Politecnico di Torino
Instituto Superior Técnico

Laurea Magistrale in Engineering and Management

Academic Year 2025/2026

Graduation Session: July 2026

Optimization of Pit Stop Strategies in Formula 1

Supervisors:

Prof. Daniel Rebelo dos Santos

Prof. Rosario Scatamacchia

Candidate:

Carolina Ribeiro Ferreira

July 2026

Declaration

I declare that this document is an original work of my own authorship and that it fulfills all the requirements of the Code of Conduct and Good Practices of the Universidade de Lisboa.

Declaration of the use of Artificial Intelligence tools

I declare that generative Artificial Intelligence (AI) and AI-assisted tools were used in this dissertation. Nevertheless, I take full responsibility for the content and integrity of this work. In compliance with the Guide for the Responsible Use of Artificial Intelligence at Técnico, I declare that the following tools were used:

1) For linguistic revision: Anthropic. 2026. Claude [Large language model]. https://claude.ai/; Google. 2026. Gemini [Large language model]. https://aistudio.google.com/

2) As a coding assistant: GitHub. 2026. GitHub Copilot [AI coding assistant]. https://github.com/features/copilot

Acknowledgments

As I reach the end of this journey and close this chapter of my academic life, I look back on a path that was long and, at times, difficult. Yet, it was ultimately wonderful, mostly because I got to share it with people who gave me every kind of memory that I will carry with me forever.

To my parents, who have always given their best so that I could become who I am today, and who gave me every tool I needed to build my own path. None of this would have been possible without them, so this achievement is also theirs.

To my brother, who I used to fight with as a kid, and who has become one of my closest friends. Thank you for playing with me, for laughing with me, and for always bringing lightness and comfort to our home.

To my boyfriend, thank you for your patience with my energy and my endless imagination, and for being my safe harbor. We're finishing this chapter together, and I couldn't be prouder of the two of us, as we now move on, side by side, into whatever comes next. And thank you for always fixing my technological problems, of which I have many.

To my family, for always giving me a second home and for wanting nothing more than my happiness.

To my godparents, who have always welcomed me with open arms and surrounded me with art and photography, whether on trips to museums or days at the beach. Thank you for enriching my life with those cultural experiences, which sparked a passion in me that only continues to grow.

To my friends, those from childhood, those I met at Técnico, and those I found in Turin, thank you for the memories that make life so much more beautiful and joyful. My life would not be the same without you.

To my professors, and especially my supervisor, thank you for your invaluable guidance, for supporting this project, and for trusting in my ability to reach the finish line.

To God, for always looking over me and guiding me through every step of the way.

Resumo

A estratégia de corrida na Fórmula 1 é um problema complexo de tomada de decisão, que exige o equilíbrio entre o desempenho do carro, a gestão da degradação dos pneus e a incerteza inerente às condições da corrida. Historicamente, as abordagens de Investigação Operacional para este problema têm sido maioritariamente determinísticas, falhando em contabilizar a volatilidade dos parâmetros em tempo real. Esta tese propõe um quadro de otimização estocástica, formulada como um modelo de Programação Linear Inteira Mista de duas fases, para determinar estratégias ótimas de *pit stop* e seleção de compostos de pneus, integrando a incerteza na degradação e na duração das paragens nas boxes, o que mostrou ser benéfico devido ao valor sistematicamente negativo do Valor da Solução Estocástica. O modelo foi calibrado com dados empíricos das épocas de 2022 a 2024 e validado através da comparação com as recomendações da Pirelli e com as estratégias reais dos vencedores das corridas. Adicionalmente, foi desenvolvido um quadro de simulação de Monte Carlo para avaliar a robustez destas estratégias perante interrupções causadas por bandeiras amarelas. Os resultados demonstram que a abordagem estocástica supera consistentemente a determinística, oferecendo estratégias mais robustas e com um melhor desempenho esperado. Esta investigação estabelece um quadro sólido para a tomada de decisão estratégica, ao quantificar o benefício da modelação explícita da incerteza, demonstrando a sua relevância competitiva no ambiente de elevado desempenho da Fórmula 1.

Palavras-chave: Fórmula 1, Otimização de Pit Stops, Programação Linear Inteira Mista, Otimização Estocástica

Abstract

Formula 1 race strategy requires teams to determine pit stop timing and tyre compound sequencing under conditions of significant uncertainty, yet existing Operations Research approaches to this problem remain largely deterministic, treating tyre degradation and pit stop duration as fixed values despite their inherent race-day volatility. This thesis addresses that gap by developing an optimisation framework that progressively incorporates uncertainty into the race strategy problem. First, a deterministic Mixed-Integer Linear Programming model is formulated and applied across all recent Grand Prix circuits of a recent Formula 1 season to determine optimal tyre sequences, stint lengths, and pit stop timings. A sensitivity analysis is utilised to recalibrate degradation parameters so that model outputs align with observed compound usage patterns. The deterministic formulation is then extended into a two-stage stochastic programming model that evaluates strategies across multiple scenarios sampled from empirical degradation and pit stop duration distributions, and its benefit is quantified through a consistently negative Value of the Stochastic Solution. Stochastic strategies are further validated against Pirelli's official pre-race recommendations and the strategies executed by actual 2024 race winners. Finally, a Monte Carlo simulation framework, running over 10000 replications, stress-tests the optimised strategies against yellow flag disruptions, confirming that anticipating a pit stop within a disruption pit window measurably reduces total race time. Together, these results demonstrate that explicitly modelling uncertainty, both in model parameters and in external race disruptions, produces race strategies that are more robust than those derived from deterministic planning alone.

Keywords: Formula 1, Pit Stop optimisation, Mixed-integer Linear Programming, Stochastic Programming, Monte Carlo simulation

Contents

List of Tables

List of Figures

Glossary

AI	Artificial Intelligence
DES	Discrete-Event Simulation
DRS	Drag Reduction System
F1	Formula 1
FIA	Fédération Internationale de l'Automobile
FP	Free Practice
GA	Genetic Algorithm
GP	Grand Prix
H	Hard Tyre
IQR	Interquartile Range
LSTM	Long Short-Term Memory
M	Medium Tyre
MGU-K	Motor Generator Unit – Kinetic
MILP	Mixed-Integer Linear Programming
MIQP	Mixed-Integer Quadratic Programming
OR	Operations Research
P5	5th Percentile
P95	95th Percentile
Q1	First Quartile
Q2	Second Quartile
Q3	Third Quartile
RNN	Recurrent Neural Networks
S	Soft Tyre
SC	Safety Car
Std	Standard Deviation
VSC	Virtual Safety Car
VSE	Virtual Strategy Engineer
VSS	Value of the Stochastic Solution
YF	Yellow Flag

Chapter 1

Introduction

Formula 1 (F1) is the highest class of international single-seater motorsport, sanctioned by the Fédération Internationale de l'Automobile (FIA). Teams design and build their own cars, and compete over a season-long calendar of Grands Prix held on circuits around the world, with points awarded at each race towards two annual championships: one for drivers and one for teams. Beyond the mechanical performance of the car, every race weekend is shaped by a layer of strategic decision-making, where teams must determine how to convert a fixed car and a fixed set of regulations into the fastest possible race time.

Therefore, winning a Grand Prix (GP) depends not only on the performance of the car and the skill of the driver, but on the decisions made by the team throughout the race weekend. Race strategy involves decisions about when to bring the car into the pit lane to change its tyres, and which tyre compound to fit in their place, since tyres degrade over the course of a race and different compounds offer different trade-offs between speed and durability. A pit stop costs valuable time on track, so a team must weigh the performance lost to a worn tyre against the time surrendered by stopping to replace it, all while the outcome of that decision is shaped by factors that cannot be fully known in advance, such as the actions of competing teams or unplanned interruptions to the race.

Consequently, race strategy sits at the intersection of two related but distinct problems: determining the best strategy to plan for under normal conditions, and determining how robust that plan remains once the uncertainty inherent to racing is taken into account. Both dimensions are strategic decisions that must be made by every team, at every race, and both are the object of the present work.

1.1 Motivation

In Formula 1, races are often decided by margins of tenths of a second, and a pit stop executed one lap too early or too late can cost track position that is rarely recovered on track, with direct consequences for both championships and, ultimately, for a team's financial standing.

From the perspective of Industrial Engineering and Management, a pit stop is a masterclass case study of process optimisation under constraints. It represents a high-speed industrial changeover, where a team of around twenty specialised technicians executes a synchronised sequence of tasks [1], on average between two and 2.5 seconds to change the four tyres, while the record stands at 1.82 seconds, set by McLaren at the 2023 Qatar Grand Prix. For an industrial engineer, this environment offers a unique opportunity to apply Operations Research (OR), incorporating stochasticity and uncertainty, and to formulate it as a mathematical optimisation problem.

However, existing Operations Research approaches to this problem remain largely deterministic, treating degradation rates and pit stop durations as fixed values despite their inherent race-day volatility. This gap between a deterministic planning tool and the genuinely uncertain environment it is meant to serve, as the core motivation for this thesis, which extends a deterministic tyre strategy optimisation framework into a two-stage stochastic formulation, and further stress tests the resulting strategies against yellow flag (YF) disruptions through Monte Carlo simulation.

1.2 Objectives and Deliverables

The main objective of this thesis is to develop an optimisation framework capable of determining the optimal pit stop timing, tyre compound sequence, and number of stops for a Formula 1 Grand Prix, first under a deterministic formulation and subsequently under a stochastic one that accounts for the inherent variability of tyre degradation and pit stop execution. The three main objectives are:

(i) Formulate and solve a deterministic Mixed-Integer Linear Programming (MILP) model to determine the optimal race strategy across the Formula 1 calendar.

(ii) Extend the deterministic formulation into a two-stage stochastic programming model, and quantify the benefit of explicitly modelling parameter uncertainty through the Value of the Stochastic Solution.

(iii) Assess the robustness of the optimised strategies under yellow flags periods, through a Monte Carlo simulation framework.

The main deliverables of this thesis are the deterministic and stochastic optimisation models developed to solve the pit stop strategy problem, the calibrated parameter set derived from the sensitivity analysis and benchmarked against empirical Formula 1 data, the validation of the results against real strategies, and the Monte Carlo simulation framework utilised to evaluate strategy robustness under race-day uncertainty.

1.3 Thesis Outline

The remainder of this work is organised as follows. Chapter 2 provides the contextualisation needed to understand Formula 1's regulatory framework, tyre compounds, race strategy typology, and sources of operational uncertainty. Chapter 3 presents the literature review, organising prior work across four methodological streams and identifying the research gap this thesis addresses. Chapter 4 introduces the optimisation model, detailing the two-stage stochastic MILP formulation developed to determine tyre sequences and pit stop timing. Chapter 5 describes the data utilised, including its sources and the derivation and calibration of the model's parameters. Chapter 6 presents the results analysis, the main focus of this thesis, covering the deterministic and stochastic model results, their validation against Pirelli's recommendations and actual race-winner strategies, and the Monte Carlo simulation of strategy robustness under yellow flag. Finally, Chapter 7 draws the work together, setting out its achievements, and the directions it opens for future work.

Chapter 2

Contextualisation

Formula 1 is the world's leading motorsport, combining cutting-edge engineering with high-stakes competitive strategy at a global scale. The sport is governed by the Fédération Internationale de l'Automobile (FIA), which establishes the technical and sporting regulations that all teams must follow. Every Formula 1 Grand Prix (GP) weekend follows a structured sequence of events, each with a distinct strategic purpose, culminating in the Sunday race.

The weekend begins with three Free Practice (FP) sessions. Both FP1 and FP2 take place on Friday, each lasting approximately sixty minutes, while FP3 occurs on a Saturday morning for the same duration. During FP1, drivers make their first contact with the circuit, establishing a baseline setup and collecting initial data on car balance and track conditions. FIA regulations also mandate that each team fields a rookie driver in at least two Free Practice sessions per season, typically during FP1. FP2 shifts focus towards longer race simulation runs, where engineers gather critical data on tyre degradation per compound, fuel-load behaviour, and overall pace across a full stint. FP3 is reserved for final setup refinements and short qualifying simulations, allowing drivers to prepare for the session that immediately follows [2]. Across all three sessions, the data collected feeds directly into the strategy models that teams build before the race, making Free Practices sessions as much about data as they are about driving.

Qualifying takes place two hours after FP3, on Saturday, and determines the starting grid for the Sunday race, in a three-segment format. The first one, Qualifying 1, runs for eighteen minutes and includes all drivers, after which the five slowest (or six in 2026 and onwards) are eliminated and assigned grid positions for the last spots. Secondly, Qualifying 2 runs for fifteen minutes, and the same rule for elimination in Qualifying 1 is applied. Finally, Qualifying 3 decides the positions among the ten remaining drivers, with the fastest time earning pole position. The grid position carries significant strategic weight, as starting at the front grants access to clean air and greater flexibility in choosing the pit window, while starting further back may force a more reactive or aggressive strategy to recover positions [3].

On Sunday, the race begins with a formation lap in which all cars leave the grid, complete one lap of the circuit, and return to their starting position, allowing drivers to bring tyres and brakes up to operating temperature before the start. The race distance is defined as the number of laps required to complete a distance in the range of 305 to 313 kilometers, subject to a maximum duration of two hours, with the exception of the Monaco Grand Prix, where the distance is only 260 kilometres due to the difficulty of overtaking on the street circuit [4].

Some weekends also have a Sprint session, occurring at approximately six circuits per season, where FP2 is replaced by a Sprint Qualification session and a shorter Sprint Race of roughly 100 kilometers, and takes place on the Saturday. Every time there is a Sprint weekend, it reduces the available practice data and alters tyre allocation rules, which adds an additional layer of complexity to strategy planning.

The results of each Grand Prix feed into two parallel championship standings across the same season: the World Driver's Championship, awarded to the driver who accumulates the most points across all races, and the Constructors' Championship, awarded to the team whose two drivers together accumulate the highest combined total points. This dual structure means that at every race, all the results carry implications at both the individual and team level, and strategic decisions on track are often made with both championships in mind simultaneously. Therefore, the current grid is often composed of ten teams, each fielding two drivers, for a total of twenty cars competing in each Grand Prix. From the 2026 season and onwards this configuration will expanded to eleven teams.

Points are awarded to the top ten classified finishers at the end of each race according to the following scale: 25 points for first place, 18 for second, 15 for third, 12 for fourth, and then ten, eight, six, four, two and one for the remaining positions down to tenth. Additionally, since the 2019 season, one bonus point is awarded to the driver who records the fastest lap during the race, as long as that driver is classified within the top ten finishers.

The season itself spans approximately 24 races distributed across the world, and runs from March to December, where each circuit has its own distinct layout, surface characteristics, and climate conditions, which are precisely what makes circuit-specific strategy optimisation both necessary and central to competitive performance.

2.1 Tyre Compounds and Allocation

The strategies made by each F1 team for each race depend on several factors, such as choosing the optimal tyre compound, which will depend on the track surface and also on the weather. Therefore, Pirelli provides a range of slick tyres, which are tyres without a tread pattern, for dry-weather from C1 to C5, and the newly introduced C6 for high-grip street circuits from 2025 [5]. However, Pirelli does not make all slick compounds available at every race. Instead, for each Grand Prix, Pirelli selects three compounds from the full range, based on factors such as track surface abrasiveness, circuit layout, expected energy loads, and historical wear data. These three selected compounds are always labelled and colour-coded as hard, medium and soft, regardless of their C-grade, hence the same label can correspond to a different compound depending on the circuit. For example, at a high-energy circuit like Silverstone (British GP), C1 is nominated as hard, C2 as medium, and C3 as soft, whereas at a smooth street circuit like Monaco, the C3 becomes the hard, C4 the medium and C5 the soft.

The hard compound, always identified by a white marking, corresponds to the lowest C-grades selected for a given weekend. These compounds are made for high-energy circuits with abrasive asphalt and high-speed corners, such as Silverstone or Spa (Belgian GP) circuit. They provide the highest durability among the three nominated compounds but require more time to reach their optimal operating temperature, denominated as the "warm-up" phase. Secondly, the medium compound, marked in yellow, is the most frequently nominated selected in every Grand Prix, because they offer the best trade-off between the longevity of the harder tyres and the peak performance of the softer ones, making it the most versatile strategic option. Then, the soft compound, represented by the red colour, corresponds to the highest C-grades, designed for street circuits such as Monaco or Baku, because they provide maximum mechanical grip on smooth asphalt and low energy loads, providing faster lap times, and therefore more advantageous for qualifying sessions [6]. However, their wear rate is higher so their durability is significantly lower, which constrains their use during the race.

Each driver receives a fixed tyre allocation for the full race weekend, shared across all practice sessions, qualifying and the race itself. Pirelli provides 13 sets of slick tyres, usually two sets of hard, three sets of medium, and eight sets of soft, and drivers who advanced to Qualifying 3 receive an additional set of soft tyres [7]. This allocation creates a resource constraint that must be factored into strategy planning, since sets consumed during practice or qualifying reduce what is available for the race.

On the other hand, when it is predicted to rain during the Grand Prix, or conditions change mid-race, Pirelli substitutes the slick tyres with the Cinturato rain tyre range, which features tread patterns designed to evacuate water and prevent aquaplaning. The intermediate tyre, identified by its green colour, is used in the transition between dry and wet conditions, since it is capable of clearing on average 30 litres of water per second at high speeds. The full wet tyre is characterised by the colour blue and is reserved for heavy rain conditions, capable of clearing 85 litres of water per second. However, their characteristics make them more susceptible to overheating and rapid degradation if the weather condition of the track starts to dry, since the tread patterns that make them effective in wet conditions generate excess heat on a dry surface. Unlike slick compounds, wet-weather tyres are not pre-selected by Pirelli, and are available at every race weekend, and the decision to deploy them is made reactively based on conditions [5].

Furthermore, the selection between all these compounds creates a multi-objective optimisation problem. Each tyre compound can be characterised by an operation window, based on the temperature ranges within which grip is maximised, and by its degradation slope across a stint. The degradation of the tyre is not only the mechanical wear where the rubber physically vanishes through contact with the asphalt, but also involves a thermal degradation, whereby the rubber undergoes chemical and physical changes that will lead to the "cliff". The cliff represents a non-linear collapse in performance where the grip levels fall significantly and the car becomes a bottleneck on the track [8]. Therefore, predicting the lap on which the cliff will occur for a given compound, under a specific track and fuel load, is one of the central challenges in race strategy.

Therefore, for each team, the goal is to maximise the time spent within the peak performance window of each compound, while complying with the mandatory pit stop required by FIA regulations.

2.2 Regulatory Constraints

The optimisation of a race is not an unbounded problem, because it is strictly constrained by FIA's regulations. The most significant change was in 2010 when mid-race refuelling was prohibited [9]. In the past, the time taken in the pit stops was a function of fuel load, while now it is dictated by tyre degradation. This single regulatory change transformed the pit stop from a fuelling operation into a pure tyre-strategy manoeuvre, making the number, timing, and compound selection of tyre changes the primary strategic variable in modern Formula 1.

During the race, it is mandatory to use at least two different dry-weather tyre compounds, imposed by FIA, to prevent a zero-stop strategy that would allow a driver to complete the race without losing time in the pit lane. Beyond the competitive rationale, this rule also serves a safety purpose. Since tyres have a specific operational lifespan, and if they are not changed, track conditions would no longer be safe. Therefore, teams are responsible for determining the best sequence of tyres to minimise the total race time, and to decide whether to pit or not at the beginning of each lap [10].

Furthermore, there are specific rules inside the pit lane for safety. There is a pit lane speed limit, typically 80 km/h, or reduced to 60 km/h for safety precautions in specific tracks [11], otherwise exceeding this limit results in a time penalty during the race, as well a financial penalty of 100€ for each km/h above the limit, up to a maximum of 1000€ [2]. Besides the speed limit, a car entering the pit lane loses between 20 and 30 seconds, depending on the length of the pit lane [9], and this time loss represents the fundamental trade-off at the heart of every strategy decision: the performance benefit of fresh tyres must outweigh the time penalty of entering the pit lane.

2.3 Race Strategy Typology

F1 races are divided into stints, which represent the number of consecutive laps driven on a single set of tyres, with the number of stints being equal to the number of pit stops plus one. The choice of how many stints to run, which compound to assign to each, and at which lap to execute each pit stop defines the core combinatorial decision that strategy engineers must solve before and during every race.

Furthermore, strategy can be either pre-planned or reactive. A pre-planned strategy is defined before the race based on simulation outputs, tyre allocation constraints, and expected degradation profiles. A reactive strategy is adjusted live during the race in response to rival pit stops, safety car periods, or unexpected tyre behaviour on track. In practice, teams always enter the race with a primary plan and two or three contingency strategies ready to deploy, since conditions rarely unfold exactly as modelled. A driver's strategy is often a response to the actions of their rivals, so the timing of transitions between stints depends heavily on competitive positioning [9].

There are two fundamental strategies in the competitive landscape, denominated as the undercut and the overcut. The undercut strategy refers to a pit that is made earlier than the rival to exploit the speed advantage of fresh tyres on the out-lap, so that when the rival subsequently pits the driver who has already pitted may emerge ahead, gaining track position and benefiting from the pace delta of new rubber. On the other hand, the overcut involves staying out on track longer than a rival who has just pitted, pushing hard on tyres that still have sufficient grip while the rival struggles through the warm-up phase of their new compound until it reaches optimal operating temperature [12]. So, the overcut is most effective when the compound in use retains enough performance to sustain fast lap times, and when the rival's new compound has a long warm-up phase, as is typically the case with harder compounds. Therefore, teams must constantly monitor the gap to the cars ahead and behind to ensure that after a pit stop the driver does not lose positions unnecessarily.

Strategy is also shaped by how easily overtaking can be achieved on a given circuit, since the value of track position is directly related to the difficulty of passing. Thus, in this context, the Drag Reduction System (DRS) has historically played an important role, which allowed a driver to activate the system, which is a rear wing opening, which reduces aerodynamic drag and gains approximately 10-12 km/h when running within one second of a rival, to help overtake or protect their position [13] [12]. However, from 2026, DRS will be replaced by the Manual Override system, which delivers a power boost from the electric Motor Generator Unit - Kinetic (MGU-K) unit when a driver is within one second of a car ahead [14].

2.4 Operational Uncertainty & Dynamic Variables

In a context like this, the strategy optimisation can be interpreted either as a deterministic model, assuming a perfect environment, or a stochastic model that accounts for unpredictable events. The uncertainty of the project comes from the stochastic variables, as events that cannot be predicted in advance, such as accidents that trigger a Safety Car (SC) or Virtual Safety Car (VSC). These events represent unpredictable interruptions that may fundamentally shift the opportunity cost of a pit stop changeover, and are precisely what separates a theoretical optimal solution from a robust real-world strategy.

2.4.1 Safety Car

The Safety Car is deployed when a major on-track incident requires the field to be neutralised under the control of a physical safety car, which enters the track and leads all competing cars at a reduced pace, and has a 47 % probability of being deployed [9]. Overtaking is prohibited for the duration of the yellow flag. All cars are queued behind the SC at a slow pace, which eliminates all time gaps between competitors, effectively compressing the field and resetting the relative positions built up over the preceding laps [9]. This compression opens a low-cost pit window, since all cars are already slowed, and the time lost by entering the pit lane is significantly reduced compared to normal racing conditions, making it an opportune moment to change tyres without surrendering as much track position as a standard stop would cost. In this case, the strategy of pitting is a trade-off because it allows the driver who pits to restart the race with new fresh tyres, but if the rival does not pit, the one who did may lose track positions upon the restart. Furthermore, with a slower pace, the tyres will lose operating temperature, which increases the risk of degradation and reduced grip in the critical laps immediately following the restart, when the pace suddenly returns to race speed. Teams must therefore decide in real time whether to pit or stay out, weighing the benefit of fresh tyre performance against the cost of surrendering track position.

2.4.2 Virtual Safety Car

On the other hand, the Virtual Safety Car that can be activated with a probability of 30% [9], is applied when there is a minor incident that requires the drivers to reduce their speed by approximately 30% and maintain a minimum lap delta time, but does not necessitate a physical car on track [15]. Since cars must slow down, it opens an efficiency window for pit stops, as the pit lane speed limits remain constant and unchanged. Therefore, under normal racing circumstances a pit stop costs around 20-30 seconds of race time, while under VSC conditions it costs 10 to 15 seconds. This reduction effectively halves the time penalty of a pit stop, creating a clear strategic opportunity for teams that react quickly. At the same time, it extends tyre life during the VSC period, as wear accumulates 50% more slowly at the reduced speed [16]. Teams that correctly anticipate or react to a VSC window can therefore gain 10 to 15 seconds over rivals who miss the opportunity.

