, 2 microcars; 21 vehicles; 64.7 t total). The combined scenario delivers −27% UFP and −12% $PM_{2.5}$, exceeding any single intervention for UFP and demonstrating that moderate, coordinated interventions yield compounding benefits. However, the asymmetry between the two pollutants is striking: the $PM_{2.5}$ benefit (−12%) is markedly smaller than the UFP benefit (−27%), despite bus modal shift alone delivering −22% $PM_{2.5}$. This atten-

uation arises because the 50% electrification component adds 2.2 tonnes of rolling mass per direction (heavier EVs replacing lighter ICE vehicles), increasing tyre-road contact forces and partially offsetting the $PM_{2.5}$ gains from speed reduction and modal shift. The electrification lever thus creates an inherent trade-off within the combined policy: it delivers the largest single-intervention UFP reduction (through regenerative braking) while simultaneously degrading $PM_{2.5}$ performance (through increased mass). This tension underscores the importance of mass-conscious electrification pathways and suggests that combining speed management with modal shift, without aggressive electrification, may deliver more balanced dual-pollutant benefits.

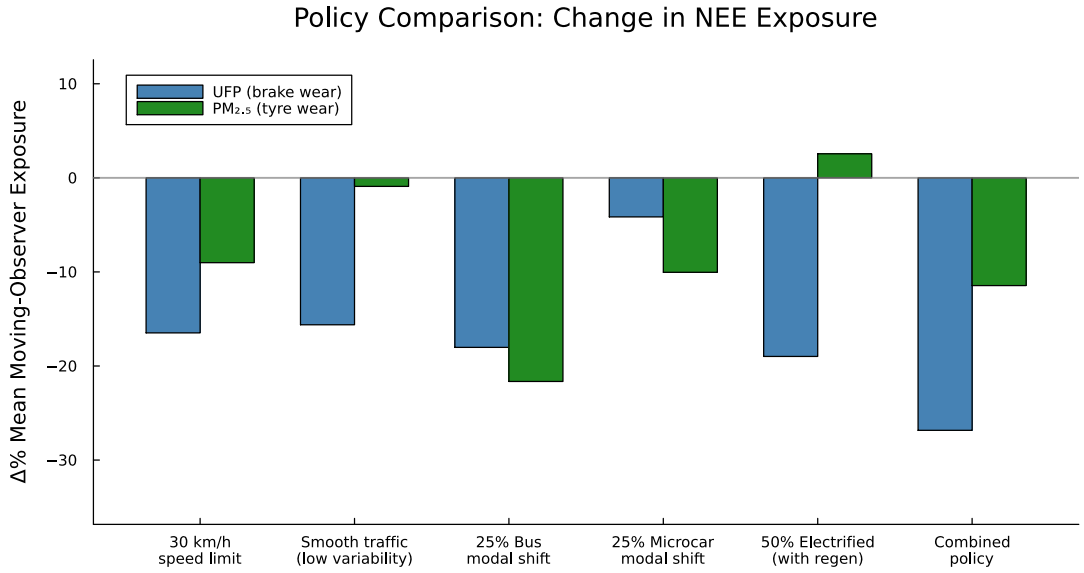

Figure 8: Policy comparison: percentage change in mean moving-observer NEE exposure relative to baseline, computed via quantile-ratios (8 seeds per scenario). UFP bars reflect brake-wear exposure; $PM_{2.5}$ bars reflect tyre-wear exposure. Single interventions cluster in a -16% to -19% UFP range (electrification with regen, bus modal shift, 30 km/h limit), except microcar substitution at 25% which yields a more modest -4% UFP. A combined policy (30 km/h + smooth traffic + 25% bus shift + 25% microcar + 50% electrification with regen) delivers -27% UFP and -12% $PM_{2.5}$, demonstrating compounding benefits from coordinated interventions.

4 Discussion

Recent advances make it feasible to combine metre-scale, driver-explicit agent-based modelling with true global sensitivity analysis at street level. This study leverages two of those advances, hyperlocal mobile sensing for calibration and polynomial-chaos surrogates for analytic Sobol indices, to pair a data-fitted ABM with a full-prior sensitivity assessment [Sudret, 2008]. In doing so, the model shifts from reproducing observations to diagnosing causal levers and ranking interventions with quantified uncertainty, an approach that, to our knowledge, remains rare for hyperlocal non-exhaust emissions.

