Finding the Balance Between Efficiency and Green:

The Pareto Frontier for Time-Sensitive B2B Logistics

Christian Martin Crivelli

Student ID: i6347028

School of Business and Economics

Maastricht University

Maastricht, Netherlands

Supervisor: Dr. Burak Can

Submission Date: June 28, 2026

Abstract

This paper aims to tackle the operational trade-off between time-based efficiency and environmental sustainability in time-sensitive B2B logistics. While traditional logistics frameworks largely rely on speed averages that overlook Newtonian-based emission spikes, this study establishes a multi-objective framework to answer: "How can we quantify the exact trade-off between travel time and CO_2 emissions using microscopic mobility data to make data-driven, operationally viable logistics decisions in time-sensitive B2B fleet management?"

With a white-box model that uses both the ISO 23795 standard and the WLTP classifier, this paper makes use of a granular dataset that includes 19 trips across Europe and Aruba and by implementing a Random Forest classifier based on key metrics to ensure data-quality for modelling. Then, a segment-based physics simulation engine was constructed to evaluate one hundred unique driver profiles scaled by a parameter ($\alpha \in [0, 1]$), Pareto frontier across variations in driver behavior and, later, vehicle cargo weight.

The empirical results reveal three distinct behavioral regimes along the estimated Pareto frontier, shifting across different optimization bounds:

- **The Eco-Friendly Bound:** Under this operational regime, average vehicle velocities decrease to approximately 85% of the baseline routing strategy. For unladen vehicles, this reduction minimizes mechanical resistance, resulting in an average network-wide emission reduction between 7.2% and 7.4%. However, this environmental gain introduces a significant travel time penalty ranging from 25% to 30%.
- **The Speed-Optimized Bound:** Conversely, prioritizing transit minimization increases speeds up to 120% of standard baselines. The resulting frontier curvature demonstrates a highly non-linear penalty structure: marginal gains in travel time savings trigger disproportionately large emission spikes, driven by aggressive acceleration phases and heightened late-braking penalties.
- **The Balanced Frontier Knee:** This segment identifies the critical operational "sweet spot" on the frontier. It represents the optimal trade-off location where

fleet managers can satisfy strict corporate Service Level Agreements (SLAs) while simultaneously capturing substantial emissions abatement.

Furthermore, this study verifies that eco-routing efficiency scales non-linearly with vehicle mass. Transitioning from an aggressive driving profile to a pure eco-profile yields a 7.74% carbon reduction under unladen conditions. Due to mass-dependent kinetic energy requirements during transient cycles, this effect compounds to a 9.36% reduction throughout the network when vehicles are fully loaded. This compounding efficiency peaks within high-density urban stop-and-go corridors, delivering localized carbon savings of up to 16.51%.

As EU ETS-2 regulations transform carbon mitigation into a financial imperative, traditional macroscopic routing proves to be insufficient. This thesis demonstrates that data-driven B2B fleet management must target the frontier knee ($\alpha \in [0.4, 0.6]$) and prioritize eco-driving behavior during peak payload allocations and within congested urban networks. Transitioning to auditable white-box frameworks ensures regulatory compliance while maintaining financial viability.

Contents

1 Introduction **5**

1.1 Eco-Routing and Eco-Driving 5

1.2 Research Objective and Question 6

2 Literature Review **6**

2.1 The Theoretical Framework 7

2.2 The ISO 23795 8

3 Data Methodology **9**

3.1 WLTP Profiles Standardization 10

3.2 Random Forest Modeling 11

3.3 Simulation Framework 11

3.3.1 Initial Vehicle and Road Set Up 11

3.3.2 Driving Profile Generation 13

3.4 The Physics Engine 14

4 Results **17**

4.1 The Effect of Cargo 18

5 Discussion **19**

5.1 Interpreting the Results 19

5.2 Comparison with Existing Literature 20

5.3 Limitations and Future Work 21

5.4 Further Considerations 22

6 Conclusion **22**

6.1 Recommendations 22

A Data Structure **26**

B Trip Information **28**

C Pareto Frontiers **29**

C.1 Unladen Trips 30

C.2 Laden Trips 31

List of Figures

1 GPS Altitude Noise Reduction via 5-Step Centered Rolling Mean for Trip 7 . 13

2 The 4 Layers of the Physics Engine 15

3 Pareto Efficiency frontiers comparing Time (s) versus CO2 emissions (g) across the unladen trip profiles. 30

4 Pareto Efficiency frontiers comparing Time (s) versus CO2 emissions (g) across the laden trip profiles. 31

List of Tables

1 Operational Speed Profiles and Vehicular Efficiency Dynamics 10

2 Operational Performance Across Extreme Driving Bounds 18

4 Summary of Dataset Attributes and Physical Profiles 26

5 Summary of Individual Trip Characteristics 28

1 Introduction

In a time where business, governments and consumers are becoming more environmentally conscious, businesses must strive to become greener. This is especially true of businesses who operate in logistics and high-velocity supply chains; these commercial fleets must balance small delivery windows and pressure to maximize emissions with a rigid regulatory and economic imperative focusing on reducing carbon emissions. The expansion of the European Union's Emission Trading System (ETS-2) has established an estimated €80 billion market for Certified Emission Reduction (CER) certificates, transforming carbon mitigation from an environmental metric into a core determinant of corporate financial viability.

At the intersection of these competing forces lies a point of friction: the trade-off between time-based efficiency and environmental sustainability. While traditional routing algorithms minimize travel time, pure eco-routing engines minimize energy demand.

1.1 Eco-Routing and Eco-Driving

To evaluate these multi-objective dynamics, a clear distinction must be established between macroscopic path selection and microscopic vehicle operation. The former refers to eco-routing, which focuses on finding the path that requires the least emissions to traverse, typically measured in the form of used fuel. And the latter refers to eco-driving, which focuses on making small changes in driving behaviour that lead to reduced fuel consumption.

Traditional eco-routing can significantly lower baseline fuel consumption by avoiding high-resistance segments, but it frequently introduces severe, unacceptable increases in total transit duration. If an eco-routing or eco-driving framework operates at a pace that is too slow, commercial dispatchers and drivers will simply reject it, destroying its real-world utility.

This is the very trade-off this thesis aims to map, transitioning from traditional time-optimal routing to pure, unconstrained eco-routing can reduce network-wide fuel consumption and carbon mass emissions by an average of 3.3% to 9.3%. However, this environmental floor consistently and severely penalizes travel times by 20% to 30%[4]. These penalties are far too costly for businesses to embrace eco-routing.