2.4.3 Fuel Mass

Moreover, the fuel mass effect acts as a continuous variable throughout the race. A car begins with approximately 110 kilograms of fuel, making it heavy and accelerating tyre wear due to the increased mechanical load on the rubber. As fuel is consumed, the car becomes progressively lighter, typically generating a pace gain of approximately 0.2 seconds per lap [10]. Consequently, the degradation slope is not constant across a stint, as a tyre fitted at the start of the race under heavy fuel loads wears faster than the same compound fitted later in the race on a lighter car. This means that the same compound will have a longer effective lifespan if fitted in the second half of the race compared to the first [17]. Models that assume a constant degradation rate, risking overestimating tyre wear late in a stint and underestimating it earlier on, can lead to suboptimal pit-windows decisions.

2.4.4 Weather Conditions

Weather is the most disruptive stochastic factor in Formula 1, representing a fundamental regime change in the operating environment. The onset of rain necessitates a transition from slick tyres to the wet-weather range, intermediates or full wets, and with it, an entirely different set of compound characteristics, degradation profiles, and lap time estimates must be applied simultaneously. The core optimisation challenge is identifying the "Crossover Point", the exact moment when the grip advantage of a rain tyre outweighs the pit stop penalty compared to remaining on slicks [17]. Misjudging this threshold by even a single lap can lead to a catastrophic loss of pace, often 10 to 15 seconds per lap, or a total failure state through aquaplaning on slicks. Furthermore, as the track dries, a second crossover back to slicks occurs. This turns the timing of the pit stop into a high-stakes exercise in risk management, where the engineer must balance the thermal degradation of rain tyres on a dry line against the mechanical grip required on the remaining wet patches [10].

2.5 Dimensions of Race Performance

The final outcome of a Formula 1 Grand Prix, and consequently of a season, is not determined by a single variable, but rather by the interaction of several distinct performance dimensions. The first is the technical and aerodynamic dimension, encompassing car development, power unit performance, aerodynamic efficiency, and reliability, which together define the car's raw competitive potential. This dimension is, to a large extent, the one rewarded by the Constructors' Championship, since it aggregates the results of both drivers on a team as a proxy for the overall performance of the car and the engineering effort behind it.

The second dimension concerns driver performance, including raw driving skills and racecraft, such as overtaking and defending track position, as well as the ability to manage the car's resources over the course of a stint, like tyre degradation, and from the 2026 season onwards, also the deployment of electrical energy under the new power unit regulations.

The third dimension is strategic decision-making, namely pit stop timing and tyre compound sequencing, together with the capacity to react to unfolding race conditions such as rival strategies, safety car periods, or changing weather. None of these dimensions operate in isolation, and none alone is sufficient to determine the race outcome: a strong car or skilled driver can still be undermined by a poor strategic decision, just as a well-executed strategy cannot compensate for a fundamentally uncompetitive car or an underperforming driver. This thesis focuses specifically on the last of these dimensions, strategic decision-making, treating a car and driver performance as fixed, exogenous inputs, rather than variables to be optimised.

Therefore, each of these three areas has, in its own right, been the subject of extensive prior research, and the dimension addressed in this work is closely related to the field of Industrial Engineering and Management.

Chapter 3

Literature Review

The optimisation of a Formula 1 race strategy represents a complex, multidisciplinary challenge characterised by high-stakes decision-making problems. Building upon established academic frameworks, this research can be approached through Discrete-Event Simulation (DES), Operations Research, Game Theory, Artificial Intelligence and Machine Learning. The literature in this field follows a clear chronological progression, driven largely by the growing availability of real-time telemetry and historical race data. Early works were primarily descriptive and deterministic, focused on evaluating pre-defined strategies under fixed assumptions. Over time, as data availability grew and computational methods advanced, the field shifted towards prescriptive and stochastic approaches capable of generating optimal strategies and accounting for the unpredictable events that characterised real race conditions.

The following sections organise this progression across the four methodological streams, identifying the contributions and limitations of each, before synthesising the remaining research gap that this thesis seeks to address.

3.1 Simulation-Based Approaches

The first methodological approach stream in the literature approached F1 strategy optimisation through Discrete-Event Simulation, treating race events as sequences of discrete occurrences that could be modelled and replayed under different input conditions.

The pioneering work in this area was conducted by Bekker and Lotz [18], who provided a fundamental framework to model on-track events such as passing manoeuvres, car failures, and pit stops. The work aims to provide a decision-support system to allow teams to compare several different scenarios before a race. However, in this case, the model is designed to evaluate specific input strategies rather than the optimal strategy to use in the race, since it is primarily descriptive and scenario-based.

Secondly, Heilmeier et al. [13] added to the simulation a lap-wise discretisation, by treating each lap as a discrete event, in order to incorporate dynamic variables such as fuel mass loss and tyre degradation coefficients. Therefore, it allows a more accurate representation of the sport. However, it still only evaluates deterministic scenarios, not representing uncertain scenarios such as safety cars, which are common.

Both works therefore share a fundamental limitation. Both can evaluate how a given strategy performs, but not generate the optimal strategy in the first place. Furthermore, neither model accounts for the stochastic events that are a regular feature of real racing conditions and that can render a preplanned strategy suboptimal within a single lap.

3.2 Operations Research and Mathematical Optimisation

As data availability in Formula 1 grew, researchers moved beyond simulations towards Operations Research, which offered the ability not only to evaluate strategies, but also to formally generate optimal ones under mathematical constraints, solving what can be defined as the “tyre compound and timing” problem. Heine and Thraves [10] utilise Dynamic Programming to determine the optimal pit window for a single driver. Their model also includes uncertainty by modelling the probability of yellow flag events to decide whether to accelerate or delay the pit stop for tyre changing. This represents a significant methodological step forward relative to the simulation literature, as it incorporates stochastic events directly into the optimisation framework rather than treating them as external disruptions.

In 2024, García [8] proposed a data-driven approach using Mixed-Integer Quadratic Programming (MIQP) to achieve a strategy decision and Bayesian Regression to estimate tyre performance during the race. Therefore, by calibrating the regression model with data from previous years’ race as “prior belief” and updating with real-time telemetry, the system generates input for the MIQP, which is then able to solve the optimal number of stints and tyre compounds in a fraction of a second. This combination of statistical learning and mathematical optimisation represents one of the most computationally complete approaches in the literature, capable of adapting to real-time race conditions while remaining grounded in a formal optimisation framework.

However, both models share a critical assumption: that each driver races in isolation, optimising their own strategy without accounting for the decisions of their rivals.

3.3 Game Theory and Competitive Interaction

A major limitation of the OR models already discussed is the assumption that drivers race in isolation. In reality, every strategic decision made by one team is simultaneously a signal to their rivals, and the optimal response depends directly on what the competitors choose to do. This interdependence is precisely the domain of Game Theory, which provides a formal mathematical framework for modelling strategic interactions between rivals.

Therefore, Aguad and Thraves [9] addressed this by using a Zero-Sum Feedback Stackelberg Game to make decisions on pit stops. In this method, the research proves that strategy is not just about tyre wear, but also competition between the “Leader” and the “Follower”. The Stackelberg framework is particularly well suited to F1, as it captures the asymmetry between the driver ahead, who sets the strategic pace, and the driver behind, who must decide whether to react or commit to an independent plan.

Furthermore, the simulations highlight the undercut and overcut and mathematical outcomes of a non-cooperative game, comparing scenarios where teams maximise their probability of winning versus minimising race time. The study also demonstrates that a strategic driver who accounts for their rival's strategy increases their winning odds by more than 15% against a non-strategic opponent. This finding has a direct practical implication: ignoring competitive interaction in a strategy model does not only reduces theoretical accuracy, but also produces strategies that are systematically exploitable by rivals who do account for it.

However, the Stackelberg model also has its own limitations. It is structured as a two-player zero-sum game, which simplifies the multi-car reality of Formula 1 racing, where a driver must simultaneously manage threats from behind and opportunities ahead across a field of twenty cars. These constraints, combined with the computational intensity of game-theoretic approaches at scale, motivated the development of Artificial Intelligence (AI) and heuristic methods capable of handling greater complexity with less restrictive assumptions.

3.4 Artificial Intelligence and Heuristics

As the limitations of both OR and Game Theory become apparent, a fourth methodological stream emerged drawing on AI and heuristic methods, which learn directly from historical data and can adapt to complexity that is difficult to model analytically.

Heilmeier et al. [16] developed a Virtual Strategy Engineer (VSE) using Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) architectures, trained on timing data across six seasons from 2014 to 2019. The system decides when to pit by overcoming the technical challenge of "imbalanced data" since pit stops are rare events relative to the total number of laps in the dataset. The LSTM architecture is particularly well suited to this problem, as it is designed to capture long-range temporal dependencies, which in this context are the relationships between tyre age, lap-time evolution, and the eventual decision to pit across the course of a race. Thus, unlike OR that requires explicit mathematical formulas, VSE relies on learned patterns and AI to adapt to race situations.

Furthermore, Bonomi et al. [17] explores the use of Genetic Algorithm (GA) to handle a massive solution space of race strategy. Each sequence of tyre and pit laps are considered "chromosome", and the GA simulates thousands of strategies to converge on the optimum one. This evolutionary approach is effective at navigating large combinatorial search spaces without requiring the problem to be convex or analytically tractable, making it well suited to the multi-compound, multi-stint structure on F1, and useful when handling multi-objective optimisation, considering stochastic events.

Finally, Choo [1] analysed features within a machine learning model to demonstrate that pit crew performance and "fresh air" windows are statistically significant predictors of race positions, emphasising that perfect tyre strategy must be supported by operational consistency. This finding reinforces a point that is easy to overlook in purely mathematical models: that the execution of a strategy is as important as its design, and that the variability in pit crew performance introduces an additional stochastic element that analytical models rarely capture explicitly.

Despite their flexibility and predictive power, AI and heuristic methods share an important limitation. They can identify when to pit and which compound to select based on learned patterns, but they offer limited interpretability, which is a critical drawback in a high-stakes environment where engineers must justify strategic decisions in real time and understand why a recommendation is being made, and not just what it is. This interpretability gap, combined with the data dependency of trained models, means that AI approaches complement rather than replace the analytical frameworks developed in the OR and Game Theory literature.

3.5 Synthesis and Research Gap

The literature shows that over the years, F1 strategy modelling has evolved from simple simulations to sophisticated optimisation frameworks, shifting from single-driver studies toward complex game-theoretic analyses. However, a remaining gap exists in the Operations Research literature. The most robust model to date, proposed by García [8], optimises tyre selection and pit windows by using a deterministic framework that treats race variables as point estimates rather than probability distributions. However, this deterministic view falls short of the reality on the track, where tyre degradation and pit stop durations are volatile and largely unknown until the lights go out. Therefore, this project bridges this gap by evolving García's work into a stochastic programming framework. By treating key performance metrics as random variables rather than static values, this model minimises expected race time across a range of scenarios. Finally, this thesis seeks to synthesise these elements to determine how the specific choice of tyre compound can be utilised as both a mechanical and a tactical tool to maximise race output.

Chapter 4

Optimisation Model

This chapter presents the mathematical formulation of the race strategy optimisation model developed to determine the optimal pit stop timing and tyre compound selection for a Formula 1 Grand Prix. Therefore, this model is formulated as a two-stage stochastic MILP problem.

The first stage involves finalising the comprehensive race strategy, including defining the compound pairs, stint lengths, and pit stop timing, which is fixed before any uncertainty is applied, reflecting the real operational constraint that F1 teams must commit to a primary strategy before the formation lap, without knowledge of how race parameters, such as tyre degradation rate, pit stop durations, and evolution of the car along the stints, will materialise on race day. Moreover, while the second stage does not introduce adaptive recourse decisions, it evaluates the cost performance of the first-stage strategy under each discrete scenario, which quantifies the expected total race time as a probability-weighted average across the scenario set. Therefore, the model is more precisely characterised as a "here-and-now" stochastic optimisation with scenario-dependent cost evaluation, where the strategy is decided once and uncertainty affects only the objective value, rather than the decisions. So, the two-stage terminology is retained throughout to reflect this structural separation between the decision phase and evaluation phase, with the understanding that no strategy adjustment is made once the uncertainty is revealed.

To define the optimisation problem, the model requires the definition of the structural index sets that characterise the race, the parameters that characterise the circuit and tyre properties, and the decision variables that represent the strategic choices available to the team. Then, the objective function is formulated to capture the track time and pit lane time, while the constraints represent both FIA sporting regulations and the physical limitations, that together define the feasible region within which the optimal strategy is sought. Additionally, certain equations were derived or adapted from the formulations presented by García [8].

4.1 Sets

To formulate the two-stage stochastic race strategy for a Formula One team as an optimisation problem, the first step is to define the sets and indices that represent the structural components of the race, the available tyre inventory, and the temporal choices available to the decision-maker.

- Race Stints ($\mathcal{I}$): Represents the set of available race stints, indexed by i, such that $i \in \mathcal{I} = \{1, 2, \ldots, S\}$, where S denotes the maximum number of stints per race, and a stint is defined as the sequence of laps driven between pit stops under the same tyre compound.
- Tyre Compounds ($\mathcal{T}$): The set of available dry tyre compounds, indexed by t, such that $t \in \mathcal{T} = \{\text{Soft}, \text{Medium}, \text{Hard}\}$, which represents the different dry compounds that Pirelli provides.
- Physical Tyre Sets ($\mathcal{S}_t$): Represents the set of specific tyre sets available in the team's allocated inventory for a given compound t, indexed by s, such that $s \in \mathcal{S}_t$.
- Stint Duration ($\mathcal{K}_t$): The set of possible stint durations, measured in laps, for a given compound t, indexed by k, such that $k \in \mathcal{K}_t = \{1, \ldots, L_t\}$, where L_t denotes the maximum possible number of laps that a single set of tyres can safely complete, or in other words, the lifespan of a tyre compound.
- Scenarios (Ω): Represents the set of discrete scenarios that represent the uncertain race conditions, indexed by ω, such that $\omega \in \Omega = \{1, 2, \ldots, W\}$, where W denotes the total number of specific scenarios under different conditions.

4.2 Parameters

The system parameters are divided into two categories: deterministic parameters, which are known and fixed prior to the race, and stochastic parameters, which vary depending on the unfolding race scenario.

Deterministic Parameters

These parameters represent the fixed structural limitations of the circuit, the physical properties of the tyres, and the probabilistic weights of the model.

- N: represents the total number of laps required to complete the race distance, which characterises each circuit.
- t_{base}: is the base lap time, representing the theoretical baseline time required to complete a single lap around the circuit, with brand new soft tyres and with no traffic.
- L_t: relating to the tyres, represents each compound lifespan, denoting the maximum safe number of laps that a specific compound can complete before structural failure.
- p_ω: the probability of occurrence for scenario ω, determining the weight of each scenario in the expected value calculation, and:

$$\sum_{\omega \in \Omega} p_\omega = 1 \tag{4.1}$$

Uncertain Parameters

To capture the real-world unpredictability in the Grand Prix, the two-stage stochastic programming approach introduces parameters that depend strictly on the realised scenario ω. Furthermore, to capture the evolution of the race, these parameters also depend on the specific race stint i:

- $P^{\omega}_{pit,i}$: denotes the pit stop time loss in scenario ω, representing the time penalty incurred by the driver when navigating the pit lane to change tyres, under FIA's rules. This parameter is uncertain to capture its inherent variability due to pit crew performance or pit lane traffic.
- β^{ω}_{t0}: quantifies the pace penalty of a tyre compound when it is fresh, with zero laps of wear, for every scenario ω. It measures the speed difference between compounds under identical conditions, and it represents the trade-off between speed and durability of the tyre.
- $\hat{\beta}^{\omega}_{t1,i}$: is the wear acceleration factor specific to tyre compound t in scenario ω during stint i, which dictates the increase in lap-time degradation as the tyre wears down. By incorporating the stint index i, the model captures vehicle physics and track evolution.

$$\hat{\beta}^{\omega}_{t1,i} = M_i \cdot \beta^{\omega}_{t1} \tag{4.2}$$

To obtain the uncertain parameter $\hat{\beta}^{\omega}_{t1,i}$, a multiplier M_i is utilised, which is a deterministic constant for each stint i, where $M_1 > M_2 > \cdots > M_S$. This multiplier represents the physical impact of fuel load and track conditions: in early stints, the car carries a heavy fuel load and operates on a track with low grip, resulting in high degradation. In later stints, the fuel load is depleted and the track is "rubbered in," significantly reducing the degradation rate of the tyre compound. Here, β^{ω}_{t1} is the base wear acceleration factor for compound t in scenario ω.

4.3 Decision Variables

This model utilises a two-stage stochastic programming approach, where the decision-making process is divided into two distinct phases, and consequently, the decision variables are categorised into first-stage and second-stage decisions. The first stage represents the definitive race strategy established by the team before the race starts, while the second stage represents the realisation of the uncertain race conditions and evaluates the time performance of the chosen strategy under each specific scenario. Therefore, the optimisation model relies on a primary binary variable to determine the optimal race strategy, along with auxiliary variables derived from the primary one to track stint activation, lap counts, and compound usage.

First-Stage Decision Variable

The primary first-stage variable is x_{ikts} that takes the value 1 if the stint i utilises the compound t for exactly k laps with set s, and 0 otherwise. This binary decision variable is the core variable of the model, as its multidimensional nature simultaneously dictates the tyre selection, the sequence of the stints, and the pit stop timing for the entire race.

To facilitate the formulation of the objective function, the variables u_i and y_i are created in order to have a linear formula to use in the MILP model.

Firstly, u_i is an integer variable representing the total number of laps driven during a specific stint for each scenario. Therefore, it is mathematically defined by the multiplication of the duration k and the binary variable x_{ikts} across all possible tyre compounds, available sets, and durations, as in the following equation:

$$u_i = \sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} k \cdot x_{ikts} \tag{4.3}$$

Moreover, y_i is a binary indicator variable that tracks the compound usage in each stint. So, it is equal to 1 if the stint i is used during the race, and 0 otherwise:

$$y_i = \sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} x_{ikts} \tag{4.4}$$

Additionally, an auxiliary binary variable, $q_i \in \{0, 1\}$ was introduced to track the occurrence of a pit stop between two consecutive stints. Specifically, q_i takes the value of 1 if and only if both stint i and $i + 1$ are active, hence q_i originates from:

$$q_i = y_i \cdot y_{i+1} \tag{4.5}$$

Lastly, the auxiliary binary variable z_t is introduced to define the regulatory compound constraint made by the FIA, which equals 1 if compound t is used at any point during the race, and 0 otherwise.

Second-Stage Decision Variables

While the strategy is fixed in the first stage, the actual time it takes to execute the strategy depends on how the race unfolds. Therefore, the second-stage variables are continuous auxiliary variables that calculate the operational cost under each specific scenario.

The primary second-stage variable is $T^{\omega}_{\text{total}}$, which represents the total race time achieved by the chosen strategy under the scenario. This calculation relies on the uncertain parameters to evaluate the two main time components: the time lost in the pit lane, dictated by the pit stop penalty ($P^{\omega}_{\text{pit,i}}$), and the time spent on track influenced by the tyre degradation ($\hat{\beta}^{\omega}_{t1,i}$).

4.4 Objective Function

The primary goal of the two-stage programming problem is to determine the optimal race strategy by minimising the expected total time required to complete the race. As a stochastic model, the objective function separates the deterministic component, representing the cost known prior to the race, from the stochastic component, which represents the expected costs that depend strictly on the unfolding race scenario. Therefore, the objective function is formulated as the sum of calculating the deterministic track time and the expected value of the uncertain parameters, yielding the minimum expected total race time, ($\mathbb{E}[T_{\text{total}}]$), weighted by their respective probability of occurrence (p_ω).

Track Time and Tyre Degradation

In order to model the total track time, the duration of an entire stint must be established. The duration of any single lap is modelled as the sum of a baseline performance constant and a variable penalty related to the tyre's state. Following a linear degradation model adapted from García [8], a temporary index j is introduced to represent a specific lap within a stint, where $j \in \{1, 2, \ldots, k\}$. Assuming the first lap has zero degradation beyond the base compound penalty, the time of a lap j is defined as:

$$Time\_of\_lap\_j = t_{base} + \beta^{\omega}_{t0} + \hat{\beta}^{\omega}_{t1,i} \cdot (j-1) \tag{4.6}$$

To identify the optimal timing for a pit stop, the time of an entire stint of k laps must be calculated. Since each lap is slower than the previous one by a fixed amount $\hat{\beta}^{\omega}_{t1,i}$, the sum of the lap times forms an arithmetic series. Therefore, by applying the mathematical rules for the sum of an arithmetic series over k laps, the cumulative time for a single stint is obtained:

$$StintTime = \sum_{j=1}^{k} \Big(t_{base} + \beta^{\omega}_{t0} + \hat{\beta}^{\omega}_{t1,i} \cdot (j-1) \Big) \tag{4.7}$$

$$StintTime = (k \cdot t_{base}) + (k \cdot \beta^{\omega}_{t0}) + \hat{\beta}^{\omega}_{t1} \cdot \left(\frac{k^2 - k}{2} \right) \tag{4.8}$$

However, in a MILP model, an integer decision variable cannot be squared, as this would create a non-linear cumulative degradation. Therefore, to overcome this limitation, a linearisation technique is utilised through the primary binary variable x_{ikts}. The components of the stint time are pre-calculated as constant parameters for every possible combination and subsequently multiplied by the decision variable.

First-Stage Component

The deterministic track time, T_{det}, represents the minimum guaranteed time the car will spend on track and is constant for any chosen stint duration k and compound t.

$$T_{det} = \sum_{i \in I} \sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} (k \cdot t_{base}) \cdot x_{ikts} \tag{4.9}$$

Second-Stage Components

The stochastic portion captures the elements that depend strictly on the realised scenario ω, so it groups the initial tyre compound penalty and the tyre wear acceleration. The first uncertain element is the track time that represents the tyre degradation, C_{tks}^{ω}. Using the previously derived arithmetic sum, the quadratic degradation cost can be pre-calculated for every possible combination of compound t, duration k, and scenario ω.

$$C_{tks}^{\omega} = (k \cdot \beta_{t0}^{\omega}) + 0.5 \cdot \hat{\beta}_{t1,i}^{\omega} \cdot (k^2 - k) \tag{4.10}$$

Moreover, the second stochastic component accounts for the time lost in the pit lane. To capture variables such as pit lane traffic, while respecting FIA rules to reduce velocity, this penalty is dependent on both the scenario ω and the specific stint i being initiated, denoted by $P_{pit,i}^{\omega}$. This penalty is incurred for every stint transition; to model it, the pit penalty for stint i is multiplied by the pit stop transition variable q_i.

Final Objective Function

Consequently, by minimising the sum of the deterministic track time and the expected value of both the stochastic degradation costs and stint-dependent pit stop penalties, the objective function drives the solver to find the fastest possible combination of tyre compounds, stint lengths, and pit stop timing:

$$\min \mathbb{E}[T_{total}^{\omega}] = T_{det} + \sum_{\omega \in \Omega} p_{\omega} \cdot \left[\sum_{i \in I} \sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} C_{tks}^{\omega} \cdot x_{ikts} + \left(\sum_{i=1}^{S-1} P_{pit,i}^{\omega} \cdot q_i \right) \right] \tag{4.11}$$

4.5 Constraint

The final stage of an optimisation model is the definition of the feasible region. The constraints in this research represent sporting regulations and physical limitations of the car. Therefore, any strategy that minimises the objective function but violates these rules is discarded by the model.

Linearization

To maintain the linearity of the MILP model, the product of the binary stint activation variable y_i and y_{i+1} is represented by the auxiliary variable q_i. To ensure that this variable correctly reflects the logic of a pit stop occurring between two stints, the following linear constraints are enforced for all $i \in 1, \ldots, S-1$

$$q_i \leq y_i \tag{4.12}$$

$$q_i \leq y_{i+1} \tag{4.13}$$

$$q_i \geq y_i + y_{i+1} - 1 \tag{4.14}$$

$$q_i \geq 0 \tag{4.15}$$

These constraints force $q_i = 1$ to ensure that the pit stop penalty is only incurred during the transition between stints, in order to prevent the model from applying penalties at the start of the race or after the final stint (S).

Stint Definition

The next constraint in this model ensures the logical structure of the race strategy. For every available stint in the model, exactly one combination of tyre compound, specific set, and duration must be selected, but only if that stint is actually activated by the model ($y_i = 1$).

$$\sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} x_{ikts} = y_i, \quad \forall i \in I \tag{4.16}$$

Furthermore, to restrict the computational boundaries, the total number of stints used in any scenario is bounded by a minimum of two, implied by the FIA rule. Additionally, a constraint was introduced to limit the number of stints by a maximum threshold (S) a value that can be adjusted by the decision-makers.

$$2 \leq \sum_{i \in I} y_i \leq S \tag{4.17}$$

Stint Consecutiveness

To ensure that all active stints are assigned consecutively from the beginning of the race, preventing the model from selecting non-sequential stints, the following constraint is enforced for all $i \in \{2, \ldots, S\}$:

$$y_i \leq y_{i-1}, \quad \forall i \in \{2, \ldots, S\} \tag{4.18}$$

This constraint is fundamental to the model's structure, as it ensures that active stints follow a strictly sequential order. This maintains the integrity of the stint-specific degradation multiplier M_i, ensuring that the high fuel load penalties associated with early stints are correctly aligned with the first segments of the race.