The calibrated ABM captures both the distributional and spatial signatures of NEEs observed at Phibsborough. In line with recent syntheses, unregulated NEEs now dominate the primary particulate burden from road transport in many settings, with brake, tyre and road-dust processes supplanting shrinking tailpipe sources [Harrison et al., 2021, Piscitello et al., 2021].

The global sensitivity analysis identifies the dominant levers through the total-order Sobol indices (ST), which capture both direct and interaction effects. Fleet composition parameters rank highest (car and e-car counts ST = 0.34, e-SUV 0.30, bus 0.29, SUV 0.27), but mean speed (ST = 0.27) and speed variability (ST = 0.23) rank alongside them, confirming that traffic dynamics are as influential as fleet mix in determining non-exhaust exposure. All first-order indices fall below 0.013, indicating that the system is fundamentally interaction-dominated: the effect of any single parameter depends on the values of the others. This elevates speed management from a secondary consideration to a first-order policy lever.

Speed management delivers the most robust technology-independent reductions: a 30 km/h limit reduces UFP by 16.5% and $PM_{2.5}$ by 9.0% relative to the calibrated baseline, while smooth traffic flow (reduced speed variability) delivers −15.6% UFP. Speed-limit interventions require no assumptions about regenerative braking efficiency, fleet turnover rates, or consumer vehicle choices. This finding aligns with international policy frameworks: the WHO's "Streets for Life" campaign recommends default 30 km/h limits in urban areas where vulnerable road users mix with motorised traffic [World Health Organization, 2021], and the UK Air Quality Expert Group's review identifies brake wear as speed-dependent and concentrated in areas with frequent deceleration [Air Quality Expert Group, 2019].

By contrast, fleet electrification, the centrepiece of most national transport decarbonisation strategies, emerges as an unreliable lever for NEE reduction. The electrification sweep (Table 3) shows that the gap between optimistic and pessimistic regenerative braking scenarios widens with penetration level: at 25% electrification, the UFP effect ranges from −11% (with regen) to +10% (without); at 50%, the spread is −19% to +13%; and at 75%, it reaches −30% to +17%. Real-world regenerative braking efficiency is poorly constrained: laboratory tests report up to 94.7% brake event coverage, but field studies show five-fold variability between driving styles, and cold temperatures and high battery charge degrade performance substantially [Storch et al., 2023, Timmers and Achten, 2016, Costagliola et al., 2024]. Current Irish incentives apply uniformly to all BEVs regardless of mass, inadvertently subsidising the purchase of heavier electric SUVs (e.g., replacing a 1,500 kg ICE car with a 2,600 kg electric SUV) that increase tyre-derived $PM_{2.5}$ regardless of braking technology. Mass-graduated incentives that steer consumers toward lighter EVs would partially mitigate this effect.

The modal shift analysis reveals two complementary mechanisms. Bus consolidation

reduces the number of independent braking trajectories, delivering substantial reductions at all shift fractions (−18% UFP at 25%, −30% at 50%), though a bus paradox emerges at 75% where the third heavy bus per direction (∼13 tonnes each) generates intense braking events that reduce the UFP benefit to −22.5%, while $PM_{2.5}$ continues to improve (−41%). Microcar substitution operates through a different channel: reducing the mass of each braking event rather than the number of events. Benefits scale consistently with adoption (−4% UFP at 25%, −26% at 50%, −30% at 75%) but with diminishing marginal returns, as the unchanged SUV fleet (12 vehicles at 2,200 kg each) increasingly dominates total brake wear. At moderate adoption (25%), bus modal shift delivers a four-fold larger UFP benefit than microcar substitution (−18% versus −4%), because removing entire vehicles from the traffic stream eliminates their braking trajectories entirely, whereas lightweighting reduces per-event emissions only in proportion to the mass change. The distributional character of exposure also differs qualitatively: bus consolidation lowers the average but introduces occasional intense spikes from heavy bus braking, while microcar substitution shifts the entire distribution leftward. Whether these distinct exposure profiles carry different health implications remains an open question.