1.2 Research Objective and Question

The primary objective of this thesis is to resolve this paradox by constructing a multi-objective Pareto optimization framework. By leveraging an empirical, high-resolution Global Navigation Satellite System (GNSS) mobility dataset alongside a deterministic, physics-based vehicle emissions model, this study maps the continuous boundary where operational transit velocity meets environmental friction. At the end of this paper, the following question should be answered: "How can we quantify the exact trade-off between travel time and CO_2 emissions using microscopic mobility data to make data-driven, operationally viable logistics decisions in time-sensitive B2B fleet management?"

2 Literature Review

Green logistics are no longer just a corporate social responsibility goal; it is a regulatory mandate. The expansion of the EU's ETS-2 emission trading framework has created an estimated €80 billion market for Certified Emission Reduction (CER) certificates[5].

Some older research focused on solving for either the fastest solution or the greenest one, but optimizing solely for emissions results in not viable time penalties. The current computer science literature is actively shifting towards multi-objective models to find more realistic, robust, and viable solutions[6].

While a multi-objective solution makes sense, it is not immediately intuitive how to map the trade-offs between time and emissions. Some current frameworks use a weighting mechanism such as *alpha* (α) or soft constraints modeled via sigmoid functions to balance conflicting link costs[7]. Though this is difficult because as time constraints are introduced, the emission reduction potential for time-based logistics drops drastically compared to unconstrained, distance-based logistics[8].

And the problem becomes even more complex when attempting to solve for a fleet, rather than just one vehicle. Fleet dispatchers cannot use binary fuel-efficient or time-efficient choices. Though while fleets introduce complexity, they allow for greater variability in the multi-objective model's result, that is because the speed cost of being eco-friendly can be diluted throughout the fleet. Still, businesses need tools and greater proof that these eco-friendly practices are viable.

A great issue in eco-routing and eco-driving literature is the prominence of macroscopic and mesoscopic models, that is, models that don't use fine, granular, second-by-second data, but instead use speed averages, or trip average data to estimate emissions[9]. Authors reveal this is especially hurtful when estimating acceleration or acceleration events at intersections, highway ramps, and traffic bottlenecks[9]].

Despite the rigor of existing studies, this paper addresses three critical gaps in the current literature: the lack of microscopic data, the omission of hard operational constraints, and the reliance on static rather than dynamic driving profiles. Even though the theoretical benefits of eco-routing are well-established, there is a lack of understanding regarding system performance under hard business-to-business (B2B) constraints, high-resolution emission modeling, and fluctuating driver behaviors. This thesis bridges these gaps by proposing a multi-objective Pareto optimization framework, underpinned by an empirical, high-resolution GNSS mobility dataset and a physics-based emission model compliant with the ISO 23795.

2.1 The Theoretical Framework

Throughout this study, there is an underlying assumption that travel time and fuel consumption are non-aligned objectives[10].

Furthermore, rather than utilizing static driving profiles, this study will make use of an α variable much like Wang, Elbery, and Rakha [2], this parameter ($\alpha \in [0,1]$) will map the continuous Pareto frontier. Defining specific α thresholds allows the algorithm to represent varying degrees of strictness for B2B Service Level Agreements. To further accentuate behavioral dynamics, a thorough standardized driving profile classification supported by the ISO 23795 will be included, this is the Worldwide Harmonized Light Vehicles Test Procedure. The reason driving behavior is so important is that systematic reviews prove that driving style and aggressiveness are dominant variables in transient emission spikes[4].

This study also uses microscopic data, as will be outlined in the data section of this paper. Because, as the literature calls for, rather than relying on inaccurate average speeds, the model calculates instantaneous power demands based on high-resolution, second-by-second GNSS data.

2.2 The ISO 23795

This thesis rejects machine learning (black-box) emission models popular in recent literature. Black-box models fail in commercial logistics because they cannot accurately adapt when physical variables—such as a heavy B2B truck's payload—change dynamically mid-route[]. And therefore, this paper uses a white-box model and the core of this research is anchored in the ISO 23795 standard for eco-routing.

This highly rigorous framework bridges the gap between theoretical computer science algorithms and the real-world physical and regulatory demands of modern commercial fleet management, as regulatory bodies require auditable, deterministic, and physically verifiable baselines, such as the one provided by the ISO framework to legally issue certifications, carbon credits or offsets.

As a foundation, the framework utilizes the following formula to calculate total vehicle fuel consumption[5]:

$$\Phi = \Phi_{\mathrm{v}} + \Phi_{\mathrm{id}} = \eta \cdot b_{\mathrm{e}} \frac{\int_0^{T} (F_A + F_B + F_C + F_D + F_E)\, v\, \Delta t}{\int_0^{T'} v\, \Delta t} \tag{1}$$

Where the standard defines the total energy demand by integrating the complete array of physical opposing forces acting on the vehicle:

F_A, F_B Acceleration and braking forces.

F_C Aerodynamic drag.

F_D Rolling friction.

F_E Slope/gravitational resistance.

Φ_{id} Baseline idling losses during standstill phases.

η and b_{e} Efficiency and engine-specific fuel consumption tuning metrics.

3 Data Methodology

The empirical base of this study is a high-resolution mobility dataset with trips collected from 2022 to 2025, recorded via Global Navigation Satellite System (GNSS) receivers on nomadic devices. The data covers European urban centers as well as routes in Aruba, covering a variety of road types, ranging from local urban streets to high-speed segments of the German autobahn. Out of the 19 total recorded trips, 11 were conducted in Aruba. This temporal and spatial window captures the granular, second-by-second driving behaviors, specifically focusing on rapid acceleration, deceleration, and idling phases, which are essential for modeling fuel consumption and carbon emissions in congested urban zones.

The data in this study was largely collected from a VW Passat Variant TSI, and while that is not true for all the data, the data was treated as if it was all collected from the same vehicle, since the other vehicles also fell in the N1 Category. This is further explored in the discussion section of this paper.

The table in Appendix A goes in-depth about the data structure, but the relevant aspects are that the dataset provides a comprehensive physical profile of each trip, including spatiotemporal variables such as speed, acceleration, altitude, kinematic ones such as speed, acceleration, and standstill time. Using these as a foundation, attributes like work, emissions, and efficiency are calculated transiently by an online server adhering to the ISO 23795 standard[5]. So while the foundation of this model is performed by empirical high-accuracy metrics, the actual work calculations, fuel consumption, and carbon emissions are Newtonian physics-based estimates rather than measurements.

The dataset was found to be highly reliable, accurate and consistent, containing no null or missing values. Minimal modifications or filters were applied to the dataset before modeling, all of which are properly outlined in this paper, this ensures that functionality of the ISO Standard, though in the methodology some modifications were made to better match the research goals, these will be adequately outlined there.