Total Race Completion

The most fundamental logical constraint in a Grand Prix is that a driver must complete the full distance of the race. Therefore, the sum of the laps driven in each stint must be equal to the total number of laps, N, required for that specific circuit.

$$\sum_{i \in I} \sum_{t \in T} \sum_{s \in S_t} \sum_{k \in K_t} k \cdot x_{ikts} = N \tag{4.19}$$

Tyre Longevity

Each tyre has a maximum structural lifespan to avoid any risk of mechanical failure. The duration of any stint is bounded by the maximum life, L_t, of the chosen compound t. So L_t is a safety threshold defined by the tyre manufacturer, but also depends on external factors.

$$k \cdot x_{ikts} \leq L_t \cdot y_i, \quad \forall i \in I, \forall t \in T, \forall s \in S_t, \forall k \in K_t \tag{4.20}$$

FIA Rule

FIA sporting regulations state that in a dry race, every driver must utilise at least two different dry-weather tyre compounds, in order to prohibit a race with zero pit stops. Therefore, the auxiliary binary variable z_t was introduced.

In order to force z_t to take the value 1 if compound t is used, a bounding constraint was created where the total number of stints using the same compound is lower than the "Big-M", which in this case is the total race laps, N.

$$\sum_{i \in I} \sum_{s \in S_t} \sum_{k \in K_t} x_{ikts} \leq N \cdot z_t, \quad \forall t \in T \tag{4.21}$$

However, to strictly ensure that $z_t = 1$ if and only if the compound t is used in at least one stint, a lower bound is required.

$$z_t \leq \sum_{i \in I} \sum_{s \in S_t} \sum_{k \in K_t} x_{ikts}, \quad \forall t \in T \tag{4.22}$$

Moreover, to ensure that in each race there is at least one pit stop involving a tyre compound change, the following constraint was created, strictly enforcing the two-compound rule.

$$\sum_{t \in T} z_t \geq 2 \tag{4.23}$$

Inventory

Last but not least, for every circuit Pirelli supplies a specific number of sets of tyres for each compound. Therefore, for the race, a team only has a few sets of the compound and every time it changes the tyres in the pit lane it must change to a new sets. Therefore, the constraint (4.24) ensures that any specific physical tyre set s is utilised for at most one stint during the entire race, preventing the reuse of a set. Meanwhile, the constraint (4.25) bounds the total number of stints for each compound t by the total number of available sets provided by Pirelli.

$$\sum_{i \in I} \sum_{k \in K_t} x_{ikts} \leq 1, \quad \forall t \in T, \forall s \in S_t \tag{4.24}$$

$$\sum_{i \in I} \sum_{k \in K_t} \sum_{s \in S_t} x_{ikts} \leq S_t, \quad \forall t \in T \tag{4.25}$$

Chapter 5

Data

To apply the optimisation model, this research utilises the data present on the Kaggle repository ”Formula 1 World Championship” from 1950 to 2024, [19]. This longitudinal dataset provides high-fidelity information across the sport, including lap-by-lap timing, pit stop durations, circuit characteristics, and specific tyre compound usage in each stint. While the dataset encompasses racing history from 1950 to 2024, this research narrows its scope to the 2022-2024 seasons. In this specific window, the technical era of Formula 1 is characterised by 457 millimetres (18-inch) wheel rims, which reduce overheating, and ground-effect aerodynamics that significantly influence tyre degradation, consequently, overall tyre diameter also increased from 670 mm to 725 mm in 2022 [20–23]. Therefore, focusing on this specific window ensures that the model parameters and inputs remain consistent and have a homogeneous impact on the final outcome. Seasons outside this window are excluded for specific operational reasons: the 2025 season is omitted as the championship had not yet concluded and no empirical data was available during the data collection phase, while the 2026 season is excluded due to massive upcoming regulatory changes. The 2026 cars will feature radically different performance characteristics, lacking the Drag Reduction System (DRS) and shifting towards optimised electrical energy and power management.

Both the primary stochastic model and its deterministic version share several core parameters. The base lap time t_{base} represents the potential pace of the track under optimal conditions, and so was derived from the Kaggle dataset by identifying the fastest time in Qualifying 3 for each race, which is also the time of the driver in the pole position. To capture the absolute minimum baseline time across these years, the base time lap was calculated by taking the minimum value recorded in Qualifying 3, between the 2022 and 2024 seasons for each circuit.

Additionally, the race distance is standardised by using the official 2024 lap counts [4], with the exception of the French Grand Prix, whose lap count is instead taken from the 2022 season, as it was not part of the 2024 calendar, as there were only minor alterations to track distance during this period. So, for this model, the values utilised are represented in Table 5.1.

Regarding the tyre inventory, for each race weekend, Pirelli provides 13 sets of dry-weather tyres, being eight soft, three medium, and two hard, plus four sets of intermediate and three sets of full-wet tyres [24]. However, accounting for the practice sessions and qualifying, the available inventory for the

Table 5.1: Race Information: Number of laps and base lap time, N and t_{base}, respectively

Race	*N* (laps)	t_{base} (s)
Bahrain Grand Prix	57	89.179
Saudi Arabian Grand Prix	50	87.472
Australian Grand Prix	58	75.915
Emilia Romagna Grand Prix	63	74.746
Miami Grand Prix	57	86.841
Spanish Grand Prix	66	71.383
Monaco Grand Prix	78	70.27
Azerbaijan Grand Prix	51	100.203
Canadian Grand Prix	70	72
British Grand Prix	52	85.819
Austrian Grand Prix	71	64.314
French Grand Prix	53	90.872
Hungarian Grand Prix	70	75.227
Belgian Grand Prix	44	103.665
Dutch Grand Prix	72	69.673
Italian Grand Prix	53	79.327
Singapore Grand Prix	62	89.525
Japanese Grand Prix	53	88.197
United States Grand Prix	56	92.33
Mexico City Grand Prix	71	75.946
São Paulo Grand Prix	71	70.727
Abu Dhabi Grand Prix	58	82.595
Qatar Grand Prix	57	80.52
Las Vegas Grand Prix	50	92.312
Chinese Grand Prix	56	93.66

race itself was standardised to four sets of soft, three sets of medium, and two sets of hard tyres.

Furthermore, tyre lifespan is a critical constraint in this optimisation problem, as it is not recommended for a vehicle to utilise a tyre beyond its operational limit. These values vary based on factors such as specific race circumstances, driving style, and track conditions [25], all of which influence thermal and structural degradation [26]. The standard longevity for C1 hard tyres typically ranges between 30 and 60 laps, as they are optimal for long stints on high-degradation tracks, while the C2 tyre longevity ranges between 25 and 50 laps. The medium compound, C3, lasts for 20 to 40 laps and is considered the most versatile. Both softer compounds possess a smaller lifespan, with the C4 ranging between 15 and 30 laps and the C5 between 10 and 20 laps [27, 28].

However, Pirelli's allocation utilises a sliding scale of five dry-weather compounds, ranging from C1 (hardest) to C5 (softest). To parameterise the model, the lifespans of C1, C2, and C3 were established by calculating the midpoint of each compound's durability range; C4 was assigned a lifespan exceeding the midpoint, and the C5 lifespan was set at the upper boundary. Consequently, the specific lifespans were established as shown in Table 5.2.

Additionally, prior to each race, Pirelli announces which three adjacent grades will be utilised. Based on these track-specific allocations, showned in Table 5.3, corresponding to the c-grade allocation for 2024 season [29], the model automatically selects the corresponding lifespan limits and attributes them to the designated hard, medium, and soft compound for that specific race.

Regarding the wear degradation factor ($\hat{\beta}^{\omega}_{t1,i}$), it is theoretically established that this value should not remain constant throughout a race. Factors such as heavy fuel loads and low-grip "green" track conditions contribute to higher tyre degradation during the opening segments of a race. Consequently, the multiplier M_i, utilised to adjust the base wear degradation $\beta^{\omega}_{t1,i}$, is expected to decrease as the race

Table 5.2: Specific lifespans, in laps, per compound grade

C-Grade	Lifespan
C1	45
C2	38
C3	30
C4	25
C5	20

Table 5.3: Track-Specific Pirelli Compound Allocations

Race	Hard	Medium	Soft
Bahrain Grand Prix	1	2	3
Saudi Arabian Grand Prix	2	3	4
Australian Grand Prix	3	4	5
Emilia Romagna Grand Prix	3	4	5
Miami Grand Prix	2	3	4
Spanish Grand Prix	1	2	3
Monaco Grand Prix	3	4	5
Azerbaijan Grand Prix	3	4	5
Canadian Grand Prix	3	4	5
British Grand Prix	1	2	3
Austrian Grand Prix	3	4	5
French Grand Prix	2	3	4
Hungarian Grand Prix	3	4	5
Belgian Grand Prix	2	3	4
Dutch Grand Prix	1	2	3
Italian Grand Prix	3	4	5
Singapore Grand Prix	3	4	5
Japanese Grand Prix	1	2	3
United States Grand Prix	2	3	4
Mexico City Grand Prix	3	4	5
São Paulo Grand Prix	3	4	5
Abu Dhabi Grand Prix	3	4	5
Qatar Grand Prix	1	2	3
Las Vegas Grand Prix	3	4	5
Chinese Grand Prix	2	3	4

progresses, reflecting the transition to depleted fuel loads and rubbered-in track surfaces. While an ideal profile would assign $M_i > 1$ for initial stints and $M_i < 1$ for final stints to reflect these physical improvements, empirical data to accurately calibrate these fluctuations was not available for this study. Therefore, a constant degradation rate was assumed, and the multiplier was set to $M_i = 1$ for all stints. It must be noted that, that under this simplification, the model does not truly capture the strategic value of stint order, because the degradation effect is effectively neutralised, and so different sequences of the same tyre compound (per example `[S-M-M]`, `[M-S-M]`, `[M-M-S]`) become mathematically equivalent, producing the same total race time. Nevertheless, the model architecture is explicitly designed to accept distinct M_i values, which allows for future integration of dynamic track evolution data, when such information becomes available. Consequently, under the assumption of $M_i = 1$, the adjusted degradation factor $\hat{\beta}^{\omega}_{t1,i}$ becomes equivalent to the base degradation parameter β^{ω}_{t1}.

The performance coefficients used in this model, β_{t0} and β_{t1}, were derived from the research of García [8]. To estimate these parameters, the study utilised multiple statistical approaches, including regression analysis, the Interquartile Method, and the Confidence Region method. This research adopted the values derived from the Interquartile Method, as it removed outliers and mitigated the influence of non-observable variables that typically led to prediction errors, while simultaneously retaining a larger sample size than the Confidence Region method. Therefore, by focusing on data distribution and percentiles, this approach ensured a robust foundation for modelling the relationship between tyre compounds and lap times.

Therefore, the parameter β_{t0}^{ω} represents the pace penalty of a tyre compound when it is fresh for the scenario ω; it quantifies the inherent speed difference between tyre compounds under identical conditions and reflects the performance trade-off between the durability and the speed of different compounds. While β_{t0}^{ω} is treated as a baseline coefficient, it is acknowledged that this value can be influenced by varying track surface textures and ambient temperatures, which affect the initial grip levels of the tyre.

Additionally, the parameter β_{t1}^{ω} denotes the wear acceleration factor, which dictates the increase in lap-time degradation as a tyre completes more laps, in every scenario. This model maintains β_{t1}^{ω} as the base wear acceleration factor for each compound, serving as the core metric for how quickly a tyre loses performance over time.

Thus, for the deterministic version of the model, both parameter values are applied as fixed constants sourced directly from the literature, while in the stochastic version they are treated as random variables. In the stochastic version, the intervals are defined based on the standard deviation (σ) reported in the literature. For the soft compound, the β_{t0} in both versions is fixed at zero as it serves as the baseline against the other compounds. The remaining β values for all compounds vary within $\pm 2 \cdot \sigma$ of their respective literature values. The exception is for the β_{t0}^{ω} of the hard compound, as it ranges from $-4 \cdot \sigma$ up to the literature value, reflecting the conclusions drawn from the deterministic results analysis. The specific values utilised for both model versions are summarised in Table 5.4.

Table 5.4: Performance and degradation coefficients for deterministic and stochastic model versions.

Compound	Deterministic		Standard Deviation		Stochastic	
	β_{t0}	β_{t1}	σ_{t0}	σ_{t1}	β_{t0}^{ω}	β_{t1}^{ω}
Soft	0.000	0.146	-	0.008	0.000	[0.130, 0.162]
Medium	0.544	0.061	0.099	0.003	[0.346, 0.742]	[0.055, 0.067]
Hard	1.511	0.031	0.116	0.004	[1.047, 1.511]	[0.023, 0.039]

The pit stop duration is a critical variable in race strategy optimisation, representing the time penalty incurred for changing tyres during a race. The data for both versions of the model was based on empirical historical records, to ensure the model reflects standard racing conditions, where outliers were removed from the analysis to prevent non-representative events from skewing the results. The outliers considered were pit stops exceeding 40 seconds.

In the deterministic version, the pit loss is treated as a fixed value for each specific race, which is calculated as the arithmetic mean of all valid pit stop durations recorded at that circuit across the analysed seasons. This provides a stable “average” baseline that represents the expected time cost for a standard pit stop at each individual track.

On the other hand, in the stochastic version, the pit stop duration (P_{pit}) is modelled as a triangular distribution defined by the parameters of minimum, maximum, and mode, as shown in Figure 5.1. This histogram aggregates data across all races within the analysed period, to ensure model consistency of this parameter. This distribution allows the model to account for the inherent variability of race-day conditions. Initially, it was hypothesised that pit lane traffic density might lead to a sequential decay factor for subsequent pit stops; however, given the lack of empirical data to quantify the impact of field dispersion on pit stop times, the model assumes a constant mode for all pit stops within a single race. This conservative approach avoids the introduction of arbitrary performance gains that cannot be validated. Furthermore, it is important to note that this assumption reflects clean racing conditions, as the model does not currently account for external disruptions such as safety car periods, virtual safety cars, or weather-induced stops, which would likely increase traffic density and contradict the expected reduction in pit stop times.

Therefore, the values for the pit stops for each race are defined in Table 5.5:

Table 5.5: pit stop duration parameters for deterministic and stochastic models.

Race	Deterministic P_{pit} (s)	Stochastic P_{pit} [min, max, mode] (s)
Bahrain Grand Prix	25.771	
Saudi Arabian Grand Prix	21.989	
Australian Grand Prix	18.877	
Emilia Romagna Grand Prix	30.964	
Miami Grand Prix	22.037	
Spanish Grand Prix	22.910	
Monaco Grand Prix	26.370	
Azerbaijan Grand Prix	21.620	
Canadian Grand Prix	24.986	
British Grand Prix	29.507	
Austrian Grand Prix	22.470	
French Grand Prix	37.085	
Hungarian Grand Prix	22.437	(13.973, 39.879, 21.814)
Belgian Grand Prix	23.548	
Dutch Grand Prix	21.269	
Italian Grand Prix	25.284	
Singapore Grand Prix	30.656	
Japanese Grand Prix	24.461	
United States Grand Prix	25.217	
Mexico City Grand Prix	23.341	
São Paulo Grand Prix	24.966	
Abu Dhabi Grand Prix	22.785	
Qatar Grand Prix	29.053	
Las Vegas Grand Prix	22.335	
Chinese Grand Prix	23.743	

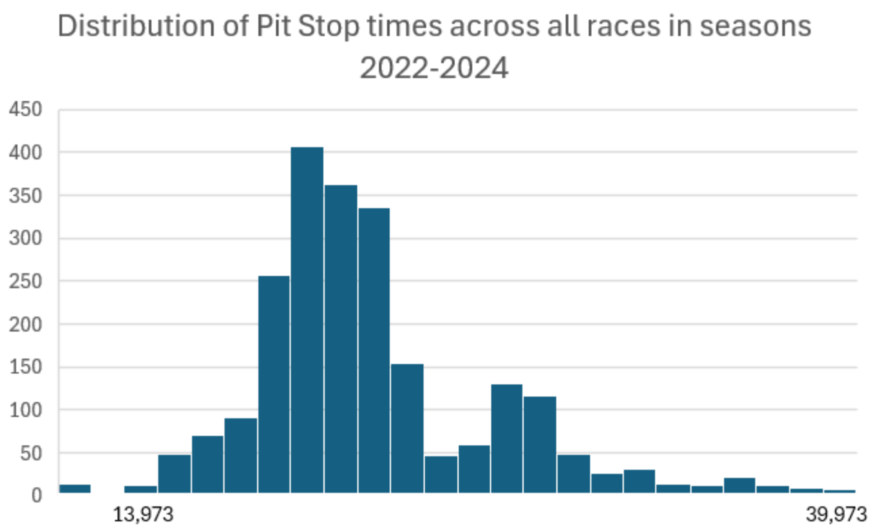

Figure 5.1: Probability density function of the triangular distribution for pit lane time loss, based on empirical historical data

Chapter 6

Results Analysis

The objective of this chapter is to evaluate the performance and optimal strategies given by the optimisation model developed. The analysis is structured into two phases.

The first phase focuses on the deterministic model, where all parameters are treated as fixed, known constants. This phase is essential for establishing a theoretical baseline. A key variable throughout the deterministic analysis is the `max stint` constraint. By varying this constraint, the analysis evaluates the sensitivity of the optimal strategy to changes in stint limits, identifying both the stable strategic core of the races and the thresholds where additional pit stops yield only marginal time gain. Although the deterministic formulation solves each individual optimisation problem within a relatively short computational time, evaluating all 25 circuits across the full range of `max stint` values (up to the total tyre inventory of nine stints) resulted in an overall execution time of approximately 20 minutes.

The second phase transitions to the stochastic model, where the key performance parameters are treated as random variables sampled from their empirical distribution across 20 discrete scenarios. This phase evaluates whether the optimal strategies identified under deterministic conditions remain robust when parameter uncertainty is introduced, and quantifies how the distribution of race outcomes shifts as a consequence. This phase is further validated against Pirelli's official pre-race strategy recommendations and the strategies executed by actual 2024 race winners, and is additionally stress-tested through a Monte Carlo simulation framework evaluating robustness under yellow flag disruptions. Compared with the deterministic formulation, the stochastic model is computationally significantly more demanding because it incorporates uncertainty in the model parameters by considering multiple scenarios, thereby increasing the problem size and computational complexity. Even with only 20 scenarios, each complete set of simulations across the 25 circuits required more than one hour of computation. Consequently, the analysis was limited to 20 scenarios, as preliminary attempts to increase the number of scenarios to 50 resulted in computational times that were considered impractical within the available computational resources.

6.1 Deterministic Model Analysis

6.1.1 Baseline Model Performance

In the initial phase of the results analysis, the deterministic optimisation model was run using the baseline parameters derived from the literature. The primary objective of this baseline run was to assess the model's behaviour across all 25 Grand Prix tracks under varying constraints of maximum allowed stints, from four to nine.

As shown in Table 6.1, which presents results only for the three and four stint range, the model demonstrated a high degree of stability, as the maximum stint constraint increased from four to nine, and the optimal strategies across the entire calendar remained consistent, obtaining the same total race time. Consequently, it reflected the model's preference for the most efficient tyre combination regardless of the additional flexibility.

Table 6.1: Model results for 25 Grand Prix tracks for 3 and 4 maximum stints.

Race	N (laps)	t_{base} (s)	P_{pit} (s)	`max stints 3`					`max stints 4`				
				Type	t Tyre Sequence	Pit Laps	Stint Lengths	Total Time (s)	Type	t Tyre Sequence	Pit Laps	Stint Lengths	Total Time (s)
Bahrain GP	57	89.179	25.771	1-Stop	`M-S`	[37]	37, 20	5197.468	1-Stop	`S-M`	[20]	20, 37	5197.468
Saudi Arabian GP	50	87.472	21.989	1-Stop	`S-M`	[20]	20, 30	4466.184	1-Stop	`S-M`	[20]	20, 30	4466.184
Australian GP	58	75.915	18.877	2-Stop	`M-M-S`	[22, 45]	22, 23, 13	4506.216	2-Stop	`M-S-M`	[23, 36]	23, 13, 22	4506.216
Emilia Romagna GP	63	74.746	30.964	2-Stop	`M-M-S`	[25, 49]	25, 24, 14	4846.004	2-Stop	`M-S-M`	[24, 38]	24, 14, 25	4846.004
Miami GP	57	86.841	22.037	2-Stop	`M-M-S`	[22, 44]	22, 22, 13	5057.516	2-Stop	`S-M-M`	[13, 35]	13, 22, 22	5057.516
Spanish GP	66	71.383	22.910	2-Stop	`S-M-M`	[15, 40]	15, 25, 26	4838.296	2-Stop	`M-M-S`	[26, 51]	26, 25, 15	4838.296
Monaco GP	78	70.27	26.370	2-Stop	`M-H-M`	[25, 53]	25, 28, 25	5651.627	3-Stop	`M-M-M-S`	[22, 44, 65]	22, 22, 21, 13	5647.911
Azerbaijan GP	51	100.203	21.620	2-Stop	`M-S-M`	[20, 32]	20, 12, 19	5206.466	2-Stop	`S-M-M`	[12, 31]	12, 19, 20	5206.466
Canadian GP	70	72.0	24.986	2-Stop	`S-M-M`	[20, 45]	20, 25, 25	5181.513	2-Stop	`M-M-S`	[25, 50]	25, 25, 20	5181.513
British GP	52	85.819	29.507	1-Stop	`M-S`	[34]	34, 18	4567.15	1-Stop	`M-S`	[34]	34, 18	4567.15
Austrian GP	71	64.314	22.470	2-Stop	`H-M-M`	[21, 46]	21, 25, 25	4713.275	3-Stop	`S-M-M-M`	[12, 32, 52]	12, 20, 20, 19	4709.046
French GP	53	90.872	37.085	1-Stop	`S-M`	[23]	23, 30	4933.094	1-Stop	`S-M`	[23]	23, 30	4933.094
Hungarian GP	70	75.227	22.437	2-Stop	`S-M-M`	[20, 45]	20, 25, 25	5402.304	2-Stop	`M-S-M`	[25, 45]	25, 20, 25	5402.304
Belgian GP	44	103.665	23.548	1-Stop	`S-M`	[16]	16, 28	4640.618	1-Stop	`S-M`	[16]	16, 28	4640.618
Dutch GP	72	69.673	21.269	2-Stop	`M-M-S`	[28, 56]	28, 28, 16	5153.094	2-Stop	`S-M-M`	[16, 44]	16, 28, 28	5153.094
Italian GP	53	79.327	25.284	2-Stop	`S-M-M`	[13, 33]	13, 20, 20	4311.227	2-Stop	`M-M-S`	[20, 40]	20, 20, 13	4311.227
Singapore GP	62	89.525	30.656	2-Stop	`M-S-M`	[24, 38]	24, 14, 24	5684.933	2-Stop	`S-M-M`	[14, 38]	14, 24, 24	5684.933
Japanese GP	53	88.197	24.460	1-Stop	`S-M`	[18]	18, 35	4776.574	1-Stop	`S-M`	[18]	18, 35	4776.574
United States GP	56	92.33	25.217	2-Stop	`M-S-M`	[22, 35]	22, 13, 21	5282.596	2-Stop	`M-M-S`	[21, 43]	21, 22, 13	5282.596
Mexico City GP	71	75.946	23.341	2-Stop	`H-M-M`	[21, 46]	21, 25, 25	5540.888	3-Stop	`M-M-M-S`	[19, 39, 59]	19, 20, 20, 12	5537.531
São Paulo GP	71	70.727	24.966	2-Stop	`M-M-H`	[25, 50]	25, 25, 21	5173.59	3-Stop	`S-M-M-M`	[12, 32, 51]	12, 20, 19, 20	5171.859
Abu Dhabi GP	58	82.595	22.785	2-Stop	`M-M-S`	[23, 45]	23, 22, 13	4901.472	2-Stop	`M-S-M`	[23, 36]	23, 13, 22	4901.472
Qatar GP	57	80.52	29.053	1-Stop	`S-M`	[20]	20, 37	4707.187	1-Stop	`S-M`	[20]	20, 37	4707.187
Las Vegas GP	50	92.312	22.334	2-Stop	`S-M-M`	[12, 31]	12, 19, 19	4711.439	2-Stop	`S-M-M`	[12, 31]	12, 19, 19	4711.439
Chinese GP	56	93.66	23.743	2-Stop	`S-M-M`	[13, 35]	13, 22, 21	5354.127	2-Stop	`M-M-S`	[22, 43]	22, 21, 13	5354.127

An analysis of the baseline results reveals that, for the vast majority, with the exception of four specific circuits, altering the maximum stint limit does not change the optimal race time. While the total race time remains constant, the specific tyre sequences and their corresponding pit windows frequently alternate.