Beyond air, the junction's NEE hotspots have implications for urban water quality. Metal-rich brake wear and rubber fragments can be mobilised by rainfall and washed into kerb inlets. Sustainable drainage systems (SuDS), such as bioretention cells and permeable paving, can intercept these pollutants before they enter the drainage network [Li and Davis, 2008, Mufidah and Soewondo, 2023, Woods Ballard et al., 2015, Kong et al., 2025]. Because the radial profile analysis (Section 2.8.3) confirms that brake wear mass concentrates within the first 10–15 m of the stop lines, the ABM can identify exactly where SuDS installations would be most effective, enabling hyper-targeted placement immediately adjacent to signalised stop lines rather than generic drainage upgrades across the entire junction.

Three extensions are natural. Exhaust emissions can be incorporated by coupling the per-vehicle speed and acceleration traces to established emission factor models (e.g., HBEFA or COPERT), enabling joint assessment of exhaust and non-exhaust contributions and quantifying how the balance shifts as the fleet electrifies. Noise can be added by coupling the same kinematic fields to tyre/road and powertrain noise models, enabling joint appraisal of air and acoustic exposures at the frontage in light of WHO and EEA guidance [World Health Organization. Regional Office for Europe, 2018, European Environment Agency, 2020]. Dust resuspension, the mechanical lifting of deposited material by tyres and vehicle wakes, can be incorporated by evolving a surface reservoir on the grid with wind- and traffic-dependent entrainment [Thorpe and Harrison, 2008, U.S. Environmental Protection Agency, 2006, Piscitello et al., 2021]; although Dublin's frequent rainfall limits the magnitude of this source, it may become relevant in extended dry spells.

Several limitations should be noted. The baseline fleet composition draws on NTA AM-peak canal cordon counts, the best publicly available modal split data for central Dublin, while the Air View calibration data were collected predominantly during the afternoon. Afternoon traffic typically features lower bus ridership and higher private-car dominance than the morning peak, so the baseline may slightly overstate bus presence relative to the measurement period; this makes the bus scenario results conservative in the sense that the modelled bus contribution is an upper bound. The dispersion model captures vehicle-wake advection via a conal kernel but does not include an ambient wind field. In the narrow street-canyon geometry of Phibsborough, vehicle-generated turbulence dominates the near-source concentrations sampled by the moving observer, so

the omission is unlikely to materially affect the calibrated emission factors. Similarly, the model does not subtract an urban background concentration; however, the calibration targets are junction-specific Air View measurements that are themselves net of any city-wide baseline, and the quantile-matching loss is sensitive to distributional shape rather than absolute offset, limiting the scope for systematic inflation of the fitted emission coefficients.

The model suggests a practical sequence for managing NEEs at signalised hotspots: first, reduce approach speeds to 30 km/h through design and enforcement that encourages smoother traffic flow, capturing the guaranteed technology-independent benefit (−16.5% UFP); second, implement moderate modal shift that reallocates passengers to buses (−18% UFP at 25% bus shift) and, where feasible, replaces ICE car trips with lighter vehicles (−4% UFP at 25% microcar adoption, rising to −30% at 75%), as identified in Sections 3.5 and 3.6; and third, as electrification proceeds, steer uptake toward lighter vehicles with effective regenerative braking and low-abrasion tyres [Timmers and Achten, 2016, Costagliola et al., 2024]. Our combined policy simulation confirms that these interventions produce compounding benefits when applied together: a 30 km/h speed limit with smooth traffic flow, 25% bus modal shift, 25% microcar substitution, and 50% electrification with regenerative braking collectively reduce UFP by 27% and $PM_{2.5}$ by 12%, exceeding any single intervention and demonstrating that moderate, coordinated action across multiple levers is the most effective strategy. While Euro 7 regulations will eventually constrain brake and tyre wear at the manufacturing level [The International Council on Clean Transportation, 2024], the fleet turnover required to realise those benefits will take a decade or more. Our results show that local authorities need not wait: speed limits, bus lanes, and targeted SuDS at ABM-identified hotspots can reduce both airborne exposure and pollutant transfer to the drainage network today [Woods Ballard et al., 2015].

Acknowledgements

The author thanks Leon Cheung for early work on this project during the RCSI Research Summer School.

References

Aclima & Google. Google Project Air View Data - Dublin City (May 2021 - August 2022), 2022. URL https://data.gov.ie/dataset/google-airview-data-dublin-city. Accessed 26th February 2026. Licensed under Creative Commons Attribution 4.0.