The only inconsistencies are with the extremely timely routes, that represent long stops, rather than long trips, this may be because of a long traffic jam, but more likely it is caused by long parking events where the device was left on. Though due to the limited amount of data, these will be treated as normal, fair trips.

More information about the trips can be found in Appendix B.

3.1 WLTP Profiles Standardization

To ensure standardized and universally comparable fuel consumption and CO_2 emission estimates, route segments within the dataset were categorized according to the Worldwide Harmonized Light Vehicles Test Procedure (WLTP) speed classifications: Low, Medium, High, and Extra-High, outlined in the table below. This categorization is essential as ISO 23795 utilizes the WLTP driving cycle as its foundational mathematical baseline[5].

Table 1: Operational Speed Profiles and Vehicular Efficiency Dynamics

Speed Profile	Speed Range	Operational Efficiency Explanation
Low	< 30 km/h	Reflects urban "Stop-and-Go" traffic where thermal efficiency is low and idling losses are high.
Medium	30–50 km/h	Typically the "efficiency sweet spot" where engine load is optimal.
High	50–70 km/h	Increasing aerodynamic drag begins to impact fuel economy.
Extra-High	> 70 km/h	Aerodynamic drag becomes the dominant force, significantly increasing CO_2 output.

A challenge arose during the processing of the Aruba data, which, unlike the European dataset, lacked a distinction between high and low-emission trips. While the European trips were pre-partitioned into these categories, the Aruba data required an alternative approach. To address this discrepancy, a Random Forest (RF) Classifier was implemented to impute missing classifications by leveraging the localized vehicle kinematics and instantaneous physical work metrics.

3.2 Random Forest Modeling

The selection of a Random Forest model was guided by its robustness and its proficiency in managing non-linear physical relationships. Given that the mathematical interactions representing inertial power demands and aerodynamic drag are inherently non-linear, the RF model effectively captures these multi-dimensional dynamics without the need for extensive feature transformation or scaling. The model also randomly samples feature subsets at each node split, effectively mitigating the issues associated with the high correlation between variables.

The Random Forest classifier was developed within a Python environment, utilizing the Scikit-learn and Pandas libraries. To assess the model's accuracy, the 5,168 labeled observations from eight distinct trips of the European dataset were divided into training and testing sets using an 80/20 stratified split. The model itself consisted of an ensemble of 100 decision trees, with Gini impurity (a metric used to measure the diversity in a dataset) serving as the primary metric for optimizing node splits, The dataset's "impurity" or "heterogeneity" is assessed using Gini impurity, a standard machine learning metric. In a specific group, it calculates the likelihood that an element would be categorized incorrectly if it were assigned a label at random based on the existing class distribution. Once the validation phase was successfully completed, the final iteration of the classifier was retrained on the entire labeled dataset; this model was then utilized to predict the hidden emission group classifications for the 4,570 unclassified data points in the Caribbean data.

3.3 Simulation Framework

This project uses a simulation designed to emulate how driving behaviors influence both total travel duration and CO2 emissions for specific real-world journeys. The simulation was done with Python, including the pandas, numpy, and networkx libraries.

3.3.1 Initial Vehicle and Road Set Up

The simulation framework integrates several foundational components to ensure physical accuracy. The vehicle is assumed to have a weight, or curb weight, of 1,500 kg and a maximum storage, or payload, capacity of 500 kg. This 1,500 kg baseline was established because a modern Passat typically weighs between 1,450 kg and 1,600 kg. The 500 kg load limit reflects a standard benchmark for this vehicle class.

To determine the precise distance between geographic coordinates, the simulation employs the Haversine formula.

$$d = 2R \cdot \arcsin\left(\sqrt{\sin^2\left(\frac{\Delta\phi}{2}\right) + \cos(\phi_1) \cdot \cos(\phi_2) \cdot \sin^2\left(\frac{\Delta\lambda}{2}\right)}\right) \tag{2}$$

where:

d	Great-circle distance between the two coordinate points (in meters).
R	Mean radius of the Earth (typically modeled as $6,371 \times 10^3$ meters).
ϕ_1, ϕ_2	Latitudes of point 1 and point 2, respectively (converted to radians).
λ_1, λ_2	Longitudes of point 1 and point 2, respectively (converted to radians).
$\Delta\phi$	Difference between the two latitudes, where $\Delta\phi = \phi_2 - \phi_1$.
$\Delta\lambda$	Difference between the two longitudes, where $\Delta\lambda = \lambda_2 - \lambda_1$.

Additionally, a road grade calculation was implemented to mitigate high-frequency noise and altitude errors inherent in GPS data. This process involves applying a five-point centered rolling average to the altitude variable—utilizing the two preceding and two succeeding data points—to estimate the vehicle's true elevation at any given moment. For any data point i, the refined altitude $h_{\text{smoothed},i}$ is calculated as:

$$h_{\text{smoothed},i} = \frac{1}{5}\sum_{j=-2}^{2} h_{i+j} = \frac{h_{i-2} + h_{i-1} + h_i + h_{i+1} + h_{i+2}}{5} \tag{3}$$

where h_i represents the raw GPS altitude recorded at index i. Using this refined altitude data, the road grade (G_i) is computed as the change in elevation over the segment distance (Δd_i) between index i and $i+1$:

$$G_i = \frac{\Delta h_i}{\Delta d_i} = \frac{h_{\text{smoothed},i+1} - h_{\text{smoothed},i}}{\Delta d_i} \tag{4}$$

To ensure stability against extreme outliers and unexpected sensor artifacts, the final calculated grade value is clipped to a predefined operational range:

$$G_{i,\text{final}} = \max\left(-0.10, \min\left(0.10, G_i\right)\right) \tag{5}$$

Below is an example of this rolling average:

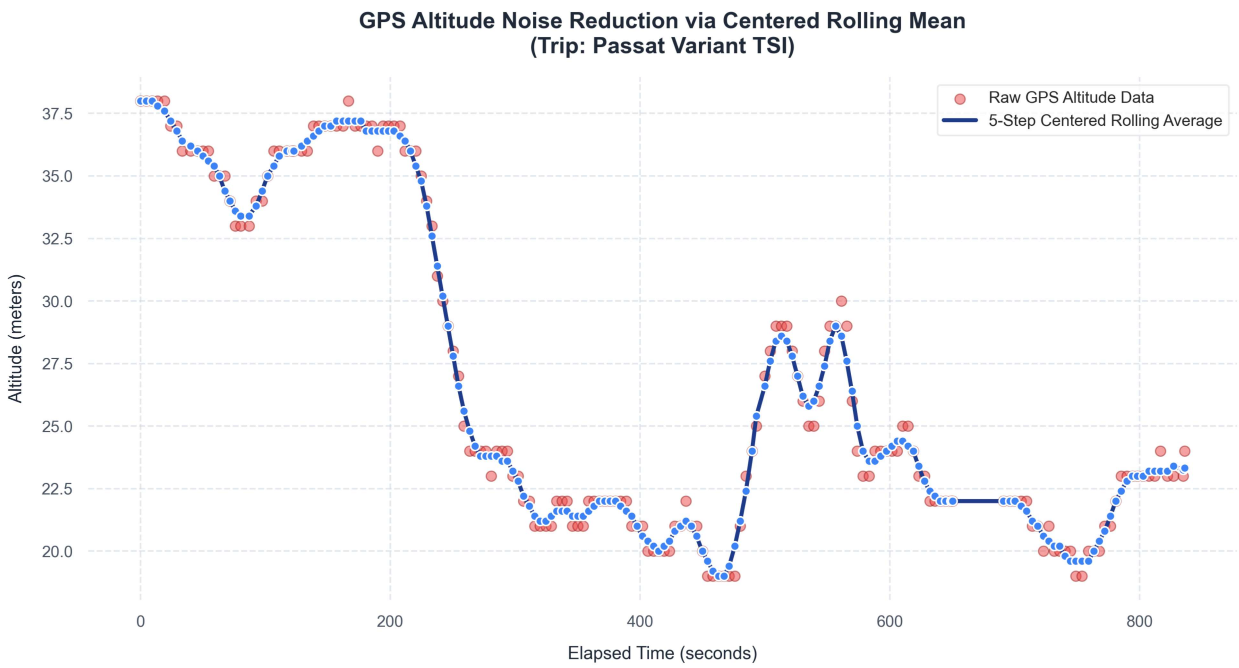

Figure 1: GPS Altitude Noise Reduction via 5-Step Centered Rolling Mean for Trip 7

3.3.2 Driving Profile Generation

The 100 used profiles range from 0 to 1, with a greater α representing a more aggressive driver. In a business case, this might be a driver that has a short delivery window, or that is late on a delivery, and therefore consumes more fuel by accelerating faster than they otherwise would. Profiles with a higher α also tend to reach greater speeds, and therefore tend to consume more fuel. On the other hand, profiles with a lower α tend to represent eco-drivers, that is, those who attempt to consume less fuel by driving slower, and accelerating calmly.

In practice this change was made by scaling the recorded speeds from the data using the following formula:

$$v_{\text{base},i} = v_{\text{recorded},i} \times (0.85 + 0.35\alpha) \tag{6}$$

where $v_{\text{recorded},i}$ is the original vehicle speed recorded in the GPS dataset at index i, and $\alpha \in [0, 1]$ represents the driver aggressiveness tuning parameter.

Meaning that at $\alpha = 0.0$, profile 1 scales down the recorded speeds from the data to 85% of their value, while $\alpha = 1.0$, profile 100 scales the original speeds up to 120% its original value.

Some of the discrepancies of the profiles are also found in their stopping behaviors. Stop Events were implemented in any place the driver of the original trip may have stopped, mimicking traffic lights, pedestrian crossings and traffic. Though the differences in results are caused by differences in the acceleration and deceleration behaviors, where profiles with greater α break harder and accelerate faster. While stops are considered, stop durations are treated as environmental constraints that are pace-independent, and therefore are ignored for the purposes of comparison, since they may differ, this is explored further in the discussion section of this paper.

3.4 The Physics Engine

The computational engine utilized in this study calculates emissions through a segment-based approach consisting of four layers. This core framework is grounded in Power Demand (PoD) tuning constants derived from ISO 23795 standards and Newtonian vehicle mechanics. The primary components of the calculation include the Baseline WLTP Lookup Rate, the Bidirectional Acceleration Penalty, the Grade Penalty, and the Idling Penalty; as shown in the figure below.

The Baseline WLTP Lookup Rate checks the profile's current speed and references the previously discussed WLTP profiles to establish a baseline emission rate in CO_2 kg/km. This lookup step anchors the simulation's foundational energy estimates in real-world regulatory benchmarks.

The Bidirectional Acceleration Penalty is a proxy for inertia force based on the ISO 23795's and Newtonian mechanics, calculated based on the changes in speed over time

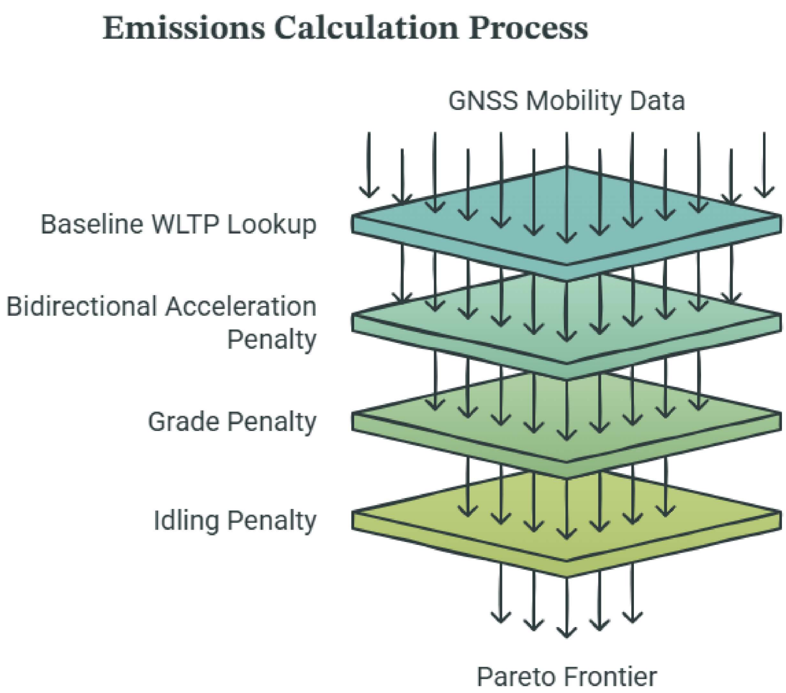

Figure 2: The 4 Layers of the Physics Engine

which determine acceleration as seen in the formula below. The model tracks the absolute acceleration value. If it breaks past the threshold of 0.5 m/s^2 , an incremental multiplier penalty is triggered. Due to the nature of absolute values, the penalty is bidirectional, meaning it penalizes both heavy, fuel-wasting acceleration spikes and hard, late braking events. These late breaking events waste kinetic energy, because rather than braking late, the driver could have allowed the car's momentum to propel it forwards while bringing the vehicle to a stop due to other resistances.