For example, in the Australian Grand Prix, sequences such as `M-M-S`, `M-S-M`, and eventually `S-M-M` were generated as the limits expanded, all yielding the exact same stint length for each compound and race duration. This variation is a direct mathematical consequence of the stint multiplier being held constant at one unit ($M_i = 1$) across all stints. Because the degradation rate for a specific compound remained identical regardless of whether it was fitted in the first or final stint, the optimisation algorithm treated the stint lengths and tyre choices as mathematically commutative. Consequently, multiple optimal pathways could be generated without incurring any variation in the final objective function value.

A further examination of the tyre sequences revealed that the medium compound was the core of the strategic output, as shown in Table 6.2, serving as the most versatile option across the widest range of circuits. The soft compound was selected in instances where a short, high-performance sprint was required, which is consistent with its real-world application as an aggressive tactical tool. However, the baseline values meant that the hard compound was never selected for maximum stint lengths of four or greater. It was only under the stricter constraint of a maximum of three stints that the strategy fundamentally changed. In four specific circuits: the Monaco, Austrian, Mexico City, and São Paulo Grands Prix, the model was restricted from performing a more aggressive three-stop strategy. In these instances, the approach shifted from an aggressive short-stint strategy to a more conservative, long-stint strategy, utilising the durability of the hard compound to compensate for the pit stop limitation.

Table 6.2: Frequency of tyre compound selection across varying maximum stint constraints.

`Max Stints`	Soft	Medium	Hard
3	21	43	4
4	25	47	0
5	25	47	0
6	25	47	0
7	25	47	0
8	25	47	0
9	25	47	0

The exclusion of the hard tyre can be explained mathematically by analysing the performance offset β_{t0} and the linear degradation rate β_{t1}. The model was operated under the objective of minimising the total race time by evaluating the breakeven point, which represents the lap at which the cumulative time loss caused by the degradation of a faster but less durable tyre exceeds the time cost associated with an additional pit stop. In most simulated races, the advantage of the soft compound was offset by its higher degradation rate, whereas the hard compound was utilised only when the stint count constraint prevented the optimisation algorithm from adopting a multi-stop strategy.

The consistent rejection of the hard compound in unconstrained stints can be explained by the conflicting interaction between the performance offset and the degradation rate obtained from the literature. To understand why the optimisation algorithm systematically rejected the hard tyre, the lap-time evolution within each stint (t) was evaluated under the baseline parameters.

As illustrated in Table 6.3, the total stint time degradation was calculated by subtracting the lap time of the first lap from that of the final lap of each stint for the three- and four-stint limit scenarios. This analysis allowed the evaluation of how much performance degradation the algorithm permitted before initiating a pit stop for a tyre change. Only a subset of race results is presented in the table as an excerpt of the complete dataset, as including all races was deemed unnecessary.

The data reveals that, for the vast majority of the races, the total time lost to degradation (Δt) in soft and medium stints ranged between one and 2.5 seconds. Conversely, the baseline initial pace assigned to the hard compound was 1.511 seconds. Therefore, a brand-new hard tyre on its first lap out of the pits was already significantly slower than a degraded tyre at the absolute end of its lifespan, as the performance gap remained in the order of approximately one second.

Table 6.3: Excerpt of total stint time degradation (Δt) for selected races under 3 and 4 maximum stint constraints.

Race	`max stints 3`			`max stints 4`		
	Stint	Tyre	Δt (s)	Stint	tyre	Δt (s)
Australian Grand Prix	1	Medium	1.281	1	Medium	1.342
	2	Medium	1.342	2	Soft	1.752
	3	Soft	1.752	3	Medium	1.281
Monaco Grand Prix	1	Medium	1.464	1	Medium	1.281
	2	Hard	0.837	2	Medium	1.281
	3	Medium	1.464	3	Medium	1.220
	N/A	N/A	N/A	4	Soft	1.752
Japanese Grand Prix	1	Soft	2.482	1	Soft	2.482
	2	Medium	2.074	2	Medium	2.074
Mexico City Grand Prix	1	Hard	0.620	1	Medium	1.098
	2	Medium	1.464	2	Medium	1.159
	3	Medium	1.464	3	Medium	1.159
	N/A	N/A	N/A	4	Soft	1.606

This paradox provided mechanical clarity on the algorithm's behaviour, particularly in races like the Monaco and Mexico City Grand Prix. When constrained to a maximum of three stints, the algorithm was mathematically forced to fit the hard compound simply to ensure the car could reach the end of the race without an additional pit stop, respecting the limitation. In these forced scenarios, the hard tyre demonstrated its inherent durability, degrading by only 0.837 seconds in Monaco and 0.62 seconds in Mexico. However, the moment the constraint was lifted, the hard tyre was immediately abandoned entirely. It was calculated that running extra stints of mediums and softs (even though they degraded by one to 2.5 seconds) was a superior mathematical trade-off, because avoiding the permanent 1.511 second pace deficit of the hard tyre saved more total race time.

Furthermore, in high-degradation scenarios where Δt exceeded the hard tyre's penalty, such as in the Japanese Grand Prix (where the soft and medium degraded by 2.482 and 2.074 seconds, respectively), the hard tyre was still rejected. In these extreme cases, it was calculated that the marginal time the hard tyre might gain at the very end of a long stint was not enough to overcome the cumulative time lost during early laps, nor did it justify a second pit lane penalty required to fit it if an alternative multi-stop strategy was available.

Overall, the exclusion of the hard compound can be explained. The literature-derived β_{t0} value being disproportionately severe, penalising the hard tyre so heavily on pace that its degradation advantage becomes mathematically irrelevant.

In modern Formula 1 racing, the hard compound is a vital strategic tool. It is often used in situations where durability and consistency over long distances—such as at the beginning of the race when the car carries a full fuel load—provide greater stability and tyre management. Furthermore, it is preferred on highly abrasive tracks or during long race stints, where minimising pit stops becomes crucial as overtaking opportunities typically arise. Therefore, a sensitivity analysis was conducted to analyse the parameter behaviours and confirm that they must be recalibrated to reflect a realistic Formula 1 environment.

6.1.2 Sensitivity Analysis

The baseline results demonstrated that the hard compound was never selected, unless constrained by the stint limit boundary condition, where the soft tyre had to be replaced by the hard compound due to its greater durability in order to complete the total race distance of the Formula 1 circuit. While this model output is mathematically consistent, it does not fully capture the strategic diversity observed in modern Formula 1, where the hard compound is frequently employed as a core element of competitive race strategies.

A critical factor in this observation is the origin of the performance parameters β_{t0} and β_{t1}. These values were derived from existing literature that utilised the same underlying Kaggle dataset. However, in those studies, the parameters were adapted specifically to align with the objectives and mathematical constraints of that research. Consequently, these literature-derived values are context-specific; they do not necessarily represent a generalised global truth of tyre behaviour, but rather a specifically optimised configuration. As a result, applying these exact parameters to the present model introduces a potential calibration bias, where the algorithm inherits the specific strategic preferences of the prior study rather than reflecting the broader tactical realities of Formula 1.

To bridge this gap and establish a more realistic baseline, a systematic sensitivity analysis was conducted. The objective of this phase was twofold: first, to quantify the model's responsiveness to variations in tyre characteristics; and second, to identify a parameter configuration that produces a balanced tyre usage profile. By independently perturbing the degradation and performance parameters, a set of coefficients was sought to allow the optimisation model to move beyond boundary conditions and replicate the compound diversity seen in actual Grand Prix weekends.

Sensitivity of Degradation Rates (β_{t1})

Tyre longevity is the primary tactical lever in Formula 1. The initial parameter set derived from the literature appeared to favour the medium compound and an aggressive medium-soft strategy to an extent that suppressed the hard tyre's utility. In this phase, the linear degradation rate was adjusted to test whether the optimisation algorithm could be steered towards a more balanced utilisation of the hard compound without compromising the model's overall logic. Additionally, the sensitivity to the medium compound was tested to determine whether the medium tyre's dominance was an inherent structural preference of the model or a direct consequence of its specific degradation rate.

Therefore, to assess the sensitivity to this parameter, the degradation rates (β_{t1}) of the compounds were examined under the following cases:

- Hard compound durability: A reduction of the hard tyre's β_{t1} by 20%, 30%, and 50% was implemented. This aimed to lower the stint cost of the hard tyre, determining at what point its increased longevity compensated for its lower raw pace.

- Soft compound aggression: A reduction of the soft tyre's β_{t1} by 20%, alongside increases of 20% and 30%, were applied. This was done to assess the model's sensitivity to the soft tyre's performance decay, identifying whether the baseline parameters were overly punitive for short-stint scenarios.
- Medium compound neutralisation: A reduction and increase of the medium tyre's β_{t1} by 20% were performed, alongside calculating the midpoint degradation value between the hard and the soft compounds, which resulted in a value of 8.9%.

Table 6.4: Compound selection frequency under perturbed degradation parameters (β_{t1}). From top to bottom, the three panels correspond to hard, soft, and medium compounds, respectively.

	20% Reduction			30% Reduction			50% Reduction		
max stint	S	M	H	S	M	H	S	M	H
3	21	43	4	21	43	4	17	41	8
4-9	25	47	0	24	46	1	18	42	7

	20% Reduction			20% Increase			30% Increase		
max stint	S	M	H	S	M	H	S	M	H
3	30	34	4	20	45	5	19	46	6
4-9	42	30	0	24	49	1	22	50	3

	20% Increase			Midpoint (0.089)			20% Reduction		
max stint	S	M	H	S	M	H	S	M	H
3	21	44	4	36	32	4	21	43	4
4-9	25	48	0	49	29	0	25	47	0

As presented in Table 6.4, the sensitivity analysis results demonstrated a high degree of strategic stability within the model. Across the majority of the perturbed parameters, the preference for the medium compound was maintained, confirming it as the strategic anchor.

The most notable observation was the resistance to changing the composition of the race strategy when degradation was modified within a range of 30%. However, in the test case where the hard tyre degradation rate was reduced by half (rendering it 50% more durable), the adoption of this compound remained marginal, increasing by only four or seven units, depending on the stint limitation. This indicated that the hard tyre's current performance deficit was too significant for its longevity alone to compensate, meaning that it continued to be categorised as a conservative boundary solution rather than a competitive pacesetter.

Moreover, the midpoint test of the medium tyre, which represented an increase of approximately 45% in degradation, revealed a threshold in the decision-making process. When the degradation efficiency of the medium tyre was neutralised, the strategy pivoted to a soft-compound dominance. This confirmed that the medium compound's dominance was not an arbitrary preference, but a direct outcome of its superior cost-per-lap efficiency. Additionally, when that efficiency was degraded, the hard tyre was not reverted to; rather, the raw performance of the soft compound was prioritised, accepting higher degradation in exchange for maintaining a competitive race time.

Therefore, these results provide two key insights for interpretation:

- Strategic robustness: The model is not prone to volatility, as its preference for the medium tyre is a logical response to the established degradation profile.
- Degradation thresholds: Modifications to the degradation rates alone are insufficient to elevate the hard compound to a standard strategy; hence, the hierarchy of preference remains the medium compound, followed by the soft, and lastly the hard.

Sensitivity of Performance Penalty (β_{t0})

While the degradation rate (β_{t1}) dictated the duration of a stint, the performance offset (β_{t0}) determined the value proposition of each compound at the start of its life. If a compound is perceived as too slow, no amount of longevity renders it optimal within a time minimisation framework. Therefore, this phase tested whether adjusting the performance gap between compounds provided the necessary incentive for the model to diversify its tyre selection. The following tests were performed:

- Standardised performance delta: The performance of the hard compound was adjusted to 1.044 (reduction of approximately 31%) to create an equidistant performance delta between the soft, medium, and hard tyres, removing potential biases in the raw speed hierarchy derived from the original literature.
- Systematic performance scaling: A progressive reduction of the β_{t0} parameter of the hard compound by 10%, 15%, 20%, and 25% was implemented to test how much raw performance the hard tyre must gain to transition from a conservative tool to a competitive primary alternative for a standard two-stop strategy.

Table 6.5: Compound selection frequency under perturbed performance offset parameters (β_{t0}) for the hard compound.

	10% Reduction			15% Reduction			20% Reduction			25% Reduction			1.044 β_{t0}		
max stint	S	M	H	S	M	H	S	M	H	S	M	H	S	M	H
3	21	43	4	15	40	10	12	38	13	8	36	17	5	35	20
4-9	23	45	2	15	40	10	12	38	13	8	36	17	5	35	20

The results of the sensitivity analysis regarding the hard compound's performance offset (β_{t0}), presented in Table 6.5, differed significantly from the β_{t1} degradation tests. While the model was highly resistant to changes in tyre longevity, it exhibited considerable sensitivity to adjustments in the initial grip offset.

The data revealed a clear relationship between the hard compound's pace deficit and its strategic viability. In the baseline scenario, the hard compound's performance penalty ($\beta_{t0} = 1.511$) rendered it practically non-viable, except in situations where three-stint constraints forced its use. A minor 10% reduction in this parameter yielded negligible changes, increasing hard compound usage to only two instances. However, the critical threshold of the optimisation function occurred between the 10% and 15% reduction intervals. At a 15% reduction, the usage of the hard compound increased to 10 stints, which demonstrated that it had become mathematically competitive against the medium tyre in unconstrained scenarios, as the resulting strategies became identical across different stint limit constraints.

As the performance penalty was further reduced—until it reached a value of 1.044 (an approximately 31% reduction)—the adoption of the hard compound increased, while a significant decline in soft compound usage was observed. This indicated that when the hard tyre possessed adequate pace, it was leveraged by the model to replace multi-stop sprint strategies with more conservative alternatives.

Selection of the Reference Model

The objective of the sensitivity analysis was to identify a calibrated parameter set that accurately reflects the strategic realities of contemporary Formula 1. To ensure the selected parameters were grounded in empirical evidence, the model's outputs were benchmarked against the actual tyre compound distribution from the 2022 to 2024 Formula 1 seasons, extracted from the Kaggle dataset:

Table 6.6: Empirical tyre compound distribution from the 2022 to 2024 Formula 1 seasons.

Compound	Total Stints	Percentage (%)
Hard	1336	39.43
Medium	1411	41.65
Soft	641	18.92

Comparing this empirical baseline to the model's uncalibrated output revealed the inaccuracy of the initial parameter distribution. The baseline model produced a highly unrealistic hierarchy of `Medium > Soft >> Hard`, almost entirely excluding the hard compound. However, the β_{t0} sensitivity analysis successfully corrected this discrepancy.

When evaluating the adjusted scenarios against the empirical data, the 25% reduction of the hard compound's performance decay emerged as the most realistic configuration. Under this calibrated scenario, the optimisation model outputs a strategy distribution of 36 medium stints, 17 hard stints, and eight soft stints across the season. Consequently, this output restores the real-world empirical hierarchy of `Medium > Hard > Soft`. Furthermore, it accurately replicates the specific usage ratio observed in the Kaggle dataset, where the hard compound usage is approximately twice that of the soft compound.

By neutralising the artificial pace penalty derived from the original literature, the optimisation model is no longer forced to avoid the hard compound. Therefore, the scenario with the 25% reduction provides the highest degree of fidelity to historical Formula 1 strategic behaviour and was thus selected as the reference model. Consequently, this reference model is utilised for all subsequent analyses, with the values presented in the Table 6.7.

Table 6.7: Beta values of the new reference model

Compound	β_{t0}	β_{t1}
Soft	0.000	0.146
Medium	0.544	0.061
Hard	1.133	0.031

6.1.3 Analysis of Optimal Strategies Under the Reference Model

The implementation of the calibrated reference model, characterised by the 25% reduction in the hard compound's pace penalty, permits a meaningful categorisation of race strategies across the Formula 1 calendar. The optimisation results are identical along the maximum stint constraint, and reveal that 14 races out of 25, resulting in 56%, have only one pit stop, while the other 44%, 11 races, opt for two stops. Within the one-stop cluster, the model shows a strong preference for the `[H-M]` (or `[M-H]`) sequences, which are selected in 13 of 14 cases, with only the Belgian Grand Prix having a `[S-M]` sequence. This distribution indicates that the model is no longer relying on mathematical boundary conditions to force strategy, but is instead autonomously identifying the optimal balance between tyre durability and pace based on the input parameters.

The strategic divergence is not arbitrary; rather, it emerges from the model's evolution of the trade-off between tyre degradation and the time loss associated with the pit lane. Three factors influence the divergence between one-stop and two-stop strategies within the same parameters: first the influence of lap duration, the effect of total race distance, and the role of the pit lane time loss. Through these correlations, it is possible to demonstrate that the model's strategic output is a direct reflection of the physical energy demands and structural constraints inherent to each circuit.

The Role of Lap Duration

The base lap time (t_{base}) is a key structural factor in determining pit stop strategies. Since t_{base} is the minimum time needed to complete a lap under optimal conditions, it serves as a proxy for the circuit's length and layout complexity. Because all races cover a fixed distance, circuits with higher t_{base} tend to be physically longer or characterised by high-speed sections that sustain elevated average velocities, while circuits with lower t_{base} tend to be shorter and more technical. Therefore, this relationship between lap duration and total race laps directly affects how tyre degradation influences the optimisation model.

Tyre degradation within a stint (Δt_{base}) is independent of t_{base}, as a pit stop is only strategically viable if the time gained from fresh tyres exceeds the time lost in the pit lane. On circuits with a high t_{base}, each lap takes a significant amount of time, meaning that degradation per lap represents a smaller percentage of the total race time. This reduces the incentive to pit for fresh tyres, which often makes a one-stop strategy more effective.

Tables 6.8 and 6.9 illustrate this effect across the 25 race calendar. Among the 14 circuits classified in the one-stop strategy, the average base lap time is approximately 88 seconds (ranging from 79.327s at Monza to 103.665s at Spa-Francorchamps). In these circuits, the hard compound stints consistently produce low Δt values despite being run for longer durations. For instance, at the Japanese Grand Prix, the hard tyres run for 29 laps and accumulate only 0.868 seconds of degradation, while at the United States Grand Prix, the hard stint covers 31 laps with 0.93 seconds of degradation. These values are substantially below the pit stop loss at their respective circuits (24.460s and 25.217s), meaning that no second stop can mathematically recover the time invested in the pit lane. The medium compound stints in these same races follow a similar pattern, with Δt values clustering in the 1.342 to 1.780 second range, remaining well within the threshold required to justify an additional stop.

Table 6.8: Stint degradation (Δt) for one-stop races under the reference model (hard β_{t0} reduced by 25%).

				Stint 1			Stint 2		
Race	N	t_{base} (s)	P_{pit} (s)	tyre	Length (l)	Δt (s)	Tyre	Length (l)	Δt (s)
Bahrain GP	57	89.179	25.771	Medium (C2)	26	1.525	Hard (C1)	31	0.930
Saudi Arabian GP	50	87.472	21.989	Medium (C3)	23	1.342	Hard (C2)	27	0.806
Miami GP	57	86.841	22.037	Hard (C2)	31	0.930	Medium (C3)	26	1.525
Spanish GP	66	71.383	22.910	Medium (C2)	29	1.708	Hard (C1)	37	1.116
Azerbaijan GP	51	100.203	21.620	Medium (C4)	24	1.403	Hard (C3)	27	0.806
British GP	52	85.819	29.507	Hard (C1)	28	0.837	Medium (C2)	24	1.403
French GP	53	90.872	37.085	Medium (C3)	24	1.403	Hard (C2)	29	0.868
Belgian GP	44	103.665	23.548	Soft (C4)	16	2.190	Medium (C3)	28	1.647
Italian GP	53	79.327	25.284	Medium (C4)	24	1.403	Hard (C3)	29	0.868
Japanese GP	53	88.197	24.461	Hard (C1)	29	0.868	Medium (C2)	24	1.403
United States GP	56	92.330	25.217	Medium (C3)	25	1.464	Hard (C2)	31	0.930
Qatar GP	57	80.520	29.053	Medium (C2)	26	1.525	Hard (C1)	31	0.930
Las Vegas GP	50	92.312	22.335	Hard (C3)	27	0.806	Medium (C4)	23	1.342
Chinese GP	56	93.660	23.743	Medium (C3)	25	1.464	Hard (C2)	31	0.930

Table 6.9: Stint degradation (Δt) for two-stop races under the reference model (hard β_{t0} reduced by 25%).

				Stint 1			Stint 2			Stint 3		
Race	N	t_{base} (s)	P_{pit} (s)	Tyre	Length (laps)	Δt (s)	Tyre	Length (laps)	Δt (s)	Tyre	Length (laps)	Δt (s)
Australian GP	58	75.915	18.877	Medium (C4)	22	1.281	Medium (C4)	23	1.342	Soft (C5)	13	1.752
Emilia Romagna GP	63	74.746	30.964	Soft (C5)	14	1.898	Medium (C4)	24	1.403	Medium (C4)	25	1.464
Monaco GP	78	70.270	26.370	Medium (C4)	24	1.403	Hard (C3)	29	0.868	Medium (C4)	25	1.464
Canadian GP	70	72.000	24.986	Medium (C4)	25	1.464	Medium (C4)	25	1.464	Soft (C5)	20	2.774
Austrian GP	71	64.314	22.470	Medium (C4)	23	1.342	Medium (C4)	23	1.342	Hard (C3)	25	0.744
Hungarian GP	70	75.227	22.437	Medium (C4)	25	1.464	Medium (C4)	25	1.464	Soft (C5)	20	2.774
Dutch GP	72	69.673	21.269	Medium (C2)	28	1.647	Soft (C3)	16	2.190	Medium (C2)	28	1.647
Singapore GP	62	89.525	30.656	Medium (C4)	24	1.403	Soft (C5)	14	1.898	Medium (C4)	24	1.403
Mexico City GP	71	75.946	23.341	Medium (C4)	23	1.342	Medium (C4)	23	1.342	Hard (C3)	25	0.744
São Paulo GP	71	70.727	24.966	Medium (C4)	23	1.342	Medium (C4)	23	1.342	Hard (C3)	25	0.744
Abu Dhabi GP	58	82.595	22.785	Medium (C4)	22	1.281	Soft (C5)	13	1.752	Medium (C4)	23	1.342

Conversely, among the 11 circuits in the two-stop strategy, nine exhibit a $t_{base} < 76$ seconds, where the laps are shorter, and require more laps, forcing drivers to run longer stints if they attempt a one-stop, which leads to higher cumulative degradation. At the Canadian Grand Prix, the soft compound covers 20 laps and reaches an accumulation of 2.774 seconds of degradation, while at the Dutch Grand Prix, the soft stint accumulates 2.190 seconds over 16 laps. In both cases, the degradation savings offered by a fresh compound outweigh the pit stop penalty, and the model correctly identifies a second stop as an optimal choice.

Moreover, it is important to note that t_{base} does not operate in isolation. The Belgian Grand Prix is an outlier within the one-stop strategy, as it is the only one with a different approach. It has the highest t_{base} (103.665 seconds), yet with only 44 laps, the Spa circuit requires a tyre sequence of [`S-M`] instead of [`M-H`], like the others, with Δt = 2.190s and 1.647s of the soft and medium compound, respectively. This is a direct consequence of the Pirelli compound allocation, which assigned `C3`, `C4`, and `C5` grades to this circuit, making the hard compound a `C3` with a lifespan of only 30 laps; this is insufficient to cover a long first stint across 44 laps without exceeding its durability threshold. The model, therefore, substitutes the soft compound for the opening sprint stint, confirming that the compound allocation constraint interacts with the lap duration effect and can locally override the dominant strategic pattern.

Overall, the analysis of t_{base} reveals that lap duration functions as a scaling mechanism for the relative cost of degradation. On long-lap circuits, the degradation magnitude remains small relative to the pit stop penalty, making one-stop strategies structurally favourable. On short-lap circuits, the compounding effect of more laps and longer required stints amplifies Δt to a level where a second stop becomes mathematically justified, even when the pit stop penalty is not negligible.

The Role of Race Distance

While the base lap time (t_{base}) determines the relative cost of degradation, the total number of laps (N) provides an essential, independent dimension of the race structure. In Formula 1, the total number of laps required to complete the race distance depends entirely on the circuit's length. Consequently, N is a reflection of individual lap length: tracks with shorter laps require more laps to cover the same distance, while tracks with longer laps require fewer.