Air Quality Expert Group. Non-exhaust emissions from road traffic. Technical report, Department for Environment, Food and Rural Affairs, London, UK, 2019. URL https://uk-air.defra.gov.uk/assets/documents/reports/cat09/1907101151_20190709_Non_Exhaust_Emissions_typeset_Final.pdf.

H. Barosova, S. Chortarea, P. Peikertova, M. J. D. Clift, A. Petri-Fink, J. Kukutschova, and B. Rothen-Rutishauser. Biological response of an in vitro human 3d lung cell model exposed to brake wear debris varies based on brake pad formulation. *Archives of Toxicology*, 92(7):2339–2351, 2018. doi:10.1007/s00204-018-2218-8.

Tomás P. Carroll, Catherine A. O'Connor, Olwen Floyd, Joseph McPartlin, Dermot P. Kelleher, Geraldine O'Brien, Borislav D. Dimitrov, Valerie B. Morris, Clifford C. Taggart, and Noel G. McElvaney. The prevalence of alpha-1 antitrypsin deficiency in Ireland. *Respiratory Research*, 12:91, 2011. doi:10.1186/1465-9921-12-91.

Central Statistics Office. Transport Environmental Indicators Ireland 2022. Technical report, CSO Ireland, November 2022a. URL https://www.cso.ie/en/releasesandpublications/ep/p-eii/environmentalindicatorsireland2022/transport/.

Central Statistics Office. Sustainable mobility and transport 2021. Technical report, Central Statistics Office, Cork, Ireland, 2022b. URL https://www.cso.ie/en/releasesandpublications/ep/p-smt/sustainablemobilityandtransport2021/electricvehicles/. Accessed: 23 October 2025.

Central Statistics Office. Census of Population 2022 – Summary Results: Employment, Occupation, Industry and Commuting. Technical report, CSO Ireland, 2023a. URL https://www.cso.ie/en/releasesandpublications/ep/p-cpsr/censusofpopulation2022-summaryresults/employmentoccupationindustryandcommuting/.

Central Statistics Office. Vehicles licensed for the first time december and year 2022. Technical report, Central Statistics Office, Cork, Ireland, January 2023b. URL https://www.cso.ie/en/releasesandpublications/ep/p-vlftm/vehicleslicensedforthefirsttimedecemberandyear2022/. Accessed: 23 October 2025.

Committee on the Medical Effects of Air Pollutants (COMEAP). Statement on the evidence for health effects associated with exposure to non-exhaust particulate matter from road transport. Technical report, Public Health England, 2020. URL https://www.gov.uk/government/publications/non-exhaust-particulate-matter-from-road-transport-health-effects.

Maria Antonietta Costagliola, Luca Marchitto, Rocco Giuzio, Simone Casadei, Tommaso Rossi, Simone Lixi, and Davide Faedo. Non-exhaust particulate emissions from road transport vehicles. *Energies*, 17(16):4079, 2024. doi:10.3390/en17164079.

Thierry Crestaux, Olivier Le Maître, and Jean-Marc Martinez. Polynomial chaos expansion for sensitivity analysis. *Reliability Engineering & System Safety*, 94(7):1161–1172, 2009. doi:10.1016/j.ress.2008.10.008.

Camelia Delcea and Nora Chiriţă. Exploring the applications of agent-based modeling in transportation. *Applied Sciences*, 13(17):9815, 2023. doi:10.3390/app13179815.

Department of Health. National healthcare quality reporting system annual report 2020. Technical report, Government of Ireland, Dublin, Ireland, 2020. URL https://www.gov.ie/en/department-of-health/collections/national-healthcare-quality-reporting-system-reports/. COPD hospitalisation rate highest among 37 OECD countries.

Department of Transport. Irish bulletin of vehicle and driver statistics 2022. Technical report, Government of Ireland, Dublin, Ireland, 2023. URL https://www.gov.ie/en/publication/f392d-bulletin-of-vehicle-and-driver-statistics/. Accessed: 23 October 2025.

European Environment Agency. Environmental noise in europe – 2020. Technical Report EEA Report No 22/2019, Publications Office of the European Union, 2020.

M. Figliuzzi, M. Tironi, L. Longaretti, A. Mancini, F. Teoldi, F. Sangalli, and A. Remuzzi. Copper-dependent biological effects of particulate matter produced by brake systems on lung alveolar cells. *Archives of Toxicology*, 94(9):2965–2979, 2020. doi:10.1007/s00204-020-02812-4.