To be clear, this bidirectional penalty, when it comes to breaks, is an operational proxy used to account for the inevitable subsequent acceleration required after dropping below optimal speed, it is not necessarily supposed to represent a fuel cost directly caused by the breaking action

$$\text{Penalty} = 1.0 + (|\text{acceleration}| - 0.5) \times \text{acceleration\_penalty} \times \text{mass\_multiplier} \quad (7)$$

where:

Penalty	The final calculated emissions penalty multiplier coefficient.
|accel|	The absolute magnitude of instantaneous acceleration or deceleration calculated as acceleration = $\Delta v/\Delta t$
0.5	The baseline physical acceleration tolerance threshold (m/s^2).
acceleration_penalty	An ISO 23795 tuning constant scaling the severity of excess acceleration.
mass_multiplier	A dynamic mass multiplier that scales the penalty based on payload weight, ranging from 0 to 1 representing the percentage of the vehicle's carrying capacity that is taken

The Grade Penalty, or slope resistance, aims to reflect the gravitational resistance when moving vertically as seen in the formula below. It is a multiplier that reflects moving up an incline and increases the fuel consumption significantly (amplified further if the vehicle's mass is high due to cargo weight). Downhill segments reduce emissions, though it is safely capped at a minimum value (0.5) so the car never magically generates fuel while descending.

$$\text{Grade Factor} = 1.0 + (\text{grade} \times \text{grade\_weight} \times \text{mass\_multiplier}) \tag{8}$$

where:

Grade Factor	The final calculated slope resistance multiplier coefficient.
grade	The localized road incline or decline, calculated as ΔElevation/Distance.
grade_weight	A constant scaling factor determining how severely a given road grade impacts fuel consumption.
mass_multiplier	A dynamic mass multiplier that scales gravitational resistance proportionally with vehicle cargo weight.

4 Results

By simulating the 100 driving profiles, across the 19 routes, unladen of any cargo, the Pareto-frontiers were mapped and the continuous multi-objective trade-off space between operational transit time and environmental impact was found. The morphology of the frontier curves remains highly consistent across both short-range urban segments as shown in trips 1 and 2 and extended commuter links like trips 9 and 10 (see Appendix C). The trade-off curve shows three distinct behavioral domains: the eco-friendly bound, the balanced frontier knee and the speed bound.

The eco-friendly bound ($\alpha \to 0.0$) focuses on minimizing the emissions, and for this purpose the vehicle speeds are limited to 85% of the baseline. And while this minimizes the impact of forces such as aerodynamic drag and minimizes bidirectional penalties, the vehicle, in most business cases, would just not be moving fast enough.

The balanced frontier knee ($\alpha \in [0.4, 0.6]$) represents the mathematical inflection point of the frontier. Here, minor changes in transit speed result in disproportionately large reductions in fuel consumption and CO_2 output. It provides an operational happy medium where commercial dispatchers can meet moderate Service Level Agreements (SLAs) while cutting emissions.

The speed bound ($\alpha \to 1.0$), scales speed to 120% of the baseline. The curve slopes upward asymptotically; attempting to shave minor seconds off the arrival time triggers massive spikes in CO_2 emissions due to heavy, energy-wasting acceleration cycles and hard braking events.

Forcing the unladen vehicle into a pure eco-driving behavior floor ($\alpha = 0.0$) drops baseline CO_2 mass emissions by an average of 7.2% to 7.4% across the network. This empirical baseline clearly quantifies the motivation behind this paper, outlining that minor emission reductions demand massive time concessions.

The effects of eco driving are most notable in smaller, low-speed, urban routes, with carbon savings of about 7.4%. That is because these routes tend to have stop and go events make up a large portion of their trip, and a lower α will tend to break softer, keep their momentum for as long as possible, and accelerate unaggressively. This paired with the speed suppression, that keeps the vehicle in the highly efficient WLTP "Low" classification,

Table 2: Operational Performance Across Extreme Driving Bounds

Metric Group	**Eco-Optimization Limit** ($\alpha = 0.0$)	**Aggressive Transit Limit** ($\alpha = 1.0$)	**Net Operational Variance**
Velocity Scaling	85% of baseline speed	120% of baseline speed	Δ 35% nominal shift
Network Travel Time	Maximum Duration (+25% to 30% average penalty)	Minimum Duration (Base SLA baseline)	Significant time penalty for green routing
CO_2 Mass Emissions	Absolute Floor (Minimum fuel burn)	Exponential Ceiling (Peak power demand)	~3.3% to 9.3% efficiency gap[10]

increases carbon-savings.

Due to the methodology's design, the travel time tends towards linearity, but the carbon emissions are fundamentally non-linear. Reducing travel time triggers disproportionately higher transient power demands due to aerodynamic drag and inertial resistance when moving towards $\alpha \to 1.0$.

4.1 The Effect of Cargo

Under unladen conditions, shifting from an aggressive transit profile ($\alpha = 1.0$) to a pure eco-profile ($\alpha = 0.0$) results in an average network-wide carbon reduction of 7.74%. However, when the vehicle operates at maximum payload capacity, this exact same behavioral shift achieves an average carbon reduction of 9.36%. This proves that eco-driving becomes mathematically more effective as vehicle weight increases. Eco-driving causes a higher percentage return on efficiency because it avoids activating the heavily penalized mass multipliers in the physics engine, that is, it costs less fuel to carry less mass, so if there is a greater mass, then eco-driving becomes more relevant.

Trip 10 displays the most dramatic vulnerability to payload compounding. For an aggressive profile, adding the cargo payload causes emissions to increase by 6.62%, but under eco-driving, this inflation is only 2.51%. Since high-density urban corridors are defined by continuous stop-and-go kinematics, the vehicle is forced to repeatedly overcome static inertia from a complete standstill. An aggressive driver carrying cargo accelerates and breaks harshly, but an eco-driver, on the other hand, protects momentum, resulting in an exceptional 16.51% carbon savings when fully laden.

5 Discussion

The primary objective of this thesis was to map the operational trade-off between trip duration and environmental impact within time-sensitive B2B logistics. By constructing a multi-objective Pareto optimization framework underpinned by the ISO 23795 standard, this study shows how driver behavior (α) and vehicle payload interact to reshape the Pareto frontier.

5.1 Interpreting the Results

At the lower bound of the frontier, vehicle speeds are structurally suppressed to 85% of the empirical baseline. This domain successfully minimizes transient power demands by mitigating aerodynamic drag and smoothing acceleration cycles. For unladen vehicles, operating at this floor yields an average CO_2 reduction of 7.2% to 7.4% across the network.