Moreover, while the total lap count (N) and the base lap time (t_{base}) often correlate, like in the Belgian Grand Prix, which has the season's longest circuit with the highest t_{base} and a low lap count, this relationship is frequently disrupted by the specific characteristics of the track layout. Circuits are not defined by length alone; the mix between high-speed straights and technical corner sections can significantly change how these two variables relate to each other. A long circuit with long straights allows for higher average speed, which keeps the t_{base} low despite the distance. Conversely, a shorter circuit defined by slow-speed corners can lead to a surprisingly high t_{base}. This is demonstrated by the Canada-Spain comparison, where the t_{base} is nearly identical with values of 72.0 and 71.383 seconds, respectively, yet a four-lap difference requires opposite strategies.

As defined in Equation (4.10), the total time cost of a stint of k laps contains a degradation term proportional to $\hat{\beta}_{t1}^{\omega} \cdot \frac{k^2-k}{2}$, where the quadratic dependence on k is the critical element determining how N influences strategy selection. Consequently, as N grows and each stint must cover more laps to avoid an additional pit stop, the degradation penalty amplifies disproportionately, where the breakeven point (the time saved by fitting a fresh set of tyres exceeds the pit stop time loss) is therefore reached at shorter stint lengths as N increases, making additional pit stops mathematically more attractive on longer circuits.

Tables 6.8 and 6.9 show this effect across the circuits with the longest number of laps on the calendar. A clear threshold emerges at $N = 70$, where every circuit at or above this value falls into a two-stop race. Moreover, in the intermediate range between $N = 58$ and $N = 70$, the strategy becomes more specific, as this range illustrates how compound allocation acts as a third moderating factor in strategy definition. The Singapore GP has a two-stop strategy despite having a high t_{base} (89.525s), because the `C3-C4-C5` allocation provides limited durability margins, where the soft stint covers only 14 laps before reaching a Δt of 1.898 seconds, while the medium stints cannot absorb the remaining distance without wearing out too much. Conversely, the Spanish GP remains a one-stop race despite its relatively high lap count. This is due to its `C1-C2-C3` compound allocation, which provides the hard compound a 45-lap lifespan, sufficient to anchor the remaining 37 laps, hence a Δt of only 1.116 seconds, enough to offset the 22.9 second pit stop penalty. Therefore, the Canada-Spain comparison is not just about t_{base} or N, but also the interaction with compound allocation, where Canada has `C3-C4-C5` and Spain has `C1-C2-C3` allocation.

Overall, the total number of laps operates as a multiplier for tyre wear cost rather than a direct rule for deciding the best pit strategy. A high lap count limits the stint lengths available for a one-stop strategy and forces the driver to deal with the rapid growth of the $k^2 - k$ penalty, lowering the breakeven threshold for an additional stop. However, as the Spain-Canada-Singapore comparison demonstrates, N interacts with both t_{base} and the Pirelli compound allocation, and none of the three factors is individually sufficient to predict the optimal strategy. Hence, the decision remains whether the total degradation saved by using fresh tyres over the remaining laps exceeds the time lost by driving through the pit lane.

The Role of Pit Stop Time Loss

While t_{base} and N determine the magnitude of the tyre wear a strategy must absorb, the pit stop time loss (P_{pit}) defines the cost that must be taken to fix that wear. In this framework, the P_{pit} has a fixed time that enters the objective function and penalises every stint transition. This value is circuit-specific, and might reflect the physical length of the pit lane, the speed limit enforced by FIA, and the velocity at which the pit crew performs the stop. Hence, unlike t_{base} and N, which shape the degradation cost, the P_{pit} sets the bar for whether it is worth fixing the degradation of the car. A higher penalty raises the bar, making it harder to justify an extra stop, while a lower penalty makes it easier.

In the first place, the two extremes of P_{pit} on the calendar are the French GP and Australian GP, with the former at the highest with 37.1 seconds and the latter the lowest at 18.9 seconds, which is a gap of approximately 18 seconds. This gap means that a second stop at France requires nearly twice the time saving from fresh tyres to be mathematically worth it compared to Australia. Therefore, France is classified as a one-stop with an [`H-M`] sequence, while Australia is a two-stop with an [`M-M-S`] sequence, bearing in mind that other sequences could be achieved as long as for the same compounds the length is the same.

This effect of time loss in the pit can be isolated by comparing the Qatar and Australian Grand Prix, which have a similar total number of laps (57 and 58), close base lap time (80.5 and 75.9 seconds), yet obtain different strategies, because the pit lane penalty difference is approximately 10 seconds. In this case, both circuits have similar levels of degradation. However, in Australia the degradation across the three stints outweighs the penalty of an extra pit stop, while in Qatar the high pit stop cost makes a two-stop strategy unrecoverable, favouring a one-stop sequence. Therefore, while other structural factors are balanced, like t_{base} and N, the P_{pit} determines the type of strategy the driver will follow.

Moreover, in the British GP, the pit stop penalty is 29.5 seconds, and despite the lap count that could otherwise support a second stop, the race remains a one-stop due to the degradation savings. The hard tyre degrades only 0.837 and the medium 1.403 seconds, which makes the total degradation savings enough to offset another pit stop penalty. Consequently, the high cost of a pit stop overrides any structural flexibility in stint lengths, forcing a one-stop regime.

However, the moderating effect of P_{pit} is limited, and a high penalty alone is insufficient to override the combined pressure of a low t_{base} and a high lap count. For example, both Singapore and Emilia

Romagna follow the same pattern, where the P_{pit} is 30.7 and 31 seconds, and both fall into the two-stop strategy with two medium stints and one soft stint. The reason is that both the allocation compounds are C3-C4-C5 over a 62 and 63 lap distances, respectively, which generate degradation savings large enough to absorb the cost of the pit stop penalty. Thus, an extra pit is advantageous because the performance benefit of swapping to fresher tyres significantly outweighs the time lost in the pits, demonstrating that compound durability also drives the strategies when P_{pit} and N reach their limits.

Taken together, the analysis of P_{pit} confirms its role as a moderating threshold rather than a final decision-maker. A high pit stop penalty raises the bar to justify an additional stop and preserves one-stop strategies on circuits near the boundary, as demonstrated by France, Qatar and Britain. A low penalty lowers the threshold and enables two-stop strategies even on circuits with moderate degradation profiles, as shown by Australia. However, when the combined pressure of the t_{base}, N, and compound allocation generates degradation savings large enough to overcome even the highest penalties on the calendar, the penalty loses its protective effect entirely.

Overall, the analysis conducted in Section 6.1.3 demonstrates that the strategic regime of each circuit emerges from the simultaneous interaction of the three structural factors analysed. None of these factors operates as a sufficient condition in isolation, hence:

- t_{base} scales the relative cost of degradation per lap, but cannot predict regime when lap counts diverge, as the Canada-Spain comparison demonstrates.
- N amplifies the quadratic degradation burden, but requires sufficiently limited compound durability to cross the breakeven threshold, as Spain's C1-C2-C3 allocation illustrates.
- P_{pit} raises or lowers the threshold that degradation savings must exceed, but is overridden when the combined pressure of the other two parameters are large enough, as Singapore and Emilia Romagna confirm.

Consequently, Pirelli's compound allocation acts as a fourth contextual condition that modulates the effect of all three, by determining the durability margins available to each stint.

6.2 Stochastic Model Analysis

Following the same methodology used in the deterministic analysis, the stochastic model was first run using the hard tyre performance range of [1.047, 1.511] seconds, defined in Chapter 5. However, this initial run resulted in the hard tyre being rarely chosen, which repeated the same calibration issue observed in the deterministic baseline. The underlying issue was structural, as the scenarios were sampled uniformly across the entire range, causing a large proportion of instances to fall within the slower end of the spectrum. In those cases, the hard tyre delivered weaker performance, making it unable to compete effectively. Consequently, the hard tyre was selected in very few scenarios, producing a tyre usage distribution of 18 soft, 46 medium, and seven hard stints, which did not match the reality of Formula 1, like presented in Table 6.6.

To address this, the upper limit of the hard tyre's performance range was reduced by 25%, resulting in a recalibrated range of [1.047, 1.133] seconds. This adjustment ensures that the sampled scenarios remain within a performance window where the hard tyre is competitive, without changing the lower limit or the distribution of the other parameters. Additionally, it is important to note that this recalibration was not done arbitrarily, but it was based on the sensitivity analysis conducted on the deterministic version. In that analysis, the 25% reduction was identified as the exact threshold at which the model begins to replicate the real-world tyre usage order of `Medium > Hard > Soft`, and so by applying this same reduction to the stochastic model, the two versions remain consistent and directly comparable. Therefore, all stochastic results presented from this point forward are based on this updated configuration.

6.2.1 Baseline Stochastic Performance

The stochastic optimisation model was run across all 25 Grand Prix circuits using the updated parameter settings. For each race, 20 scenarios were sampled from the empirical distribution of β_{t0}^{ω}, β_{t1}^{ω}, and P_{pit}^{ω}. The main results are presented in Table 6.10.

Table 6.10: Main stochastic optimisation results across 25 Grand Prix circuits.

Race	N (laps)	t_{base} (s)	Expected Pit Stop Time (s)	Type	Tyre Sequence	Pit Window	Stint Length	Expected Total Time (s)
Bahrain GP	57	89.179	24.651672	2-Stop	S-M-M	[13, 35]	13, 22, 22	5182.037
Saudi Arabian GP	50	87.472	25.555414	1-Stop	H-M	[24]	24, 26	4456.807
Australian GP	58	75.915	24.466324	2-Stop	M-S-M	[23, 35]	23, 12, 23	4501.371
Emilia Romagna GP	63	74.746	24.790392	2-Stop	M-S-M	[25, 38]	25, 13, 25	4818.579
Miami GP	57	86.841	25.393946	2-Stop	M-S-M	[22, 35]	22, 13, 22	5043.725
Spanish GP	66	71.383	25.197014	2-Stop	S-M-M	[13, 40]	13, 27, 26	4821.248
Monaco GP	78	70.270	24.564690	2-Stop	M-H-M	[25, 53]	25, 28, 25	5618.636
Azerbaijan GP	51	100.203	25.589913	1-Stop	M-H	[24]	24, 27	5198.196
Canadian GP	70	72.000	25.507601	2-Stop	M-H-M	[22, 47]	22, 25, 23	5161.138
British GP	52	85.819	25.795430	1-Stop	H-M	[27]	27, 25	4551.225
Austrian GP	71	64.314	25.197173	2-Stop	M-M-H	[22, 44]	22, 22, 27	4694.418
French GP	53	90.872	25.194259	1-Stop	H-M	[30]	30, 23	4900.539
Hungarian GP	70	75.227	24.517201	2-Stop	M-S-M	[25, 45]	25, 20, 25	5388.865
Belgian GP	44	103.665	25.208043	1-Stop	H-M	[24]	24, 20	4632.928
Dutch GP	72	69.673	25.216916	2-Stop	S-M-M	[14, 43]	14, 29, 29	5132.290
Italian GP	53	79.327	25.091349	1-Stop	H-M	[30]	30, 23	4292.755
Singapore GP	62	89.525	24.730849	2-Stop	S-M-M	[13, 37]	13, 24, 25	5656.143
Japanese GP	53	88.197	25.373363	1-Stop	M-H	[24]	24, 29	4763.475
United States GP	56	92.330	25.005663	1-Stop	M-H	[25]	25, 31	5260.778
Mexico City GP	71	75.946	25.782944	2-Stop	M-M-H	[20, 41]	20, 21, 30	5518.640
São Paulo GP	71	70.727	25.389442	2-Stop	M-H-M	[22, 49]	22, 27, 22	5149.133
Abu Dhabi GP	58	82.595	25.298755	2-Stop	S-M-M	[11, 34]	11, 23, 24	4891.615
Qatar GP	57	80.520	25.360276	2-Stop	S-M-M	[11, 34]	11, 23, 23	4682.472
Las Vegas GP	50	92.312	25.303292	1-Stop	H-M	[27]	27, 23	4702.367
Chinese GP	56	93.660	25.041688	1-Stop	H-M	[30]	30, 26	5337.895

A first observation concerns the expected time lost during a pit stop. While in the deterministic model the values were fixed and ranged between 18.877 seconds at the Australian GP and 37.085 seconds at the French GP, the stochastic model settled between 24.5 and 25.8 seconds for every circuit, regardless of the actual pit lane characteristics. This convergence is a direct consequence of the modelling assumption introduced in Chapter 5. where a single global triangular distribution was applied uniformly to all circuits. Although the dataset includes individual pit stop times for the 2022 to 2024 season, circuit-specific details such as pit lane length or speed limits were not available, and consequently, a single global distribution was applied to represent the general pit stop process. This method was considered preferable to assigning arbitrary circuit-specific values, as it relies on the entire empirical dataset rather than on unsupported assumptions. The expected value is found by computing $\frac{minimum+maximum+mode}{3}$, which, given the empirical parameters, produces a consistent central tendency across circuits. Consequently, the extreme penalties that influenced the deterministic model are pulled towards the middle. Therefore, this changes how much the pit stop penalty impacts the overall strategy.

Regarding tyre usage, the stochastic model shows a total of 10 soft stints, 40 medium stints, and 15 hard stints across the 25 races, as summarised in Table 6.11.

Table 6.11: Total tyre usage distribution from the stochastic model.

Compound	Total Stints	Percentage (%)
Soft	10	14.9
Medium	40	59.7
Hard	15	22.4

This distribution maintains the `M > H > S` hierarchy established in the earlier deterministic model. The hard tyre makes up 22.4% of all stints, while the medium tyre remains the most popular, as it is the most versatile, at 59.7%. The ratio of the hard tyre to soft tyres is 1.5, while still below the empirical ratio of approximately 2.0, this represents a significant improvement over the uncalibrated baseline.

Moreover, in terms of race strategies, the model identifies 10 one-stop races and 15 two-stop races across the calendar, with no circuit requiring more stops. Compared to the deterministic reference model, which produced 14 one-stop and 11 two-stop races, the stochastic model exhibits a systematic shift towards more multi-stop strategies. This shift is a direct consequence of the P_{pit}^{ω} homogenisation, as these penalties converge towards the distribution mean of approximately 25 seconds, which lowers the breakeven threshold for a second stop and makes two-stop strategies mathematically viable at circuits where the deterministic penalty had previously outweighed their potential benefits.

6.2.2 Impact of Parameter Uncertainty on Optimal Strategies

Despite the overall stability already described, some races resulted in optimal strategies that differed from the deterministic model. These changes fall into three categories: races where the number of pit stops increased from one to two, races where the tyre choice changed, and races where the stop count and tyre selection remained the same but with different stint lengths. For the remaining races that kept the same strategy, tyre selection, and stint lengths, but with different tyre sequences, the results are analytically identical to the deterministic model, because the model utilises a constant degradation multiplier with a value of one, for all stints. Therefore, the order in which the tyres are utilised does not affect the total race time. However, if the model were updated to include changing fuel loads and track conditions, the sequence would become strategically important, because early and late stints would have different degradation profiles.

Therefore, Table 6.12 provides a side-by-side comparison of the deterministic and stochastic strategies for all 25 circuits. As shown in the table, most circuits produce identical or equivalent strategies in both versions, which confirms that the deterministic model is generally robust.

Table 6.12: Side-by-side comparison of deterministic and stochastic optimal strategies for all 25 circuits.

	Deterministic				Stochastic			
Race	Type	Pit Window	Tyre Sequence	Stint Lengths	Type	Pit Window	Tyre Sequence	Stint Lengths
Bahrain GP	1-Stop	[26]	M-H	26, 31	2-Stop	[13, 35]	S-M-M	13, 22, 22
Saudi Arabian GP	1-Stop	[23]	M-H	23, 27	1-Stop	[24]	H-M	24, 26
Australian GP	2-Stop	[23, 35]	M-S-M	23, 12, 23	2-Stop	[23, 35]	M-S-M	23, 12, 23
Emilia Romagna GP	2-Stop	[24, 38]	M-S-M	24, 13, 26	2-Stop	[25, 38]	M-S-M	25, 13, 25
Miami GP	1-Stop	[31]	H-M	31, 26	2-Stop	[22, 35]	M-S-M	22, 13, 22
Spanish GP	1-Stop	[29]	M-H	29, 37	2-Stop	[13, 40]	S-M-M	13, 27, 26
Monaco GP	2-Stop	[25, 53]	M-H-M	25, 28, 25	2-Stop	[25, 53]	M-H-M	25, 28, 25
Azerbaijan GP	1-Stop	[24]	M-H	24, 27	1-Stop	[24]	M-H	24, 27
Canadian GP	2-Stop	[25, 50]	M-M-S	25, 25, 20	2-Stop	[22, 47]	M-H-M	22, 25, 23
British GP	1-Stop	[28]	H-M	28, 24	1-Stop	[27]	H-M	27, 25
Austrian GP	2-Stop	[23, 46]	M-M-H	23, 23, 25	2-Stop	[22, 44]	M-M-H	22, 22, 27
French GP	1-Stop	[30]	H-M	30, 23	1-Stop	[30]	H-M	30, 23
Hungarian GP	2-Stop	[25, 45]	M-S-M	25, 20, 25	2-Stop	[25, 45]	M-S-M	25, 20, 25
Belgian GP	1-Stop	[16]	S-M	16, 28	1-Stop	[24]	H-M	24, 20
Dutch GP	2-Stop	[16, 44]	S-M-M	16, 28, 28	2-Stop	[14, 43]	S-M-M	14, 29, 29
Italian GP	1-Stop	[24]	M-H	24, 29	1-Stop	[30]	H-M	30, 23
Singapore GP	2-Stop	[13, 37]	S-M-M	13, 24, 25	2-Stop	[13, 37]	S-M-M	13, 24, 25
Japanese GP	1-Stop	[29]	H-M	29, 24	1-Stop	[24]	M-H	24, 29
United States GP	1-Stop	[25]	H-M	25, 31	1-Stop	[25]	H-M	25, 31
Mexico City GP	2-Stop	[23, 46]	M-M-H	23, 23, 25	2-Stop	[20, 41]	M-M-H	20, 21, 30
São Paulo GP	2-Stop	[22, 49]	M-H-M	22, 27, 22	2-Stop	[22, 49]	M-H-M	22, 27, 22
Abu Dhabi GP	2-Stop	[11, 35]	S-M-M	11, 24, 23	2-Stop	[11, 35]	S-M-M	11, 24, 23
Qatar GP	1-Stop	[26]	M-H	26, 31	2-Stop	[11, 34]	S-M-M	11, 23, 23
Las Vegas GP	1-Stop	[27]	H-M	27, 23	1-Stop	[27]	H-M	27, 23
Chinese GP	1-Stop	[25]	H-M	25, 31	1-Stop	[30]	H-M	30, 26

Transitioning from One-Stop to Two-Stop Strategies

There are four circuits that shift from a one-stop to a two-stop strategy: Bahrain, Miami, Spain, and Qatar Grand Prix. In the deterministic model, the pit stop time loss was high enough to make a second stop too costly, while in the stochastic model, this is replaced by an average penalty of about 25 seconds, which lowers the breakeven point, making the time saved by using a new fresh set of tyres more valuable than the time lost in the pit lane.

At the Qatar GP, the deterministic P^{ω}_{pit} was 29.053 seconds, which favoured a one-stop `[M-H]` strategy. In the stochastic model, the expected P^{ω}_{pit} is 25.36 seconds, which represents nearly a four-second reduction, making a two-stop `[S-M-M]` strategy the optimal choice. Similar shifts occurred at the other three circuits. These flips confirm that the homogenisation of the pit stop penalty is the main driver of strategic change. Tracks with penalties much higher or lower than this threshold, such as the French GP and Australian GP, remained unchanged because their strategy is less sensitive to these shifts.

Compound Substitution within Stable Strategy Types

Moreover, the circuits of Canada and Belgium kept the same number of stops but swapped the soft compound for the hard compound. In both cases, this happened because the expected pace penalty for the hard tyre was lower than the fixed value used in the deterministic model, almost close to the lower limit of the β^{ω}_{t0} range.

At the Canada circuit, the deterministic model utilised a strategy of `[M-M-S]`, while in the stochastic it used a `[M-H-M]` strategy, where the hard tyre pace penalty reduced from 1.1311 to 1.057, becoming a more competitive alternative to the soft. Combined with a lower degradation cost, the hard tyre delivered a lower total cost across the sampled scenarios, resulting in an accumulated degradation of only 0.669

seconds, compared to 2.774 seconds when utilising the soft compound. Additionally, at the Belgian GP, a similar shift occurred, where the strategy [S-M] with a Δt of 2.190 and 1.647 seconds, respectively, switched to a [H-M] with a Δt of 0.577 and 1.113 seconds.

Therefore, in this version of the model, all one-stop strategies only utilise the medium and the hard compound, while the soft compound is now exclusively reserved for two-stop strategies. This structural change is justified by the soft compound's higher degradation cost, so it is not viable for the long, sustained stints required by a one-stop strategy. Instead, the model restricts the soft compound to short stints, typically lasting between 11 and 14 laps, which allows the driver to benefit from its high initial performance without reaching the point of terminal degradation. The exception to this rule is the Hungarian GP where a two-stop strategy [M-S-M] is employed with stint lengths of 25, 20, and 25 laps. For this race the compound allocation is C3-C4-C5, with the medium compound (C4) reaching its maximum lifespan, 25 laps. With a total race distance of 70 laps, a pure three-stint medium strategy would require each stint to cover approximately 23-24 laps, which is within the lifespan limit. However, a pure [M-M-M] sequence is excluded by the FIA two-compound rule, which obligates the use of at least two different dry-weather compounds. The hard compound, being C3 with a lifespan of 30 laps, could theoretically satisfy this constraint, but its pace penalty and degradation profile make it a more expensive option per lap than the soft compound over a short sprint stint. Consequently, the model identifies the soft compound as the optimal compound, covering 20 laps. It should be noted, however, that these lifespan values are not absolute, and in real-world racing, these limits depend on several external factors and conditions, and could be extended for a few laps based on track conditions, the driver, or the risk profile of the team's decision-maker, whether risk-averse, risk-neutral, or risk-seeking in their strategy execution.

Stint Length Redistribution

Beyond changes in stop counts or tyre choices, many stable races showed a shift in the optimal number of laps for each stint, as shown in Table 6.13. This happens across most of the season, confirming that the uncertain parameters not only change which strategy to choose, but also how the race distance is shared among compounds. In all cases, the redistribution is the result of the combined effect of parameter shifts across all three compounds simultaneously, where some compounds become cheaper and others more expensive.

The most common pattern is an increase in laps of the hard compound, due to the reduction of the hard compound's pace penalty, which lowers the cost per lap and makes longer hard stints optimal. At the Austrian Grand Prix, an improved hard pace and lower degradation ($\Delta\beta_{t0} = -0.083$ and $\Delta\beta_{t1} = -0.005$) extends the hard stint from 25 to 27 laps at the expense of the medium stints. Moreover, at the Mexico City and French Grand Prix, the hard stints gain five and six laps, respectively, due to the improvement of the hard compound, since $\Delta\beta_{t0} < 0$, and the simultaneously increase in medium compound cost, as $\Delta\beta_{t0} > 0$.

Additionally, the same combined logic applies to races where the medium stints gain laps. At the British GP, the medium gains one lap because its improvement ($\Delta\beta_{t0} = -0.091$) outweighs the hard's improvement ($\Delta\beta_{t0} = -0.062$), making the medium the cheaper compound.

Table 6.13: Expected parameter deviations per compound for all 25 circuits, computed as the difference between the stochastic expected value and the deterministic fixed value from Table 6.7.