Miriam E. Gerlofs-Nijland, Bas G. H. Bokkers, Henriette Sachse, Jan J. E. Reijnders, Mats Gustafsson, A. John F. Boere, Paul F. H. Fokkens, Daan L. A. C. Leseman, Klaus Augsburg, and Flemming R. Cassee. Inhalation toxicity profiles of particulate matter: a comparison between brake wear with other sources of emission. *Inhalation Toxicology*, 31(3):89–98, 2019. doi:10.1080/08958378.2019.1606365.

Theodoros Grigoratos and Giorgio Martini. Non-exhaust traffic related emissions – brake and tyre wear PM. Technical Report JRC89231, Joint Research Centre, European Commission, 2014. URL https://publications.jrc.ec.europa.eu/repository/handle/JRC89231.

Sashikanth Gurram, Amy L. Stuart, and Abdul R. Pinjari. Agent-based modeling to estimate exposures to urban air pollution from transportation: Exposure disparities and impacts of high-resolution data. *Computers, Environment and Urban Systems*, 75: 22–34, 2019. doi:10.1016/j.compenvurbsys.2019.01.002.

Roy M Harrison, James Allan, David Carruthers, Mathew R Heal, Alastair C Lewis, Ben Marner, Tim Murrells, and Andrew Williams. Non-exhaust vehicle emissions of particulate matter and VOC from road traffic: A review. *Atmospheric Environment*, 262:118592, 2021. doi:10.1016/j.atmosenv.2021.118592.

Friederike Hülsmann, Regine Gerike, and Matthias Ketzel. Modelling traffic and air pollution in an integrated approach – the case of munich. *Urban Climate*, 10:732–744, 2014. doi:10.1016/j.uclim.2014.01.001.

Ihab Kaddoura, Ricardo Ewert, and Kai Martins-Turner. Exhaust and non-exhaust emissions from today's and future road transport: A simulation-based quantification for Berlin. *Transportation Research Procedia*, 62:696–702, 2022. doi:10.1016/j.trpro.2022.02.086.

Hanna L. Karlsson, Anders G. Ljungman, Jens Lindbom, and Lennart Möller. Comparison of genotoxic and inflammatory effects of particles generated by wood combustion, a road simulator and collected from street and subway. *Toxicology Letters*, 165(3): 203–211, 2006. doi:10.1016/j.toxlet.2006.04.003.

Jiwon Kong, Seongeom Jeong, Jieun Lee, and Sanghyun Jeong. Permeable pavement blocks as a sustainable solution for managing microplastic pollution in urban stormwater. *Science of The Total Environment*, 966:178649, 2025. doi:10.1016/j.scitotenv.2025.178649.

Daniel Krajzewicz, Stefan Hausberger, Peter Wagner, Michael Behrisch, and Mario Krumnow. Second generation of pollutant emission models for SUMO. In Michael Behrisch and Melanie Weber, editors, *Modeling Mobility with Open Data: 2nd SUMO Conference 2014*, pages 203–221. Springer, 2015. doi:10.1007/978-3-319-15024-6_12.

Ho Yeung Lam, Sisitha Jayasinghe, Kiran D. K. Ahuja, and Andrew P. Hills. Agent-based modelling in active commuting research: a scoping review. *Critical Public Health*, 35 (1):66–83, 2025. doi:10.1080/09581596.2025.2483473.

Houng Li and Allen P. Davis. Heavy metal capture and accumulation in bioretention media. *Environmental Science & Technology*, 42(14):5247–5253, 2008. doi:10.1021/es702681j.

M. J. Lubczyńska, R. L. Muetzel, H. El Marroun, X. Basagaña, M. Strak, W. Denault, V. W. V. Jaddoe, M. Hillegers, M. W. Vernooij, G. Hoek, T. White, B. Brunekreef, H. Tiemeier, and M. Guxens. Exposure to air pollution during pregnancy and childhood, and white matter microstructure in preadolescents. *Environmental Health Perspectives*, 128(2):027005, 2020. doi:10.1289/EHP4709.

Ed Manley and Tao Cheng. Exploring the role of spatial cognition in predicting urban traffic flow through agent-based modelling. *Transportation Research Part A: Policy and Practice*, 109:14–23, 2018. doi:10.1016/j.tra.2018.01.020.