The efficacy of this bound is heavily localized; low-speed urban corridors (such as Trips 1 and 2 in Musberg) experience the most pronounced benefits (see Appendix C.1 and C.2). This occurs because the suppressed velocity keeps the vehicle operationally locked within the highly efficient WLTP "Low" classification (¡30 km/h), minimizing the catastrophic thermal efficiency losses and high idling overhead typical of urban "stop-and-go" cycles. However, the accompanying 25% to 30% increase in network travel time renders this extreme bound nonviable for traditional time-sensitive B2B operations.

Another finding comes in the form of reaffirming what had already been shown. Late and harsh braking acts as a severe kinetic energy sink; aggressive drivers frictionally dissipate momentum that an eco-driver would have leveraged to overcome rolling resistance and aerodynamic drag during coasting phases. As outlined in the result section of this paper, B2B logistics operators should aim for the frontier knee, that is, $\alpha \in [0.4, 0.6]$. Within this operational sweet spot, minor concessions in transit speed yield disproportionately large dividends in carbon mitigation. It allows fleet dispatchers to protect strict SLAs while simultaneously capturing the bulk of achievable emissions reductions.

A critical contribution of this study is the result that eco-driving efficiency scales non-linearly with vehicle mass. That is, as mass increases, aggressive driving becomes more expensive, or even wasteful. Under unladen conditions, moving from an aggressive profile ($\alpha = 1.0$) to a pure eco-profile ($\alpha = 0.0$) nets a 7.74% carbon reduction. When the vehicle is simulated at its maximum payload capacity of 500 kg, this exact behavioral shift achieves a network-wide average carbon reduction of 9.36%.

5.2 Comparison with Existing Literature

The anticorrelated relation between speed and emission reductions that pure eco-routing consistently penalizes travel time by 20% to 30%[4] is explicitly mirrored in this study's eco-friendly bound. Furthermore, our findings align with[8], confirming that when hard time constraints are introduced, the emission reduction potential of time-based logistics drops drastically.

However, this thesis may advance existing knowledge in two distinct ways:

1. Rejection of Macroscopic Averages: Unlike traditional frameworks criticized for relying on static trip speed averages[9] , this paper's utilization of second-by-second GNSS data captures the transient emission spikes caused by localized micro-kinematics, such as highway ramps and intersections.

2. White-Box Models: While recent computer science literature leans heavily toward machine learning models for emission estimation[7], such models fail in comparison to the deterministic, white-box methodology of the ISO 23795 standard. As this framework ensures that carbon savings are fully auditable, verifiable, and legally compliant with global climate policies like the EU's ETS-2 framework.

5.3 Limitations and Future Work

While great effort was made to make a robust and comprehensive study, the model is imperfect, and future work should take this into account when researching eco-driving and eco-routing.

First, the constraints are a problem with the data, and while the data was of great quality, 19 trips are not nearly enough to fine tune the model. Future work should keep collecting and compounding trips. Another data-related issue was the lack of real-traffic, which the model adapted to by assuming the traffic was represented by stopping events. While this solution worked for this thesis, future work should aim to stimulate live traffic and driving decisions. Another issue with the model is that no reliable speed-limit information was available for all the roads, so at some points drivers, especially the ones tending towards a greater α nay have surpassed speed limits, which is, of course, unadvised. Further shortcomings are the lack of inclusion of traffic light timings, which aggressive profiles may accelerate to match; the lack of temperature and weather conditions, which may affect maximum speeds, breaking behaviors and friction and resistance physics.

Another small issue with the data, and subsequently the model is that it treats all trips as if they were driven by a Passat Variant TSI, which is true for many of the trips, but not all. A related issue is that little is known about the vehicle's cargo, though the assumption is that the cargo is always negligible, therefore the data may be right-skewed and the simulated vehicle may tend towards being slower than it actually is. On top of that, most business fleets do not use a passenger vehicle, rather they rely on vans and trucks, but there was no such data available. Furthermore, it is unclear if different fuel types, such as gas, diesel, electricity; and vehicle types, such as traditional, hybrid, or electric; affect the simulation.

Additionally, the GNSS receiver is subject to movement while inside the vehicle, so additional elevation data should be used in conjunction with the collected data. Also some assumptions were made, these include the assumption that there are no additional energy costs, such as AC or lighting; or perishables in the cargo.

5.4 Further Considerations

Weights are likely dynamic in a real B2B situation, that is, the vehicle's total mass may vary throughout its course, based on what is delivered or picked up. This should be accounted for in future work.

6 Conclusion

This thesis has successfully helped map the complex, multi-objective trade-off space between operational transit time and environmental sustainability within time-sensitive B2B logistics. By constructing a multi-objective Pareto optimization framework grounded in the deterministic, white-box ISO 23795 standard, this research quantified how driver behavioral dynamics and vehicle payload interact to alter the morphology of the Pareto frontier.

Furthermore, this study mathematically verified that the efficiency of eco-driving scales non-linearly with vehicle mass. Shifting from an aggressive transit profile to a pure eco-profile nets a 7.74% carbon reduction under unladen conditions, but achieves a 9.36% network-wide reduction when the vehicle is simulated at its maximum payload capacity. This payload compounding effect is most pronounced in high-density urban corridors defined by continuous stop-and-go kinematics, where an eco-driver protects vehicle momentum and mitigates heavy inertial penalties, achieving up to 16.51% carbon savings when fully laden.

As green logistics transitions from a voluntary corporate social responsibility goal to a strict regulatory mandate, data-driven frameworks like the one developed here become essential. Utilizing an auditable, deterministic baseline compliant with global climate policies like the ISO 23795 standard ensures that emissions data is credible.

6.1 Recommendations

For any practitioners or policymakers, the recommendations will be made clear in this section. First, fleets should target the frontier knee, the metaphorical sweet spot, that dictates the points at which the engine is at its most efficient, and the carbon emissions are still manageable. Implement the frontier as a part of fleet management, to reduce the total emissions, and to disperse the time losses with the reduced fuel consumptions.

Furthermore, focus on eco-driving in congested roads and dense urban areas, or when the cargo is at its heaviest, as these situations are when the effects of this practice are felt the most. And finally, policymakers should incentivize the adoption of deterministic, white-box auditing tools based on international standards like ISO 23795.

References

[1] Kyoungho Ahn and Hesham A. Rakha. "Network-wide impacts of eco-routing strategies: A large-scale case study". In: *Transportation Research Part D: Transport and Environment* 25 (2013), pp. 119–130. ISSN: 1361-9209. DOI: 10.1016/j.trd.2013.09.006. URL: https://www.sciencedirect.com/science/article/pii/S1361920913001259.

[2] Jinghui Wang, Ahmed Elbery, and Hesham A. Rakha. "A real-time vehicle-specific eco-routing model for on-board navigation applications capturing transient vehicle behavior". In: *Transportation Research Part C: Emerging Technologies* 104 (2019), pp. 1–21. ISSN: 0968-090X. DOI: 10.1016/j.trc.2019.04.017. URL: https://www.sciencedirect.com/science/article/pii/S0968090X18316437.