	Soft		Medium		Hard	
Race	$\Delta\beta_{t0}$	$\Delta\beta_{t1}$	$\Delta\beta_{t0}$	$\Delta\beta_{t1}$	$\Delta\beta_{t0}$	$\Delta\beta_{t1}$
Bahrain GP	0.0000	$+0.0018$	$+0.1727$	$+0.0024$	-0.0371	-0.0064
Saudi Arabian GP	0.0000	-0.0152	-0.0489	$+0.0037$	-0.0011	-0.0056
Australian GP	0.0000	-0.0112	-0.1072	$+0.0027$	-0.0241	$+0.0023$
Emilia Romagna GP	0.0000	$+0.0063$	-0.0634	$+0.0027$	-0.0806	-0.0030
Miami GP	0.0000	-0.0096	-0.1317	-0.0047	-0.0311	$+0.0033$
Spanish GP	0.0000	-0.0092	$+0.0439$	-0.0011	-0.0138	$+0.0064$
Monaco GP	0.0000	$+0.0150$	-0.0810	$+0.0032$	-0.0321	-0.0019
Azerbaijan GP	0.0000	$+0.0119$	-0.1034	-0.0006	-0.0010	$+0.0044$
Canadian GP	0.0000	-0.0037	-0.0173	-0.0016	-0.0756	-0.0013
British GP	0.0000	$+0.0088$	-0.0905	0.0000	-0.0617	-0.0059
Austrian GP	0.0000	$+0.0002$	-0.1105	-0.0055	-0.0831	-0.0052
French GP	0.0000	$+0.0121$	$+0.0945$	-0.0058	-0.0648	-0.0046
Hungarian GP	0.0000	$+0.0089$	$+0.1033$	-0.0037	-0.0784	$+0.0032$
Belgian GP	0.0000	-0.0004	$+0.1666$	-0.0055	-0.0611	-0.0047
Dutch GP	0.0000	$+0.0104$	-0.1400	$+0.0040$	-0.0396	$+0.0055$
Italian GP	0.0000	-0.0156	$+0.0160$	$+0.0042$	-0.0034	$+0.0011$
Singapore GP	0.0000	-0.0003	-0.0335	-0.0057	-0.0571	-0.0052
Japanese GP	0.0000	$+0.0124$	-0.1793	$+0.0055$	-0.0274	$+0.0021$
United States GP	0.0000	-0.0047	$+0.1622$	$+0.0008$	-0.0205	$+0.0050$
Mexico City GP	0.0000	$+0.0095$	$+0.0945$	-0.0052	-0.0526	-0.0035
São Paulo GP	0.0000	-0.0090	-0.0528	-0.0013	-0.0020	$+0.0073$
Abu Dhabi GP	0.0000	-0.0063	-0.1234	-0.0024	-0.0780	-0.0057
Qatar GP	0.0000	$+0.0145$	-0.1783	-0.0042	-0.0398	$+0.0012$
Las Vegas GP	0.0000	$+0.0110$	$+0.0571$	$+0.0018$	-0.0115	$+0.0063$
Chinese GP	0.0000	-0.0118	-0.0508	-0.0057	-0.0671	$+0.0045$

Overall, this happens because the expected values of pace and wear of each compound are a balance of cost, and any redistribution of stint lengths is determined by the relative competitiveness of all three compounds, rather than by changes to a single tyre. Consequently, the optimisation model reallocates laps according to the combined effect of the parameter shifts across the entire tyre set.

Expected Total Race Time

A consistent pattern appears when comparing the expected total race times of the stochastic model to those of the deterministic model across all stable races. In every single case, the stochastic expected total time is lower, with differences ranging from 4.8 seconds at the Australian GP to 28.8 seconds at the Singapore GP, as shown in Table 6.14. This systematic reduction is driven by two separate factors. First, the uniform triangular pit stop distribution produces an expected penalty of approximately 25.2 seconds for every circuit. Second, as shown in Table 6.13, the expected values for the performance parameters under the stochastic scenarios are lower than the fixed deterministic values for most of the circuits. It is important to note that these performance parameters are expressed as pace penalties ($\Delta\beta$), where a negative value indicates that the tyre performs faster than the deterministic baseline. For instance, the hard compound's pace penalty ($\Delta\beta_{t0}$) is negative in 24 out of 25 circuits, and the medium compound's ($\Delta\beta_{t0}$) is negative in 15 out of 25, meaning that the sampled scenarios on average assume better tyre performance compared to the deterministic baseline. Since both factors reduce the cost of the strategy, the stochastic expected total time being universally lower is a direct consequence of how the parameter distributions were constructed.

Table 6.14: Total race time comparison between stochastic and deterministic models for all 25 circuits, and between the stochastic model and a deterministic model that utilises the expected values (T^*), including circuit-specific and expected pit stop time loss. Δ Total Time is computed as E[Total Time] $-$ Total Time.

	Deterministic		Stochastic		Comparison		
Race	P_{pit} (s)	Total Time (s)	E[P_{pit}] (s)	E[Total Time] (s)	T^* (s)	Δ Total Time (s)	VSS (s)
Bahrain GP	25.771	5192.489	24.652	5182.037	5190.56	-10.452	-8.519
Saudi Arabian GP	21.988	4465.012	25.555	4456.807	4466.29	-8.205	-9.485
Australian GP	18.877	4506.216	24.466	4501.371	4512.90	-4.845	-11.531
Emilia Romagna GP	30.964	4846.004	24.790	4818.579	4832.68	-27.425	-14.096
Miami GP	22.037	5055.488	25.394	5043.725	5053.95	-11.763	-10.223
Spanish GP	22.909	4837.306	25.197	4821.248	4843.44	-16.058	-22.195
Monaco GP	26.370	5641.043	24.565	5618.636	5633.56	-22.407	-14.922
Azerbaijan GP	21.620	5203.344	25.590	5198.196	5206.00	-5.148	-7.806
Canadian GP	24.986	5181.513	25.508	5161.138	5178.79	-20.375	-17.654
British GP	29.507	4565.436	25.795	4551.225	4555.52	-14.211	-4.293
Austrian GP	22.470	4704.755	25.197	4694.418	4698.69	-10.337	-4.270
French GP	37.085	4928.643	25.194	4900.539	4913.60	-28.104	-13.064
Hungarian GP	22.437	5402.304	24.517	5388.865	5408.38	-13.439	-19.519
Belgian GP	23.548	4640.618	25.208	4632.928	4643.30	-7.690	-10.369
Dutch GP	21.269	5153.094	25.217	5132.290	5157.29	-20.804	-25.004
Italian GP	25.284	4304.957	25.091	4292.755	4306.63	-12.202	-13.876
Singapore GP	30.656	5684.933	24.731	5656.143	5668.24	-28.790	-12.092
Japanese GP	24.461	4774.244	25.373	4763.475	4772.36	-10.769	-8.886
United States GP	25.217	5277.143	25.006	5260.778	5282.92	-16.365	-22.140
Mexico City GP	23.341	5532.369	25.783	5518.640	5536.42	-13.729	-17.780
São Paulo GP	24.966	5165.071	25.389	5149.133	5164.56	-15.938	-15.431
Abu Dhabi GP	22.785	4901.472	25.299	4891.615	4899.31	-9.857	-7.695
Qatar GP	29.053	4702.207	25.360	4682.472	4691.46	-19.735	-8.986
Las Vegas GP	22.334	4707.358	25.303	4702.367	4713.96	-4.991	-11.588
Chinese GP	23.743	5350.149	25.042	5337.895	5347.99	-12.254	-10.091

The impact of each factor changes from track to track, which can be seen by looking at races where only one of these factors has a significant effect. At the United States Grand Prix, the pit stop time losses are similar in both models, with a difference of only 0.2 seconds, yet the expected total race time still drops by 16.4 seconds. Therefore, this time difference can be justified by the β parameters, where the expected pace penalty for the hard tyre is approximately 0.020 seconds lower than the deterministic value, which saves about 0.6 seconds over the 31 laps. Conversely, at the French GP, where the pit stop loss drops by approximately 11.9 seconds and the expected total time is reduced by 28.1 seconds, the β parameters account for the remaining drop of 16 seconds.

Consequently, the expected total race times in the stochastic model should not be interpreted as exact predictions of how long a real race will take, but instead, as the best possible result based on the specific parameter distributions chosen.

6.2.3 Value of the Stochastic Solution

To formally quantify the benefit of incorporating parameter uncertainty directly into the optimisation process, the Value of the Stochastic Solution (VSS) was computed for each of the 25 circuits. The VSS is defined as the difference between the expected total race time obtained from the stochastic model, E[Total Time], and the reference value T^*, which represents the total race time obtained by solving the model deterministically utilising the expected values of the uncertain parameters, β_{t0}^{ω}, β_{t1}^{ω}, and P_{pit}^{ω}, as fixed constants. Therefore:

$$\text{VSS} = \mathsf{E}[T_{\text{total}}] - T^* \tag{6.1}$$

Across all 25 circuits, the VSS has negative values, confirming that $\mathsf{E}[T_{\text{total}}] < T^*$ in every case, as presented in Table 6.14. The average VSS of -12.86 seconds highlights the practical significance of accounting for uncertainty, as a time difference of this magnitude in Formula 1 can be sufficient to alter finishing positions and ultimately affect race outcomes.

The magnitude of this improvement is significant in the context of Formula 1, ranging from -4.27 to -25.044 seconds per race. In relative terms, this represents a performance gain of approximately 0.09% to 0.5% of the total race time. Given that a gap of 20 seconds can easily determine the difference between a podium finish and a mid-field position, these results highlight the substantial competitive advantage provided by a robust stochastic approach.

This outcome proves that the stochastic model consistently outperforms the deterministic version by avoiding the limitations of point-estimate optimisation. The formulation of T^* reduces the full distribution of each uncertain parameter to a single point, the mean, and optimises the strategy as if this value will occur with certainty. Conversely, the stochastic model evaluates the objective function across all 20 discrete scenarios simultaneously, allowing the optimal first-stage strategy to be selected based on its aggregate performance across the entire scenario set, rather than its performance under a single average condition.

Furthermore, as all circuits were modelled under the same global distributions, as described in Chapter 5, the variation in the VSS magnitude does not reflect differences in the uncertainty, since the same distributions were sampled for every race. Consequently, this variation can be explained by the structural characteristics of each circuit, such as the base lap time, the total number of laps, and Pirelli's compound allocation. Thus, the magnitude of the VSS value is a reflection of how a specific race's structure interacts with the sampled data.

Overall, the fact that VSS remains negative across the entire calendar provides strong empirical validation for the two-stage stochastic programming approach adopted in this model. It confirms that explicitly modelling uncertainty in tyre wear and pit stops, rather than relying on point estimates, yields measurably better expected outcomes. This reinforces the practical value of the stochastic framework over the deterministic baseline established earlier.

6.3 Comparison of Strategies

In order to compare the relevance and accuracy in a real-world setting the optimal strategies generated by the stochastic model are compared against the three different benchmarks. The first is the deterministic model, which utilises fixed values to serve as a baseline, to allow a clear measure of how much the uncertain parameters impact the results. The second benchmark is Pirelli's official pre-race strategy recommendations, which reflect the professional judgement of tyre manufacturers. These recommendations are based on Pirelli's extensive knowledge of historical degradation patterns, track characteristics, the physical properties of the tyres (which are not fully disclosed), and specific tyre allocations selected for each race. The third benchmark is the strategy actually used by the winner in selected 2024 Formula 1 races. To ensure a fair comparison, only races that took place in clean conditions are considered, meaning races that did not have yellow or red flags, or rainy weather, since these factors would distort the results by deviating from the model's assumptions. Therefore, each of these three comparisons serves a specific goal: the deterministic comparison to measure the effect of uncertainty; the Pirelli comparison to check whether the results align with expert advice, and the winner comparison to test if the model produces strategies that are actually capable of winning in real racing conditions.

6.3.1 Alignment with Pirelli's Pre-Race Strategy Predictions

Prior to every race, Pirelli releases official strategy graphics that outline two to four potential race plans for each circuit. These suggestions, which include specific tyre sequences and recommended pit windows, are based on the track's compound allocation, historical degradation data, and empirical knowledge of each track's strategic landscape, and reflect the tyre manufacturer's expert insight into all the characteristics. Therefore, the comparison between stochastic strategy and pre-race Pirelli strategy is evaluated along two different dimensions: whether the model's optimal strategy type and compound sequence corresponds to one of Pirelli's suggested options, and whether the model's pit window falls within Pirelli's recommended lap interval. Because Pirelli provides various one-stop and two-stop strategies, comparisons are made against the most similar viable option rather than a single "correct" approach.

Therefore, the selected races provide a comprehensive overview by covering all three compound allocation ranges and both primary strategy types. The chosen circuits are as follows:

- Italian GP: Features a one-stop optimal strategy and utilises the `C3-C4-C5` compound range.
- British GP: Features a one-stop optimal strategy and utilises the `C1-C2-C3` compound range.
- Dutch GP: Features a two-stop optimal strategy and utilises the `C1-C2-C3` compound range.
- Hungarian GP: Features a two-stop optimal strategy and utilises the `C3-C4-C5` compound range.
- Monaco GP: Features a two-stop optimal strategy and utilises the `C3-C4-C5` compound range; this is a highly significant event where a one-stop strategy is traditionally the standard approach.
- Saudi Arabian GP: Features a one-stop optimal strategy and utilises the `C2-C3-C4` compound range.
- Miami GP: Features a two-stop optimal strategy and utilises the `C2-C3-C4` compound range.

Italian Grand Prix

Table 6.15: Pirelli's pre-race strategy recommendations for the Italian Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [30]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	[M-H]	19-25
Option 2	Two-stopper	[M-H-H]	11-16, 30-36
Option 3	One-stopper	[H-M]	27-33
Stochastic	One-stopper	[H-M]	30
C3 - Hard C4 - Medium C5 - Soft			

The Italian GP represents the strongest alignment case between the stochastic model and Pirelli's recommendations. For this race, the tyre allocation is C3-C4-C5 with estimated lifespans of 30, 25, and 20, respectively. Pirelli provided three primary strategies: a one-stop [M-H] strategy with a pit window between laps 19 and 25, a one-stop [H-M] strategy with a pit between laps 27 and 33, and a two-stop [M-H-H] with stops occurring between laps 11-16 and 30-36. The stochastic model identified a one-stop [H-M] strategy as optimal, prescribing a pit stop at lap 30 with a stint length of 30 laps on the hard compound, followed by a 23 lap stint on the medium compound.

On the stop count dimension, the model aligns perfectly with Pirelli's primary recommendation, as both identify the one-stop strategy as optimal. Regarding the compound sequence, the model's [H-M] sequence matches Pirelli's second one-stop option exactly, with the designated pit window at lap 30 falling precisely within the recommended range of laps 27 to 33. Furthermore, it is important to note that the [M-H] sequence is also mathematically viable within this simulation, since the model assumes constant degradation multipliers, and therefore the tyre sequences are commutative, yielding an identical total race time regardless of the stint order. Consequently, the calculated 23 lap stint of the medium compound remains within the operational window of the first one-stop Pirelli's strategy. Finally, with all three dimensions of alignment, the stop count, tyre sequence and pit window satisfied, this race constitutes the most strongly validated strategy among the circuits analysed.

British Grand Prix

Table 6.16: Pirelli's pre-race strategy recommendations for the British Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [31]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	[S-M]	18-24
Option 2	One-stopper	[S-H]	16-22
Option 3	One-stopper	[M-S]	30-36
Option 4	One-Stopper	[M-H]	24-30
Stochastic	One-stopper	[H-M]	27
C1 - Hard C2 - Medium C3 - Soft			

The British Grand Prix presents a unique case where the stochastic model's primary output requires interpretation through the lens of the model's underlying assumptions. Pirelli's pre-race strategy for the 2024 Silverstone event proposed four distinct one-stop strategies, combining all compounds in [S-M], [S-H], [M-S], and [M-H] sequences. The stochastic model identifies a one-stop [H-M] strategy, with a pit stop at lap 27 and stint lengths of 27 and 25 laps, respectively.

Although the [H-M] sequence is not explicitly recommended in Pirelli's provided data, the model's results demonstrate a strong strategic alignment with the [M-H] option, when accounting for tyre degradation influence over each stint. Assuming commutativity between tyre sequences ($M_i = 1$), the model's strategy is mathematically equivalent to an [M-H] approach with a pit stop at lap 25. Furthermore, since the transition at lap 25 for the [M-H] sequence and lap 27 for the [H-M] sequence yield an equivalent expected total race time, the model's optimal solution remains robustly aligned with Pirelli's recommendation. Consequently, the model's internal logic mirrors the performance profile of Pirelli's [M-H] strategy, which targets a pit window between laps 24 and 30.

The analytical insight here is that the model converges on the same compound pairing (C1 and C2), confirming that the hard and medium compounds represent the optimal performance window for this circuit. The slight shift in the recommended sequence arises from the simplification of $M_i = 1$, which does not account for the transient effects of heavy initial fuel loads and the evolution of track grip. Consequently, the model correctly identifies the optimal tyre pair and timing parameters, validating its predictive capability despite the inherent constraints of a linear degradation model.

Dutch Grand Prix

Table 6.17: Pirelli's pre-race strategy recommendations for the Dutch Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [32]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	[S-H]	24-30
Option 2	One-stopper	[S-M]	28-33
Option 3	One-stopper	[H-S]	43-48
Option 4	Two-Stopper	[S-M-S]	20-26, 47-53
Stochastic	Two-stopper	[S-M-M]	14, 43
C1 - Hard C2 - Medium C3 - Soft			

For the Dutch Grand Prix, Pirelli provided three one-stop strategies, and one two-stop strategy, being [S-M-S] with pit windows in laps between 20 and 26, and between 47 and 53. On the other hand, the stochastic model identifies a two-stop [S-M-M] with pit stops at laps 14 and 43, as optimal.

While there is a divergence in the final compound selection, this can be attributed to the model's objective function in relation to tyre degradation parameters. Pirelli's [S-M-S] strategy is designed to exploit the soft compound's peak performance for a high-pace burst at the end of the race, likely assuming a specific track evolution profile where grip levels increase significantly in the final laps. Conversely, the model's selection of an [S-M-M] sequence suggests a strategic preference for stability and durability in the latter two-thirds of the race. By opting for a second stint on the medium compound, the model demonstrates that, under its calculated stochastic degradation parameters, the medium tyre provides a more favourable cumulative time, effectively mitigating the risk of late-race degradation that would be associated with the soft compound.

Hungarian Grand Prix

Table 6.18: Pirelli's pre-race strategy recommendations for the Hungarian Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [33]

Strategy	Type	Compound Sequence	Pit Window
Option 1	Two-stopper	[M-H-M]	16-21, 45-51
Option 2	Two-stopper	[M-H-H]	13-18, 41-47
Option 3	Two-stopper	[S-H-H]	9-14, 38-44
Option 4	Two-Stopper	[S-H-M]	9-14, 43-49
Stochastic	Two-stopper	[M-S-M]	25, 45
C3 - Hard C4 - Medium C5 - Soft			

The Hungarian Grand Prix favours multi-stop strategies due to its high thermal degradation. Pirelli's 2024 guidance provided four two-stop options, predominantly utilising the medium and the hard compound, with limited soft usage, restricted to short stints due to rapid thermal degradation. The stochastic model outputs a two-stop [M-S-M] strategy, with pit stops at laps 25 and 45.

The model aligns with the professional consensus on the two-stop configuration, and its final pit window at lap 45 falls within the range of Pirelli's recommendations. A notable divergence is the model's utilisation of the soft C5 compound for a 20 lap stint, whereas Pirelli's recommendation limits the soft stint to shorter durations, of maximum of 14 laps, to avoid performance degradation. This discrepancy can be explained by the model's distinct durability parameters, where the lifespan of the C5 is defined at 20 laps, and C4 and C3 at 25 and 30 laps, respectively.

The model's internal durability constraints for the soft compound are more permissive than those used by Pirelli. Consequently, the model identifies the soft tyre as a viable option for a 20 lap middle stint, whereas professional race engineers prioritise the medium or hard compound for longer durations to maintain pace consistency. This validates the model's ability to optimise within the constraints provided to it. This demonstrates that the model's outputs are a direct function of its calibrated lifespan constraints, providing a consistent logical basis for the strategies generated.

Monaco Grand Prix

The Monaco Grand Prix presents a unique strategic environment, defined by its extreme difficulty in overtaking, which shifts the strategic focus from track position racing to pit stop optimisation. Pirelli's 2024 strategies focused exclusively on one-stop strategies, largely centred on maximising the life of the C3 (hard) compound to bridge the majority of the race distance. In contrast, the stochastic model outputs a two-stop [M-H-M] strategy, with pit stops at laps 25 and 53.

Table 6.19: Pirelli's pre-race strategy recommendations for the Monaco Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [34]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	[M-H]	24-34
Option 2	One-stopper	[H-M]	47-57
Option 3	One-stopper	[S-H]	17-27
Stochastic	Two-stopper	[M-H-M]	25, 53
C3 - Hard C4 - Medium C5 - Soft			

The divergence between the model's and Pirelli's one-stop benchmark is fundamentally rooted in a tyre lifespan constraint. Pirelli's recommendation relies on the hard compound's ability to endure stints of approximately 50 laps, while the model is constrained to a maximum of 30 laps for the C3 compound. Consequently, the model is mathematically precluded from replicating the one-stop endurance profile. Instead, the model identifies the two-stop configuration as the optimal path given its internal durability parameters.

In 2025, the FIA introduced a regulation mandating the use of all three dry-weather compounds, effectively requiring a minimum of two pit stops per race. This initiative aimed to enhance strategic decision-making and incentivise on-track position battles, as the narrow streets of Monaco render traditional overtaking nearly impossible. Consequently, teams were forced to optimise their strategies under these more complex constraints. While the model's optimised strategy respects the two-stop requirement, it does not incorporate all three compounds. However, this limitation is mitigated by the 2026 regulatory reversal, which abolished the mandatory three-compound rule and once again allowed teams to utilise two-compound strategies, thereby broadening the spectrum of viable strategic options.

While the model does not replicate the real-world 2024 one-stop strategy, it successfully identifies how a two-stop race would optimise performance within the constraints of defined tyre durability. This highlights the model's robustness, where it does not simply copy the pre-race consensus, but it optimises within its own rigorous physical constraints.

Saudi Arabia Grand Prix

Table 6.20: Pirelli's pre-race strategy recommendations for the Saudi Arabia Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [35]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	[M-H]	19-25
Option 2	One-stopper	[S-H]	15-21
Option 3	One-stopper	[M-S]	28-35
Option 4	Two-Stopper	[S-M-S]	10-16, 30-36
Stochastic	One-stopper	[H-M]	24
C2 - Hard C3 - Medium C4 - Soft			

For the Jeddah Corniche Circuit, Pirelli in 2024 suggested three one-stop permutations and a two-stop strategy utilising the soft compound for the initial and final sprint phases. The stochastic model outputs a one-stop strategy starting with the hard compound, followed by a medium compound, with a tyre change at lap 24.

While the model identifies a one-stop outcome that aligns with the professional consensus, there is a distinct difference in compound prioritisation. Pirelli's first option is a one-stop strategy `[M-H]` with a pit window in laps between 19 and 25, which favours a longer hard stint to ensure durability, whereas the model suggests a longer medium stint with 26 laps. Under the model's assumption of a constant tyre degradation multiplier, the `[H-M]` sequence is equivalent to the `[M-H]` sequence, reasembling with the first Pirelli suggestion. However, in this configuration, the model recommends a pit stop at lap 26, slightly exceeding Pirelli's window of 19 to 25 laps. This divergence highlights the difference between an optimisation model that calculates the absolute minimum race time based on the inputs, like degradation parameters, and professional strategy, which often incorporates a safety margin for tyre management.

Miami Grand Prix

Table 6.21: Pirelli's pre-race strategy recommendations for the Miami Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [36]

Strategy	Type	Compound Sequence	Pit Window
Option 1	One-stopper	`[M-H]`	15-21
Option 2	One-stopper	`[H-M]`	35-41
Option 3	One-stopper	`[S-H]`	11-17
Option 4	One-Stopper	`[H-S]`	38-44
Stochastic	Two-stopper	`[M-S-M]`	22, 35
`C2` - Hard `C3` - Medium `C4` - Soft			

The Miami Grand Prix features a distinct compound allocation of `C2`, `C3`, and `C4` for the hard, medium, and soft compounds, respectively. Pirelli's 2024 pre-race suggests a set of one-stop strategies, primarily utilising the medium-hard or soft-hard sequence, with longer hard stints. Conversely, the stochastic model identifies a two-stop `[M-S-M]` strategy as optimal, with pit stops at laps 22 and 35, thereby producing medium stints of 22 laps each and a soft stint of 13 laps.

The discrepancies between the model's recommendation and Pirelli's one-stop strategy are primarily driven by the tyre lifespan constraints imposed within the stochastic model. Due to the assumption that $M_i = 1$, the model's selected `[M-S-M]` strategy is mathematically equivalent to the `[S-M-M]` and `[M-M-S]` configurations, which align with Pirelli's third and fourth options that prioritise a soft sprint at the beginning or end of the race. Furthermore, the model's 13 lap soft stint perfectly aligns with Pirelli's recommended performance window of 11 to 17 laps for the `C4` compound. However, to achieve a one-stop strategy using the hard-soft tyre pairing, the remaining race distance would require the hard compound to complete 44 laps of the 57 lap race. Additionally, even if the soft stint were extended to Pirelli's upper limit of 17 laps, the hard compound would still need to cover 40 laps. Therefore, neither scenario is feasible, as the model restricts the C2 compound to a maximum lifespan of 38 laps. Consequently, the model is unable to achieve a one-stop solution and instead replaces the hard stint with two 22 lap medium stints, which cover the remaining race distance while satisfying all durability constraints.

6.3.2 Race Winner Strategy Alignment

To assess the practical validity of the stochastic optimisation model, a comparison between its recommendations and the strategies actually employed by the winner of each Grand Prix of 2024 was made. For this analysis, data was utilised from the Kaggle repository "F1 Stint Data with Aggression Scores", which provided detailed information on pit stop timings, tyre compounds, and stint lengths for every driver across all races from 2018 to 2024, [37]. The comparison evaluates whether the winner's compound sequence and pit window align with the model's optimal recommendation, and whether the winner's race time and the winner's expected race time validate the model's performance prediction.