Qi Meng, Jonathan Liu, Jiaqi Shen, Irish Del Rosario, Pascale S. J. Lakey, Manabu Shiraiwa, Jason Su, Scott Weichenthal, Yifang Zhu, Farzan Oroumiyeh, Suzanne E. Paulson, Michael Jerrett, and Beate Ritz. Fine particulate matter metal composition, oxidative potential, and adverse birth outcomes in Los Angeles. *Environmental Health Perspectives*, 131(10):107012, 2023. doi:10.1289/EHP12196.

Douglas C. Montgomery, Elizabeth A. Peck, and G. Geoffrey Vining. *Introduction to Linear Regression Analysis*. Wiley, 5 edition, 2012. ISBN 978-0-470-54281-1.

Fathiya Mufidah and Prayatni Soewondo. Microplastic removal from road stormwater runoff using lab-scale bioretention cell. *Journal of Engineering and Technological Sciences*, 55(5):524–537, 2023. doi:10.5614/j.eng.technol.sci.2023.55.5.2.

National Transport Authority. Canal cordon report 2022: Report on trends in mode share of vehicles and people crossing the canal cordon 2006-2022. Technical report, National Transport Authority, Dublin, Ireland, May 2023a. URL https://www.nationaltransport.ie/wp-content/uploads/2023/06/NTA-Canal-Cordon-Report-2022-Final.pdf. Accessed: 23 October 2025.

National Transport Authority. 2022 m50 cordon report: Report on inbound people movements across the m50 cordon. Technical report, National Transport Authority, Dublin, Ireland, November 2023b. URL https://www.nationaltransport.ie/wp-content/uploads/2024/11/2022-M50-Cordon-Report.pdf. Accessed: 23 October 2025.

Seyed Hossein Hosseini Nourzad, Dario D. Salvucci, and Anu Pradhan. Computational modeling of driver distraction by integrating cognitive and agent-based traffic simulation models. In *Computing in Civil and Building Engineering (2014)*, pages 1885–1892, 2014. doi:10.1061/9780784413616.234.

Organisation for Economic Co-operation and Development. Non-exhaust particulate emissions from road transport: An ignored environmental policy challenge. Technical report, OECD Publishing, Paris, 2020.

Bart Ostro, Aurelio Tobias, Xavier Querol, Andrés Alastuey, Fulvio Amato, Jorge Pey, Noemí Perez, and Jordi Sunyer. The effects of particulate matter sources on daily mortality: a case-crossover study of Barcelona, Spain. *Environmental Health Perspectives*, 119(12):1781–1787, 2011. doi:10.1289/ehp.1103618.

Amelia Piscitello, Carlo Bianco, Alessandro Casasso, and Rajandrea Sethi. Non-exhaust traffic emissions: Sources, characterization, and mitigation measures. *Science of The Total Environment*, 766:144440, 2021. doi:10.1016/j.scitotenv.2020.144440.

David Q. Rich, Wangjian Zhang, Shao Lin, Stefania Squizzato, Sally W. Thurston, Edwin van Wijngaarden, Daniel Croft, Mauro Masiol, and Philip K. Hopke. Triggering of cardiovascular hospital admissions by source specific fine particle concentrations in urban centers of New York State. *Environment International*, 126:387–394, 2019. doi:10.1016/j.envint.2019.02.018.

Hyesop Shin. Quantifying the health effects of exposure to non-exhaust road emissions using agent-based modelling (ABM). *MethodsX*, 9:101673, 2022. doi:10.1016/j.mex.2022.101673.

Hyesop Shin and Mike Bithell. TRAPSim: An agent-based model to estimate personal exposure to non-exhaust road emissions in central Seoul. *Computers, Environment and Urban Systems*, 99:101894, 2023. doi:10.1016/j.compenvurbsys.2022.101894.

R. B. Smith, S. D. Beevers, J. Gulliver, D. Dajnak, D. Fecht, M. Blangiardo, M. Douglass, A. L. Hansell, H. R. Anderson, F. J. Kelly, and M. B. Toledano. Impacts of air pollution and noise on risk of preterm birth and stillbirth in London. *Environment International*, 134:105290, 2020. doi:10.1016/j.envint.2019.105290.