[3] Weiliang Zeng, Tomio Miwa, and Takayuki Morikawa. "Prediction of vehicle CO2 emission and its application to eco-routing navigation". In: *Transportation Research Part C: Emerging Technologies* 68 (2016), pp. 194–214. ISSN: 0968-090X. DOI: 10.1016/j.trc.2016.04.007. URL: https://www.sciencedirect.com/science/article/pii/S1361920916306009.

[4] Min Zhou, Hui Jin, and Wenshuo Wang. "A review of vehicle fuel consumption models to evaluate eco-driving and eco-routing". In: *Transportation Research Part D: Transport and Environment* 49 (2016), pp. 203–218. ISSN: 1361-9209. DOI: 10.1016/j.trd.2016.09.008. URL: https://www.sciencedirect.com/science/article/pii/S1361920916306009.

[5] International Organization for Standardization. *Intelligent transport systems — Extracting trip data using nomadic and mobile devices for estimating CO2 emissions — Part 1: Fuel consumption determination for fleet management.* Standard ISO 23795-1:2022. Geneva, Switzerland: International Organization for Standardization, 2022. URL: https://www.iso.org/standard/76971.html.

[6] Ahmed Fahmin et al. "Eco-Friendly Route Planning Algorithms: Taxonomies, Literature Review and Future Directions". In: *ACM Comput. Surv.* 57.1 (2024). ISSN: 0360-0300. DOI: 10.1145/3691624. URL: https://doi.org/10.1145/3691624.

[7] Xianan Huang and Huei Peng. *Eco-Routing based on a Data Driven Fuel Consumption Model*. 2018. arXiv: 1801.08602 [stat.AP]. URL: https://arxiv.org/abs/1801.08602.

[8] Anmol Pahwa and Miguel Jaller. "Evaluating private and system-wide impacts of freight eco-routing". In: *Transportation Research Part D: Transport and Environment* 130 (2024), p. 104170. ISSN: 1361-9209. DOI: https://doi.org/10.1016/j.trd.2024.104170. URL: https://www.sciencedirect.com/science/article/pii/S1361920924001275.

[9] Ana Aguiar et al. "MobiWise: Eco-routing decision support leveraging the Internet of Things". In: *Sustainable Cities and Society* 87 (2022), p. 104180. ISSN: 2210-6707. DOI: 10.1016/j.scs.2022.104180. URL: https://www.sciencedirect.com/science/article/pii/S2210670722004930.

[10] Hesham Rakha and Kyoungho Ahn. *Developing an eco-routing application*. Technical Report N14-07. Sponsored by the United States. Department of Transportation. Office of Research and Special Programs. Virginia Polytechnic Institute and State University, Jan. 2014. URL: https://rosap.ntl.bts.gov/view/dot/26994.

A Data Structure

Table 4: Summary of Dataset Attributes and Physical Profiles

Variable	Unit of Measure	Explanation
Time	Milliseconds (ms)	The absolute timestamp when the data row was recorded
Latitude	Degrees (°)	The geographic GPS coordinates of the vehicle
Longitude	Degrees (°)	The geographic GPS coordinates of the vehicle
Distance	Meters (m)	The distance traversed during that specific timestamp sample interval
Altitude	Meters (m)	The vehicle's elevation above sea level
Accuracy	Meters (m)	The estimated error margin or precision radius of the GPS signal
Speed	Meters per Second (m/s)	The instantaneous speed of the vehicle
Acceleration	Meters per Second Squared (m/s^2)	The vehicle's instantaneous change in velocity over time
Stand Still Time	Seconds (s)	The cumulative time the vehicle has spent idling or at a complete stop
Acceleration Work	Joules (J)	The energy expended accelerating the vehicle mass
Aerodynamic Drag Work	Joules (J)	The energy expended to push past air resistance

Variable	Unit of Measure	Explanation
Grade Work	Joules (J)	The energy changes due to moving up or down slopes
Roll Work	Joules (J)	The energy lost due to friction between the vehicle tires and pavement
Stand Still Work	Joules (J)	The energy expended while the vehicle is running but stationary
Fuel	Liters (L)	The absolute volume of fuel consumed during that sampling window
CO_2 Emissions	Kilograms (kg)	The Carbon Dioxide emitted during that sampling window
ECE Acceleration	Percentage (%)	Metrics comparing the trip's work profiles relative to ECE Framework
ECE Aerodynamic Drag	Percentage (%)	Metrics comparing the trip's work profiles relative to ECE Framework
ECE Stand Still Time	Percentage (%)	Metrics comparing the trip's work profiles relative to ECE Framework
ECE Work	Percentage (%)	Metrics comparing the trip's work profiles relative to ECE Framework

This table provides a summary of attributes within the dataset, which are a comprehensive physical profile of each trip. Attributes like work, emissions, and efficiency are calculated transiently by an online server based on the mobility data sent from the nomadic devices. The calculations are based on basic physics, and the formula details can be found in the ISO 23795 standard. But these are not real-time measures, rather physics-based estimations.

The Economic Commission for Europe (ECE) driving cycles are fundamentally linked to how international vehicle emissions standards are benchmarked.

B Trip Information

The following table presents a detailed summary of the individual trips recorded within the dataset, including their location, total distance, elapsed time, and total CO_2 emissions.

Table 5: Summary of Individual Trip Characteristics

Trip	**Location**	**Distance (km)**	**Time (s)**	**CO_2 (g)**
Trip 1	Musberg	5.759	736	604.5
Trip 2	Musberg	5.629	712	602.1
Trip 3	Rotenburg	13.35	745	1867.5
Trip 4	Rotenburg	10.759	951	1303.9
Trip 5	Rotenburg	10.184	810	1372.9
Trip 6	Rotenburg	9.58	751	1304.2
Trip 7	Rotenburg	10.978	836	1421.9
Trip 8	Punta Brabo	6.013	3573	1867.7
Trip 9	Bubali	11.458	7881	3379.1
Trip 10	Oranjestad	12.631	6682	4579.6
Trip 11	Schotlandstraat	65.446	61,169	10,393.6
Trip 12	Bubali	5.46	6855	1654.8
Trip 13	Frankfurt	7.197	750	2227.5
Trip 14	Bubali	10.078	5590	2983.5
Trip 15	Schotlandstraat	6.358	58,893	4888.8
Trip 16	Bubali	10.884	9787	3681.5
Trip 17	Oranjestad	6.086	69,703	957.7
Trip 18	Bubali	13.833	6449	3619.6
Trip 19	Schotlandstraat	14.919	4869	2619.3

C Pareto Frontiers

For terseness purposes, some frontiers will be omitted. In this appendix, you may find the Pareto frontiers for trips 1, 9, 10, 11 and 13, accompanied by their laden alternatives.