The race selection for this comparison was based on strict criteria. Only races conducted entirely under clean conditions without safety car deployments, virtual safety car periods, red flags, or adverse weather that would require wet-weather compounds were considered. These constraints minimise the divergence between the model's idealised assumptions of uninterrupted green-flag racing and actual race conditions, allowing for a more direct comparison between predicted and realised strategy. Therefore, five races were chosen that satisfied these criteria, based on a thorough review of official race reports on the Formula 1 website: the Dutch GP [38], the Italian GP [39], the Spanish GP [40], the Las Vegas GP [41], and the Monaco GP [42], though Monaco requires methodological clarification.

The Monaco Grand Prix was included because it is a race that heavily relies on pit-stop strategy. However, the race was affected by a red-flag stoppage on lap one following a multi-car collision before the first laps could be completed. After the restart, the remainder of the race proceeded under clean conditions. Crucially, when the race was red-flagged, all drivers elected to change their tyres during the restart procedure, treating the stoppage as an opportunity to adjust their strategy. In Monaco, where the layout makes on-track overtaking extremely difficult, the majority of drivers utilised the mandatory tyre change to execute their only pit stop, rather than attempting to pass competitors on track. The winner, like most drivers, utilised the red-flag period as the strategic pit stop, meaning the race was conducted as a one-stop strategy.

For each race, the winner's actual race strategy is compared against both the stochastic model's optimal strategy and the winner's expected race time computed with the model's expected β values for each compound at that circuit and the stochastic expected P_{pit}, applied to the winner's actual tyre sequence and pit timing via the model's objective function formulation. This recalculated expected total time is then compared against the winner's actual race time to quantify the deviation and assess whether the actual racing experience produced an outcome better or worse than the model's expected value under identical strategic choices. Therefore, the results are presented in Table 6.22.

Table 6.22: Model expected strategy and performance metrics compared against the winner's actual race strategy for selected Grand Prix.

	Dutch GP	Italian GP	Spanish GP	Las Vegas GP	Monaco GP
Stint 1 compound	Medium	Medium	Soft	Medium	Medium
Stint 1 length	28	15	17	12	1
Stint 2 compound	Hard	Hard	Medium	Hard	Hard
Stint 2 length	44	38	27	20	77
Stint 3 compound	-	-	Soft	Hard	-
Stint 3 length	0	0	22	18	0
Pit 1 lap	28	15	18	12	1
Pit 1 time	25.217	25.091	25.197	25.303	24.565
Pit 2 lap	-	-	44	32	-
Pit 2 time	0	0	21.942	25.303	0
Base time	69.673	79.327	71.383	92.312	70.270
N laps	72	53	66	50	78
Expected Total Race Time	5160.214	4310.126	4825.810	4732.988	5676.034
Expected Model's Total Race Time	5132.290	4292.755	4821.248	4702.367	5618.636
Winner's Actual Total Race Time	5445.519	4480.727	5300.227	4925.969	8595.554

As observed from Table 6.22, a consistent pattern emerges across all five selected races. The winner's actual total race time is systematically higher than the expected total race time computed by applying the model's stochastic parameters to the winner's actual strategy. This gap is a structural consequence of the model's idealised assumptions, as it optimises under clean racing conditions with no traffic, no wheel-to-wheel defending, and no external disruptions. In practice, actual lap times are influenced by a wide range of factors that the model does not account for, including ambient and track temperature, wind, humidity, and the progressive evolution of car weight as fuel is consumed. Beyond physical conditions, driver-specific factors such as individual tyre management style, physical fatigue over the race distance, and the tactical demands of defending or attacking positions on track all contribute to lap time deviations that a degradation-based model cannot capture.

Furthermore, the beta parameters utilised in the model were calibrated from a statistical average across multiple seasons and drivers, which represent a central tendency rather than the precise degradation profile of a specific race day. Consequently, the expected total race times should be interpreted as the lower bounds under the model's parameter assumptions rather than predictions of actual race outcomes. However, the magnitude of this gap is not constant across circuits, and varies according to how much real-world complexity was present in each specific race, as discussed in the following subsections.

Dutch Grand Prix

The Dutch GP presents a fundamental divergence between the stochastic model's recommendation and the strategy employed by the race winner, Lando Norris. The model identifies a two-stop `[S-M-M]` strategy with pit stops at laps 14 and 43 as optimal, favouring a short soft sprint stint length. In contrast, Norris executed a one-stop `[M-H]` strategy, pitting at lap 28 and completing the race with a 44 lap second stint on the hard compound, prioritising durability and track position over the performance benefit of a third compound change. These two different approaches represent different strategic priorities; the model minimises cumulative degradation cost through an additional pit stop, while the real-world strategy prioritises minimising pit lane time loss and maintaining track position.

When the model's stochastic parameters are applied to the winner's actual strategy, the expected total race time is 5160.214 seconds, compared to the model's own optimal expected race time of 5132.290 seconds. This difference of approximately 28 seconds suggests that under the model's degradation assumptions, the two-stop strategy would have yielded a marginally faster theoretical race time. However, this advantage must be interpreted carefully, as it does not account for the track position cost of an additional pit stop or the competitive dynamics of the race.

The gap between the winner's actual race time of 5445.519 seconds and the expected race time under his strategy is approximately 285 seconds. Moreover, Norris executed an actual pit stop time of 21.536 seconds, approximately four seconds below the model's expected pit stop penalty of 25.217 seconds for this circuit, a saving that partially reduces the gap but is evidently insufficient to explain this divergence in the race time results. From the race report, the race was defined by sustained wheel-to-wheel competition between Norris and Verstappen from the very first lap, with Verstappen jumping Norris at the start and building a lead before Norris mounted a progressive fightback, eventually passing Verstappen with a DRS-assisted move down the start/finish straight on lap 18. Verstappen himself reported over the radio that his car did not respond to his inputs mid-stint, suggesting both drivers were managing tyre and car behaviour well outside the clean degradation profile the model assumes. Consequently, these combined factors, alongside the systematic beta parameter limitations discussed previously, account for the gap between the theoretical expected time and Norris's actual race outcome.

Italian Grand Prix

The Italian GP presents a closer alignment between the stochastic model's recommendation and the strategy employed by the race winner, Charles Leclerc. The model identifies a one-stop `[H-M]` strategy as optimal, with a pit stop at lap 30, resulting in stint lengths of 30 and 23 laps, respectively. Leclerc executed a one-stop `[M-H]` strategy, pitting at lap 15 with stint lengths of 15 and 38 laps. Both approaches agree on the fundamental strategic choice of a one-stop race utilising the hard and medium compounds, representing a strong alignment in strategic priorities. The divergence lies in the pit window and the compound stint length, where Leclerc opted for a 15 lap medium stint, while the model recommended a 30 lap medium stint. When the model's stochastic parameters are applied to the winner's actual strategy, the expected total race time is 4310.126 seconds, compared to the model's own optimal expected time of 4292.755 seconds, a difference of approximately 17 seconds. This gap is due to the difference in stint lengths between the two strategies, where Leclerc's initial 15 lap stint results in a shorter medium stint and a longer 38 lap hard compound second stint compared to the model's recommendation, and since both strategies utilise the same compound pairing under a constant degradation multiplier, the cumulative degradation cost accumulates differently across the stint lengths, producing this marginal difference in total expected race time rather than reflecting any fundamental strategic superiority of one approach over the other.

Additionally, Leclerc's actual race time of 4480.727 seconds and the expected race time under his strategy produces a gap of approximately 171 seconds, which is considerably smaller than the gap observed at the Dutch Grand Prix. His pit stop was executed in 24.551 seconds, within one second of the model's expected penalty of 25.091 seconds, meaning that the pit stop does not influence the divergence in race-time results. According to the official race report, Monza produced a race that was competitive at the front between Leclerc, Piastri, and Norris, and that was largely defined by tyre management and strategic patience rather than intense wheel-to-wheel pressure on the race leader. Leclerc spent the majority of his long hard stint in clean air at the front, managing his tyres carefully to hold off the McLarens closing on fresher rubber in the final laps. This type of controlled, largely uncontested race leadership is the closest real-world scenario to the isolated clean-air conditions the model assumes, which explains why the Italian Grand Prix produces one of the smallest divergences between theoretical and actual race performance across the circuits analysed.

Spanish Grand Prix

The Spanish Grand Prix presents an interesting partial alignment between the stochastic model's recommendation and the strategy employed by the race winner, Max Verstappen; however, it also shows the biggest gap between expected and actual times. The model identifies a two-stop `[S-M-M]` strategy as optimal, with pit stops at laps 13 and 40, so with stint lengths of 13, 27 and 26 laps, respectively. Verstappen also executed a two-stop strategy, utilising a `[S-M-S]` sequence and pit stops at laps 17 and 44, which results in stint lengths of 17, 27, and 22 laps stint lengths. Both approaches agree on the same type of race, an identical opening soft stint and medium second sprint with the same stint length. The sole divergence lies in the final stint compound selection, where the model opts for a second medium stint prioritising stability and lower cumulative degradation cost, while Verstappen utilised a soft tyre for a final 22 lap stint to defend against Norris's closing gap. When the model's stochastic parameters are applied to the winner's actual strategy, the expected total race time is 4829.065 seconds, compared to the model's own optimal expected time of 4821.248 seconds, a difference of approximately eight seconds. Given that both strategies share an identical structure for two-thirds of the race, this gap is entirely due to the degradation cost of Verstappen's final soft stint, which carries a higher cumulative degradation penalty.

The gap between Verstappen's actual race time of 5300.227 seconds and the expected race time under his strategy is approximately 471 seconds, the largest gap observed across all four races under the clean conditions. Verstappen recorded actual pit stop times of 21.544 seconds and 22.340 seconds, both approximately four and three seconds below the model's expected penalty of 25.197 seconds, representing a combined saving of roughly 7 seconds relative to the model's assumption. While this partially reduces the gap, it is evidently insufficient to explain a divergence of this magnitude. As discussed in the opening of this section, the beta parameters represent a statistical average across multiple seasons rather than race-specific conditions, and physical factors such as track temperature, wind, and fuel load evolution all contribute to this divergence. According to the official race report, the race was characterised by an unusually high level of competitive intensity from the opening lap, with Russell launching from fourth to lead at the start, Norris squeezing Verstappen onto the grass at Turn 1, and Sainz and Leclerc making contact at the same corner. Throughout the race, Verstappen faced sustained pressure from Norris, who reduced the gap to just 2.3 seconds by lap 65, requiring active defensive driving and tactical pace management that inevitably produced lap times above the theoretical baseline. The combined effect of these competitive dynamics, compounded by the beta parameter limitations, explains why the Spanish Grand Prix produces the largest gap between expected and actual race times among the circuits analysed, with the exception of the Monaco GP.

Las Vegas Grand Prix

In the Las Vegas GP, the winner George Russell, opted for a two-stop `[M-H-H]` strategy, pitting at laps 12 and 32, while the model's optimal strategy is a one-stop `[H-M]` strategy with a pit stop at lap 27. These approaches represent different strategic priorities, where the model minimises cumulative pit lane cost by avoiding an additional stop, while Russell's team opted for fresher rubber more frequently to maintain pace on the circuit.

Applying the model's stochastic parameters to Russell's actual strategy yields an expected total race time of 4732.988 seconds, approximately 30 seconds above the model's own optimal of 4702.367 seconds. This difference derives from the additional pit stop in Russell's strategy, because under the model's parameters the degradation savings from two shorter and fresher hard stints do not fully recover the time cost of the extra pit lane entry, making the one-stop solution marginally more efficient on purely theoretical terms.

Russell's actual race time of 4925.969 seconds exceeded the expected race time under his strategy by approximately 193 seconds. His two pit stops were executed in 21.072 and 20.975 seconds, thereby being approximately 5 seconds faster than the model's expected penalty of 25.303 seconds, producing a combined saving of roughly 10 seconds that partially closes the gap but falls far short of explaining it. The official race report published by the F1 team states that the event took place under unusually cool and slippery conditions, with Leclerc suffering a dramatic tyre performance cliff on his mediums by lap 10, illustrating how far actual degradation behaviour deviated from a standard seasonal average. These atypical thermal conditions are precisely the type of circuit-specific variability that the model's beta parameters, calibrated on multi-season averages, cannot reproduce. Combined with Hamilton's charge from tenth on the grid through the field and the sustained competitive dynamics at the front between four drivers, including the winner Russel, the actual lap times accumulated well above the clean-air theoretical baseline, accounting for the remaining divergence.

Monaco Grand Prix

The Monaco Grand Prix requires a distinct methodological treatment compared to the other circuits analysed, due to the exceptional circumstances that defined the 2024 race. As already outlined, Monaco was included despite the red flag stoppage on lap one, because the remainder of the race proceeded under entirely clean conditions. Crucially, all drivers utilised the red flag period to change tyres, converting what would have been a strategic pit stop into a neutralised stoppage, and the race was subsequently run as a one-stop encounter from the restart.

The stochastic model identifies a two-stop `[M-H-M]` strategy as optimal, with pit stops at laps 25 and 53, while Leclerc's actual race strategy was a one-stop, opening with the medium compound for a single lap, followed by the hard compound for the remaining 77 laps. As also already discussed, in the Section 6.3.1 the model's two-stop recommendation at Monaco is a direct consequence of the tyre lifespan constraint, making a one-stop solution mathematically infeasible within the model's framework. In reality, as the official race report confirms, teams exploited the red flag stoppage to execute their mandatory compound change without losing racing time, enabling a one-stop profile that the model cannot replicate under normal pit stop conditions.

When the model's stochastic parameters were applied to Leclerc's actual strategy, the expected total race time was 5676.034 seconds, compared to the model's own optimal expected time of 5618.636 seconds, a difference of approximately 57 seconds. This gap reflects the higher cumulative degradation cost of Leclerc's extremely long 77-lap hard compound second stint, which far exceeds the compound's modelled lifespan and accumulates substantially more degradation than the model's balanced two-stop split, confirming that the one-stop profile is structurally penalised within the model's parameter framework regardless of how it is executed.

The gap between Leclerc's actual race time of 8595.554 seconds and the expected race time under his strategy of 5676.034 seconds amounts to approximately 2919 seconds, almost 49 minutes, which is by far the largest divergence observed across all five races, but entirely explained by the red flag stoppage itself. According to the official race report, the race was interrupted before the completion of lap one and remained neutralised for an extended period while barrier repairs were carried out. Leclerc's recorded pit stop time of 2358.026 seconds, approximately 40 minutes, directly reflects this stoppage period, which is naturally captured in the total race time but bears no relation to the model's pit stop penalty assumption of approximately 24.5 seconds. Consequently, the gap of approximately 2919 seconds between the actual and expected race time is not a reflection of beta parameter limitations or unmodelled racing dynamics, but is entirely attributable to the red flag neutralisation period, which falls entirely outside the scope of a clean-condition optimisation model. For this reason, the Monaco Grand Prix cannot be used to draw meaningful conclusions about the model's predictive accuracy in the same way as the other four circuits analysed, and serves instead as an illustration of the model's boundary conditions when exceptional race interruptions occur.

6.4 Monte Carlo Simulations

The deterministic and stochastic optimisation models developed in Sections 6.2 and 6.1 share a common structural assumption of a race being conducted entirely under green flag conditions, with no external disruptions. While this simplification allows for a rigorous and manageable formulation of the tyre compound and pit stop timing problem, it also creates a notable gap relative to the operational reality of a Formula 1 Grand Prix. A yellow flag is one of the most frequent uncertain events in modern racing, capable of fundamentally altering the cost-benefit calculation of a pit stop within a single lap. This event imposes a mandatory velocity reduction of approximately 30% across the affected laps, which simultaneously penalises the driver's race time,and reduces the effective time loss of entering the pit lane, creating a narrow window of strategic opportunity that teams must identify and exploit in real time. The interaction between the timing of this event and the pre-planned pit window is therefore one of the most critical dimensions of strategic robustness that the preceding model has not yet addressed.

To quantify this interaction, an extended analysis was conducted through a Monte Carlo simulation framework applied directly to the optimal strategies identified in Section 6.2. Rather than re-formulating the MILP model to incorporate the yellow flag event as an additional stochastic variable, the simulation treats the optimal strategy as a fixed input and evaluates its performance across 10000 independent replications. In each replication, the occurrence of a yellow flag, its starting lap, and its duration are drawn from their respective probability distributions, and the total race time is recomputed. This approach therefore tests the robustness of the pre-assigned strategy rather than seeking a new optimal solution.

The probability of a yellow flag occurring in a given replication is derived from the probabilities of a safety car or a virtual safety car occurring, defined in Section 2.4, which are 47% and 30%, respectively. Since the simulation adopts the assumption that at most one such event can occur during a single race, these two events are treated as mutually exclusive, and their summed probabilities represent the probability of a yellow flag occurring, which in this case is 77%. This single occurrence assumption, while a simplification relative to real races that can experience multiple safety car or virtual safety car periods, allows the simulation to isolate the effect of a single disruption on a fixed strategy without compounding the analysis with sequential pit window reevaluations.

The simulation introduces a set of parameters and stochastic inputs that extend beyond those already present in the optimisation model. The degradation and performance parameters (beta parameters) sampled in each replication follow the same distributions presented in Table 5.4, incorporating the recalibration of the hard compound's performance offset defined in Section 6.2, thereby ensuring consistency with the reference stochastic model used throughout this chapter. Table 6.23 summarises these new inputs alongside their distributional assumptions and values.

Table 6.23: Fixed parameters, stochastic inputs, and derived quantities for the Yellow Flag simulation.

Parameter	Description	Value / Formula
w	Pit window half-width in laps	[3, 5, 7]
p_{YF}	Probability of a Yellow Flag occurring	0.77
N	Number of Monte Carlo replications	10000
δ	Velocity reduction during Yellow Flag	30%
P_{pit}^{green}	Pit stop duration under green flag (s)	$\mathcal{T}(13.973,\ 21.814,\ 39.879)$
P_{pit}^{yellow}	Pit stop duration under Yellow Flag (s)	$\mathcal{U}(10,\ 15)$
j	Lap on which the Yellow Flag begins	$\mathcal{U}\{1,\ N_{laps}\}$
d_{YF}	Duration of the Yellow Flag in laps	$\mathcal{U}\{3, 6\}$
$Pit\ window$	Interval within which an early pit stop can be triggered, defined based on the original pit lap suggested by the stochastic model, ℓ_i	$[\ell_i - w,\ \ell_i]$
d_{eff}	Effective Yellow Flag duration in laps	$\min(d_{YF},\ N_{laps} - j + 1)$
ℓ_i^{eff}	Effective length of stint i after possible pit lap adjustment in laps	$\begin{cases} j, & \text{pit changed} \\ \ell_i, & \text{otherwise} \end{cases}$
P_{pit}^{used}	Pit stop time loss applied in the replication (s)	$\begin{cases} P_{pit}^{yellow}, & j \in [\ell_i - w, \ell_i] \\ P_{pit}^{green}, & \text{otherwise} \end{cases}$

The simulation introduces a set of parameters and stochastic inputs that extend beyond those already present in the optimisation model. The degradation and performance parameters (beta parameters) sampled in each replication follow the same distributions presented in Table 5.4, incorporating the recalibration of the hard compound's performance offset defined in Section 6.2, thereby ensuring consistency with the reference stochastic model used throughout this chapter. Table 6.23 summarises these new inputs alongside their distributional assumptions and values.

When there is a yellow flag period, an additional time penalty, Π_{YF}, is incurred due to the mandated velocity reduction, δ, over the effective duration of the disruption, computed as:

$$\Pi_{YF} = t_{base} \cdot \frac{\delta}{1-\delta} \cdot d_{eff} \tag{6.2}$$

The total time spent on track across all stints, T_{track}, is obtained by applying the same stint time formulation utilised throughout the optimisation model, where the cumulative time of a stint of k laps on compound t is given by the Equation 4.8, such that:

$$T_{track} = \sum_{i \in I} StintTime(\ell_i^{eff},\ t_i) \tag{6.3}$$

Finally, the expected total race time for each replication combines the additional penalty if there is a yellow flag period, the time spent on track, and the time spent on the pit lane:

$$T_{total} = T_{track} + \Pi_{YF} + \sum_{i \in I} P_{pit}^{used} \tag{6.4}$$

To enable a consistent comparison across replications and cases, the deviation from the optimal stochastic solution is defined as:

$$\Delta T = T_{total} - T_{otimo} \tag{6.5}$$

Where T_{otimo} denotes the optimal race time obtained from the stochastic optimisation model in Section 6.2 for the corresponding strategy. This metric quantifies the additional time cost incurred in each replication relative to the strategy's expected performance under the stochastic model.

It is important to note that, for the one-stop strategies, the pit lane time is either P_{pit}^{yellow} or P_{pit}^{green}. For the two-stop strategies, one pit stop is always associated with P_{pit}^{green}, whereas the other may be associated with either P_{pit}^{green} or P_{pit}^{yellow}.

Based on this, every replication is classified into one of three mutually exclusive cases. The first case is `No YF`, which means that no yellow flag occurred and the race runs under standard green flag conditions. The second case, `YF in window` represents the case where a yellow flag starts on a lap that falls within the pre-planned pit window, and the driver pits early to exploit the reduced pit stop penalty, still respecting the compound lifespan. Lastly, `YF out of window` represents the case where a yellow flag occurs outside the pit window and the original pit lap is retained. All subsequent analysis in this section is structured around the comparison between the three cases for both the one-stop and two-stop representative circuits.

6.4.1 One-Stop Strategy

The Japanese Grand Prix was selected as the representative one-stop circuit for this analysis. The race covers 53 laps around the Suzuka Circuit, with a `C1-C2-C3` compound allocation. The stochastic model identified as optimal a one-stop strategy `[M-H]` compound sequence as optimal, with the first stint covering 24 laps, and the second covering the remaining 29 laps, thereby fixing the pit stop on lap 24. The pit window is therefore defined over the intervals [21,24], [19,24] and [17,24] for $w = 3$, $w = 5$, and $w = 7$ respectively. Under the stochastic parameters, the model achieved an expected total race time of 4763.475 seconds following this strategy.

The simulation was run for each value of w within a single execution loop, producing 10000 replications per window width. The empirical probability of a yellow flag occurring in a Formula 1 Grand Prix was set at 77%, consequently, the model replicated across the three pit width scenarios a no yellow flag share of approximately 22 to 24%, which matches closely with the theoretical complement of 23%, with a small remaining variation attributable to the sampling noise of the Monte Carlo replications. As the window width increases, the percentage of yellow flags occurring inside the pit window increases, conversely to the yellow flags outside the pit window, which decreases. This shift happens because of the mechanical effect of widening the window. As each additional lap is added to the left boundary of the interval, a larger share of the uniformly distributed yellow flag start lap is captured, which causes out of window replications to be progressively reclassified as inside the window. Even with the highest pit window, the simulations account for less than 12% of cases falling inside the pit window. This confirms that the probability of a yellow flag aligning with the pit window remains low regardless of how the window is defined, simply because the window only covers at most eight laps out of the entire 53 laps of the circuit.

w	Case	Percentage (%)
	`No YF`	24.15
`w=3`	`YF in window`	6.16
	`YF out of window`	69.69
	`No YF`	23.14
`w=5`	`YF in window`	8.64
	`YF out of window`	68.22
	`No YF`	22.28
`w=7`	`YF in window`	11.49
	`YF out of window`	66.23

Table 6.24: Distribution of replications by pit window width, w, and replication scenarios.

The detailed statistical distributions for all three values of w, including the mean, standard deviation (std), minimum, maximum, quartiles, and the fifth and ninety-fifth percentiles (P5 and P95, respectively) for both T_{total} and ΔT, are presented in full in Appendix B. However, for the following analysis, only ΔT is discussed, as both follow identical distributional patterns, and the absolute time is simply shifted by the optimal baseline value.

Starting with the no yellow flag scenario, the boxplot in Figure 6.1 shows that the data cluster tightly at the bottom of the ΔT scale, close to the dashed line that marks the expected total race time obtained from the model, across all three window widths. This is consistent with the fact that under clean racing conditions, the only source of variability is the β_{t0}^{ω}, β_{t1}^{ω}, $P_{pit}^{green,\omega}$, since no yellow flag penalty is applied, which is confirmed by the mean values of ΔT of 10.37, 10.23, and 10.11 seconds, for $w = 3$, $w = 5$, and $w = 7$ respectively, with the standard deviation remaining nearly constant at 6.45, 6.48, and 6.30 seconds. These small differences across the three window widths are expected, given that the no yellow flag case is mechanically independent of w.