Society of the Irish Motor Industry. 2022 new car registrations marginally ahead, 81% uplift in electric vehicles. Technical report, Society of the Irish Motor Industry, Dublin, Ireland, January 2023. URL https://www.simi.ie/en/news/

2022-new-car-registrations-marginally-ahead-81-uplift-in-electric-vehicles. Accessed: 23 October 2025.

Lukas Storch, Christopher Hamatschek, David Hesse, Felix Feist, Thomas Bachmann, Philipp Eichler, and Theodoros Grigoratos. Comprehensive analysis of current primary measures to mitigate brake wear particle emissions from light-duty vehicles. *Atmosphere*, 14(4):712, 2023. doi:10.3390/atmos14040712.

Bruno Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964–979, 2008. doi:10.1016/j.ress.2007.04.002.

X. Sun, X. Luo, C. Zhao, B. Zhang, J. Tao, Z. Yang, W. Ma, and T. Liu. The associations between birth weight and exposure to fine particulate matter (PM2.5) and its chemical constituents during pregnancy: A meta-analysis. *Environmental Pollution*, 211:38–47, 2016. doi:10.1016/j.envpol.2015.12.022.

The International Council on Clean Transportation. Euro 7: The New Emission Standard for Light- and Heavy-Duty Vehicles in the European Union. Technical report, ICCT, March 2024. URL https://theicct.org/publication/euro-7-emission-standard-ldv-hdv-eu-mar24/. Policy Update.

Alistair Thorpe and Roy M Harrison. Sources and properties of non-exhaust particulate matter from road traffic: a review. *Science of the Total Environment*, 400(1-3):270–282, 2008. doi:10.1016/j.scitotenv.2008.06.007.

Victor RJH Timmers and Peter AJ Achten. Non-exhaust pm emissions from electric vehicles. *Atmospheric Environment*, 134:10–17, 2016. doi:10.1016/j.atmosenv.2016.03.017.

Panagiotis G. Tzouras, Lambros Mitropoulos, Eirini Stavropoulou, Eleni Antoniou, Katerina Koliou, Christos Karolemeas, Antonis Karaloulis, Konstantinos Mitropoulos, Marilena Tarousi, Eleni I. Vlahogianni, and Konstantinos Kepaptsoglou. Agent-based models for simulating e-scooter sharing services: A review and a qualitative assessment. *International Journal of Transportation Science and Technology*, 12(1):71–85, 2023. doi:10.1016/j.ijtst.2022.02.001.

U.S. Environmental Protection Agency. AP-42, Fifth Edition, Volume I, Section 13.2.1: Paved Roads. Technical report, Office of Air Quality Planning and Standards, U.S. Environmental Protection Agency, November 2006. URL https://www.epa.gov/air-emissions-factors-and-quantification/ap-42-section-1321-paved-roads-related-information-0.

Ondine S. von Ehrenstein, Hilary Aralis, Myles Cockburn, and Beate Ritz. In utero exposure to toxic air pollutants and risk of childhood autism. *Epidemiology*, 25(6):851–858, 2014. doi:10.1097/EDE.0000000000000150.

Alice M. Wood, Roy M. Harrison, Sean Semple, Jon G. Ayres, and Robert A. Stockley. Outdoor air pollution is associated with disease severity in α_1-antitrypsin deficiency. *European Respiratory Journal*, 34(2):346–353, 2009. doi:10.1183/09031936.00087908.

B. Woods Ballard, S. Wilson, H. Udale-Clarke, S. Illman, T. Scott, R. Ashley, and R. Kellagher. *The SuDS Manual (C753)*. CIRIA, London, 2015. URL https://www.ciria.org/CIRIA/Books/Free_publications/C753F.aspx.

World Health Organization. Streets for life: #love30 campaign. https://www.who.int/campaigns/un-global-road-safety-week, 2021.

World Health Organization. Regional Office for Europe. *Environmental noise guidelines for the European Region*. WHO Regional Office for Europe, 2018. ISBN 978-92-890-5356-3. URL https://www.who.int/europe/publications/i/item/9789289053563.

Lu Zhou, Yingmin Tao, Huichu Li, Yue Niu, Liang Li, Haidong Kan, Juan Xie, and Renjie Chen. Acute effects of fine particulate matter constituents on cardiopulmonary function in a panel of COPD patients. *Science of The Total Environment*, 770:144753, 2021. doi:10.1016/j.scitotenv.2020.144753.