C.1 Unladen Trips

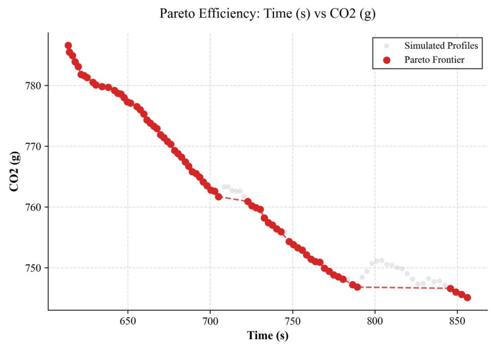

(a) Trip 1

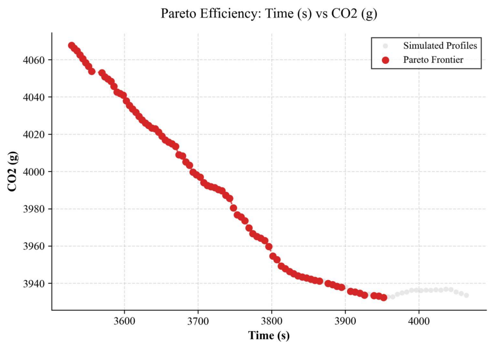

(b) Trip 9

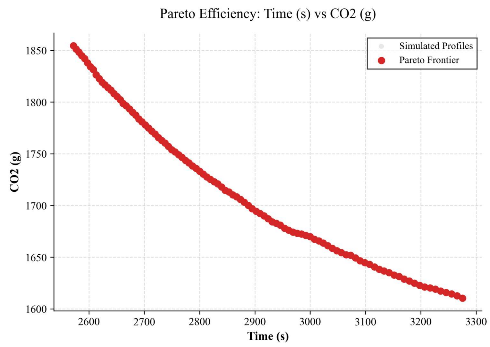

(c) Trip 10

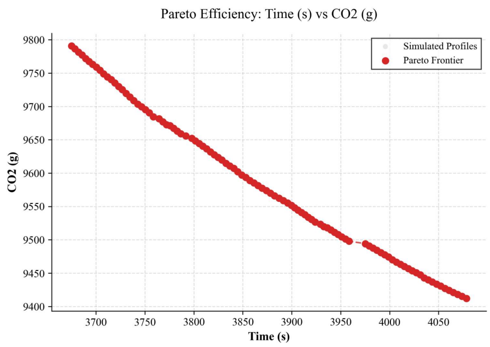

(d) Trip 11

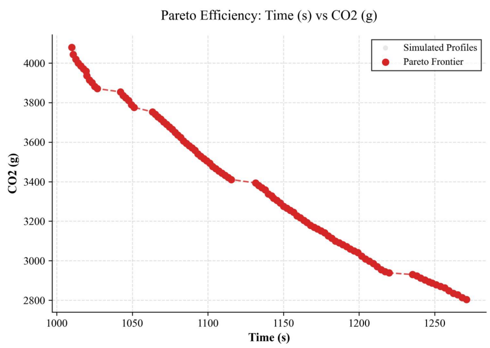

(e) Trip 13

Figure 3: Pareto Efficiency frontiers comparing Time (s) versus CO2 emissions (g) across the unladen trip profiles.

C.2 Laden Trips

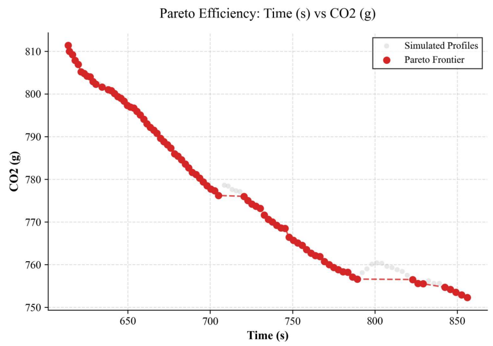

(a) Trip 1

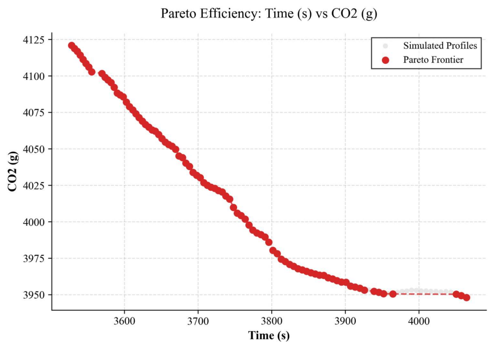

(b) Trip 9

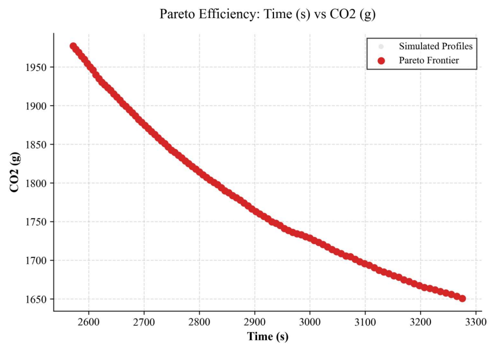

(c) Trip 10

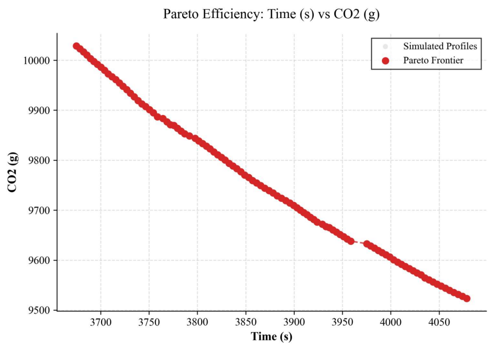

(d) Trip 11

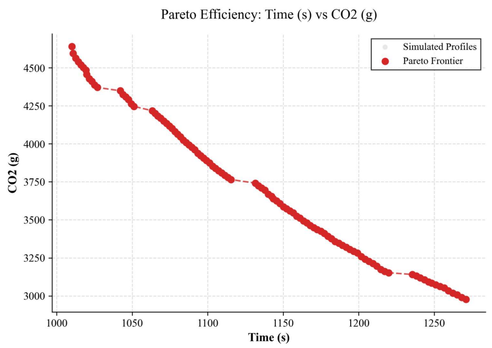

(e) Trip 13

Figure 4: Pareto Efficiency frontiers comparing Time (s) versus CO2 emissions (g) across the laden trip profiles.