Furthermore, the boxplot contains some outliers point under the dashed line, which represent the negative values of the difference between total times, which is also confirmed by Table 6.25, which indicates that in a subset of replications, the combination of sampled beta parameters and pit stop duration was favourable enough to produce a total race time lower than the reference.

w	Mean	Std	Min	Max	Q1	Q2	Q3
`w=3`	10.37	6.45	-9.20	29.14	5.86	10.03	14.86
`w=5`	10.23	6.48	-8.44	30.80	5.53	9.95	14.55
`w=7`	10.11	6.30	-7.56	29.50	5.67	9.95	14.33

Table 6.25: Statistical summary of ΔT for the no yellow flag (`No YF`) scenario across different window widths. Q2 (second quartile) also represents the median.

Secondly, a different pattern emerges from the case where a yellow flag starts inside the pit window. From the same boxplot, the green box, as wee as the red box, are considerably higher than the no yellow flag scenario, which is explained by the yellow flag penalty with the additional uncertainty introduced by the randomly sampled duration and starting lap. Table 6.26 shows that several of these statistics, like the mean of 167.30, 165.91, and 168.41 seconds for the $w = 3$, $w = 5$, and $w = 7$, respectively, the standard deviation, around 42 seconds in all three cases, and the extremes, have a small variability, since a yellow flag falling inside the window carries roughly the same cost independently of the width of the window.

However, an analysis of the quartiles provides a more complex distribution pattern. The first quartile (Q1) decreases from 143.25 seconds at the smallest width to approximately 118 seconds at higher window widths. Conversely, the third quartile (Q3) remains stable between $w = 3$ and $w = 5$, increasing by approximately 26 seconds when the window reaches its maximum width. Consequently, the interquartile range (IQR) expands as w increases, despite the relative stability of the median values (154.60, 155.33, and 157.10). The boxplot illustrates this expansion, with the $w = 7$ box exhibiting a greater vertical span compared to the other configurations, despite all three cases having near identical whiskers and central tendency. Therefore, this can be explained by the fact that increasing the pit window does not only change how costly the penalty of a yellow flag is, but also incorporates a higher range of j laps, and by extension, the effective duration of the yellow flag period; consequently, the IQR expands. Finally, across all three window widths, the box also remains consistently below the yellow flag out of window case at the same width, which reflects the pit loss mitigation provided by P_{pit}^{yellow} when compared to standard green flag penalties.

w	Mean	Std	Min	Max	Q1	Q2	Q3
`w=3`	167.30	41.72	99.98	233.28	143.25	154.60	193.53
`w=5`	165.91	42.45	101.89	236.01	117.33	155.33	191.76
`w=7`	168.41	42.90	101.75	235.01	119.21	157.09	218.48

Table 6.26: Statistical summary of ΔT for the yellow flag inside the pit window (`YF in window`) scenario across different window widths. Q2 also represents the median.

Lastly, the yellow flag out of window case represents the most frequent outcome across all three window widths, as previously established in Table 6.24, and its statistical profile in Table 6.27 reflects the compounded cost of a Yellow Flag disruption without the benefit of a discounted pit stop. In this scenario, the driver incurs the full velocity reduction penalty associated with the yellow flag while still paying the standard green flag pit loss, since the disruption occurs outside the interval within which the pit stop could be brought forward. Consequently, both the boxplot in Figure 6.1 and the underlying statistics indicated in the table highlight that this case produces the highest median cost of the three, exceeding the median of the yellow flag that start inside the pit window case by approximately 11 to 13 seconds at each window width. This difference is consistent with the pit time discount, and its persistence across all three values of $w = 3$, $w = 5$, and $w = 7$ reinforces the conclusion that the benefit of an in window pit stop is not a consequence of a specific window configuration, but rather a structural feature of the simulation.

Moreover, this case displays a pattern of stability across all window widths, observed from the table and boxplot. This consistency arises because the case relies on an exclusionary logic that draws from a consistent subset of replications, regardless of the window definition. Thus, while the pit window dictates the volume of cases classified outside the pit window, the underlying statistical distribution of those outcomes remains unchanged.

w	Mean	Std	Min	Max	Q1	Q2	Q3
`w=3`	173.87	47.65	34.35	257.40	129.91	168.47	208.44
`w=5`	172.94	47.95	33.55	256.66	129.13	166.92	206.92
`w=7`	173.91	47.90	32.12	256.98	130.13	169.25	207.45

Table 6.27: Statistical summary of ΔT for the yellow flag outside the pit window (`YF out of window`) scenario across different window widths. Q2 also represents the median.

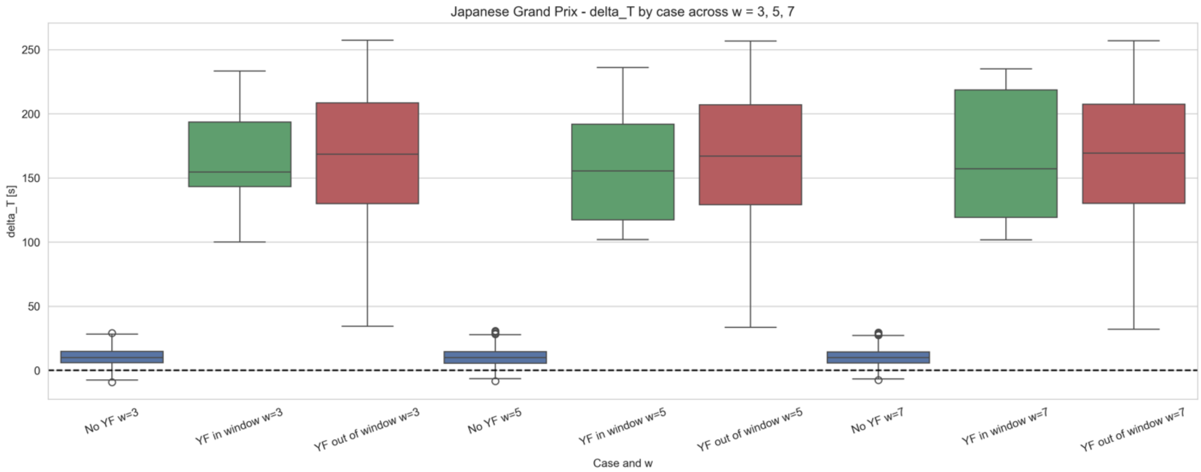

Figure 6.1: Distribution of total race time (ΔT in seconds) across the three pit window scenarios, for the Suzuka circuit.

The scatter plot in Figure 6.2 provides a visualisation of how the occurrence of a yellow flag determines the total race time across all affected replications, and its structural patterns are consistent across all three values of $w = 3$, $w = 5$, and $w = 7$. The green dots represent the yellow flags that occur inside the pit window, which are visibly below the red dots in the same lap range, confirming that pitting under the yellow flag when it falls within the window consistently produces a lower expected total race time than remaining on the original strategy under equivalent conditions. The vertical separation between green and red dots in the window region is consistent with the median differences already quantified in Tables 6.26 and 6.27, on the order of approximately 11 to 14 seconds, since the exact gap for any given pair of points also depends on the specific duration of the yellow flag period, betas, and pit stop timing of each replication. This separation quantifies the effective benefit of the reduced pit stop penalty when the strategic opportunity is correctly exploited.

Furthermore, the red dots display a clear banded horizontal structure across all laps outside the window, which corresponds directly to the four possible values of $d_{YF} \in \{3, 4, 5, 6\}$ laps resulting from the discrete uniform distribution. Since the velocity reduction penalty Π_{YF} is proportional to d_{eff}, each value of d_{YF} produces a distinct level of total race time penalty, and so the four resulting bands are visible throughout the scatter plot. Consequently, this confirms that the duration of the yellow flag is a primary driver of total race time variability outside the pit window.

Finally, the red dots in the final laps of the race, approximately from lap 49 onwards, drop sharply to the lowest observed values of T_{total}, reaching the minimum value of approximately 4797 seconds. This occurs because when the yellow flag starts in the final laps of the race, d_{eff} is limited by the remaining lap count, and the velocity reduction penalty becomes progressively smaller as fewer laps remain to apply it, rather than disappearing outright, since a yellow flag starting on the very last laps can still contribute a small but non-zero penalty.

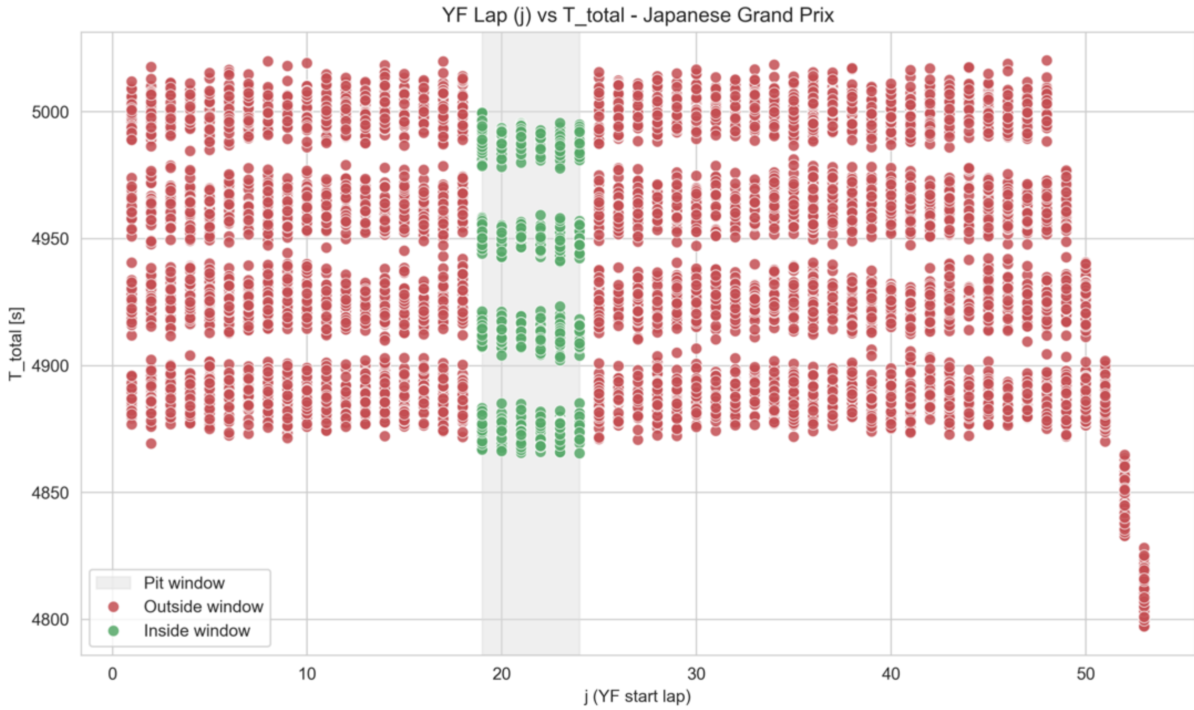

Figure 6.2: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Suzuka circuit when $w = 5$

.

Figure 6.2 plots every replication where a yellow flag occurs with the x-axis showing lap j and the y-axis showing T_{total}, coloured by inside and outside window classification, with the pit window band shaded over laps 19 to 24, for $w = 5$. The equivalent scatter plots for $w = 3$ and $w = 7$ follow the same structural pattern, differing only in the width of the shaded pit window band and the number of green dots it contains, and are presented in the Figures B.2 and B.3 in the Appendix B.

6.4.2 Two-Stop Strategy

To extend the analysis to a two-stop strategic environment, the Dutch Grand Prix was selected. According to the stochastic optimisation results presented in Table 6.10, the optimal strategy for this circuit is a two-stop `[S-M-M]` sequence, with pit stops scheduled at laps 14 and 43, resulting in stint lengths of 14, 29, and 29 laps for the `soft`, `medium`, and `medium` compounds, respectively. The expected optimal total race time under clean conditions is 5132.29 seconds. Additionally, for this case, two pit windows are defined around each planned pit stop lap, and for a window width of $w = 3$, the first pit window is on laps 11 to 14 and the second on laps 40 to 43. Consequently, for $w = 5$, the windows extend to laps nine to 14 and 38 to 43, respectively. Finally, for $w = 7$, the first window covers laps seven to 14 and the second covers laps 36 to 43.

As already established in the one-stop analysis, the no yellow flag share remains close to the theoretical 23% regardless of window width, and widening the window increases the share of yellow flags falling inside it, while decreasing those falling outside. However, in the two-stop strategy this increase is considerably more pronounced, rising from 8.33% at $w = 3$ to 16.94% at $w = 7$, which is a direct consequence of the strategy having two pit windows instead of one. So, the combined coverage of both intervals expands from 8 laps at $w = 3$, to 12, and then to 16 laps for $w = 7$, out of the 72 laps of the circuit. Therefore, as this combined coverage grows, so does the probability of yellow flags starting on a lap that is inside the pit window proposed, which explains why the in window case nearly doubles across the tested range, a rate of increase not observed in the single window case.

w	Case	Percentage (%)
$w = 3$	`No YF`	23.07
	`YF in window`	8.33
	`YF out of window`	68.60
$w = 5$	`No YF`	22.94
	`YF in window`	13.17
	`YF out of window`	63.89
$w = 7$	`No YF`	22.96
	`YF in window`	16.94
	`YF out of window`	60.10

Table 6.28: Distribution of replications by pit window width, w, and replication scenarios.

Like in the previous analysis, the full results for both T_{Total} and ΔT are in the Appendix B.

Regarding the case where no yellow flag occurs during the race, the mean ΔT remains close to 29 seconds and the standard deviation close to 12 seconds across all three window widths, confirming, as expected, that this case is independent of w. However, both values are considerably higher than those observed for the equivalent case in the one-stop strategy, where the mean and standard deviation remained close to ten and six seconds, respectively. This gap can be primarily attributed to the pit stop time modelling assumption, since the same sampled pit stop duration is applied to both pit stops within a two-stop race, which amplifies its contribution to the overall variability compared to a one-stop race where it is applied only once.

w	Mean	Std	Min	Max	Q1	Q2	Q3
$w = 3$	29.29	11.95	0.44	67.26	20.37	28.26	37.09
$w = 5$	29.15	12.40	-2.55	67.89	19.87	27.80	38.13
$w = 7$	28.90	11.93	-5.35	68.41	20.27	28.27	36.90

Table 6.29: Statistical summary of ΔT for the no yellow flag (`No YF`) scenario across different window widths. Q2 also represents the median.

Turning to the case where the yellow flag period starts inside one of the two pit stop windows, this scenario proves considerably more stable across the three window widths than its one-stop counterpart. While the median increases by a modest and consistent margin at each step, and the lower quartile drifts upward, both by around two seconds. However, the major instability occurs in the upper quartile, increasing by approximately ten seconds from $w = 3$ to $w = 5$ and then stabilising when transitioning to

$w = 7$. Therefore, this trend demonstrates that while variability increases with window width in a single-stop strategy, the same effect is not present in a two-stop strategy, where the impact remains limited regardless of width.

This increased stability is a direct result of utilising two pit windows instead of one. Each widening step adds laps to both intervals simultaneously, effectively doubling the pool of replications that shift into the in window classification at every increment of $w = 3$, $w = 5$, or $w = 7$. Therefore, this broader intake of replications reduces the sampling noise, which explains why the median and lower quartile remain almost steady.

w	Mean	Std	Min	Max	Q1	Q2	Q3
$w = 3$	149.63	33.91	86.15	216.07	120.04	150.75	173.97
$w = 5$	152.04	34.85	89.19	220.90	120.48	153.98	183.60
$w = 7$	152.88	34.38	88.20	221.13	122.91	155.72	184.26

Table 6.30: Statistical summary of ΔT for the `YF in window` scenario across different window widths. Q2 also represents the median.

Lastly, for the same reasons established in the one-stop analysis, the yellow flag start outside of the pit window case remains highly stable across all three values of width, because this classification draws from a consistent subset of replications regardless of how the window itself is defined, since widening $w = 3$, $w = 5$, or $w = 7$ only reclassifies replications at the margin into the in window case, without altering the underlying distribution of those that remain outside it.

However, unlike the one-stop strategy, the gap between the two yellow flag cases narrows considerably as $w = 3$, $w = 5$, and $w = 7$ increases, where pitting early during the yellow flag saves approximately nine seconds relative to the out of window case at $w = 3$, while this advantage drops to approximately four and a half seconds at $w = 5$ and stabilises at around five seconds at $w = 7$.

This narrowing is consistent with a measurable increase in track time cost associated with anticipating the pit stop under a wider window. As the window widens, a yellow flag is more likely to trigger a pit stop further from the originally optimised lap, which shifts the surrounding stint lengths away from their optimal balance and introduces a small additional degradation cost. This growing cost partially offsets the benefit of the reduced yellow flag pit stop penalty, which explains why the advantage of pitting early narrows and then stabilizes as w increases.

w	Mean	Std	Min	Max	Q1	Q2	Q3
$w = 3$	159.37	38.54	33.46	243.76	129.02	159.53	192.16
$w = 5$	158.49	38.33	33.37	241.67	129.00	158.46	190.42
$w = 7$	159.28	38.48	34.27	243.41	129.90	160.71	191.50

Table 6.31: Statistical summary of ΔT for the `YF out of window` scenario across different window widths. Q2 also represents the median.

Overall, the boxplot analysis reveals a systematically higher mean and greater variability across all three cases in the two-stop strategy relative to the one-stop strategy. This difference is primarily explained by a structural feature of the pit stop time modelling, since a single pit stop duration is sampled from the triangular distribution and applied to every pit stop within the same race, and a two-stop strategy incurs this same random value twice. Consequently, the variability contributed by the pit stop time to the total race time is amplified in the two-stop case, which was confirmed directly by comparing the standard deviation of the applied pit stop time between the two circuits, where approximately six seconds for the one-stop strategy against approximately 12 seconds for the two-stop strategy, closely matching a twofold increase.

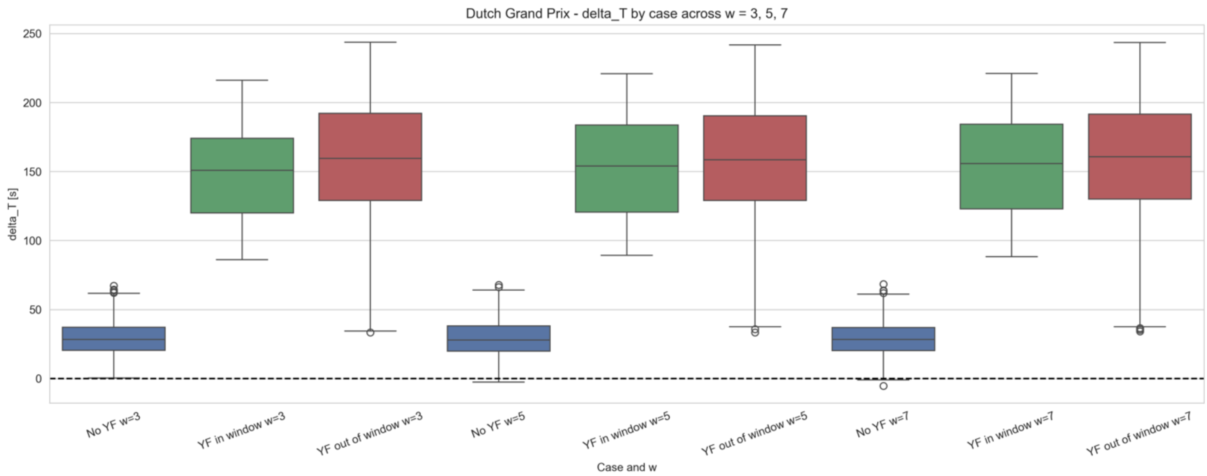

Figure 6.3: Distribution of total race time (ΔT in seconds) across the three pit window scenarios, for the Zandvoort circuit.

The scatter plot in Figure 6.4 illustrates every replication in which a yellow flag occurred, plotting the total race time T_{total} against the yellow flag start lap j, for $w = 5$. Within each of the two shaded pit window bands, green and red points are not uniform across the full vertical range. Thus, in the upper part of each band, the green dots remain visibly below the red ones, reflecting the gap already identified between the third quartiles of the in window and out of window cases. Conversely, in the lower part of each band, the green and red points occupy a very similar range of values, which is consistent with the first quartiles of both cases being close to one another. However, the visual boundaries for $d_{YF} \in \{3, 4, 5, 6\}$ laps appear less distinct than in the one-stop case, with increased blending between adjacent bands. This is consistent with the greater overall variability previously observed for the two-stop strategy.

Finally, towards the final laps of the race, the red points also drop sharply to the lowest observed values of total race time, mirroring the effect that also occurred in the one-stop case. Since the effective duration of the yellow flag period is capped by the number of laps remaining in the race, a yellow flag starting close to the finish applies a progressively smaller velocity reduction penalty, even though it does not disappear entirely.

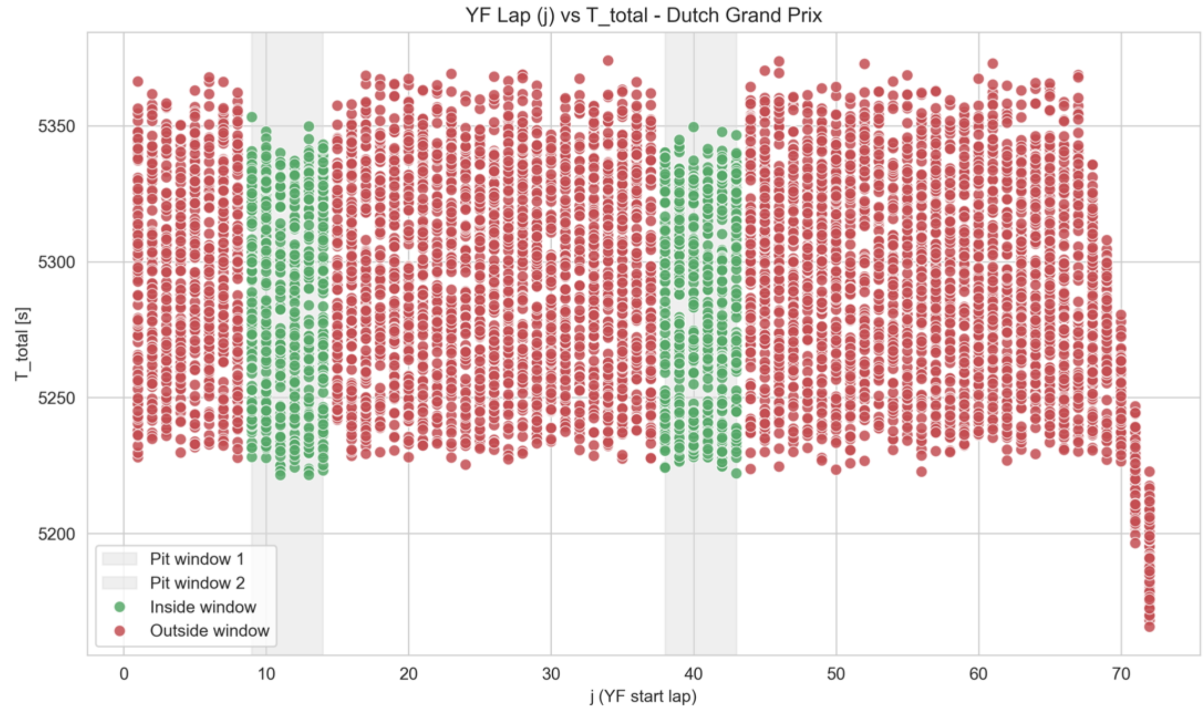

Figure 6.4: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Zandvoort circuit when $w = 5$

.

6.4.3 Synthesis of Yellow Flag Robustness Results

To conclude, when applying the yellow flag disruption to the optimal strategies identified in Section 6.2, both circuits show that only a small share of replications, at most 12% even at $w = 7$, exhibit a yellow flag falling within the pit window, confirming that the alignment between an unplanned disruption and a pre-planned window remains inherently unlikely regardless of strategy type. However, the two representative circuits respond differently to the widening of this window. In the one-stop strategy, increasing *w* produces a clear expansion of the interquartile range for the yellow flag in-window case, because a wider window incorporates a broader range of possible yellow flag starting laps *j*, and by extension a wider range of effective yellow flag durations, without a second pit window to average this variability out. However, in the two-stop strategy, this same widening instead produces a comparatively stable distribution across $w = 3$, $w = 5$, and $w = 7$, because the presence of two pit windows means that each increment in *w* enlarges both intervals simultaneously, doubling the pool of replications reclassified as in-window at each step and thereby reducing the sampling noise that would otherwise be visible in a single-window case. Additionally, the two-stop strategy exhibits a systematically higher mean and greater variability across all three cases when compared to the one-stop strategy, because the pit stop duration is sampled once per race and applied to both stops, effectively doubling its contribution to the overall variance relative to a strategy with a single pit stop. Overall, these results indicate that the number of planned pit stops, rather than the width of the pit window itself, is the more decisive factor in determining how a strategy withstands an unplanned yellow flag disruption.

Chapter 7

Conclusions

This thesis set out to address the gap identified in the literature review, which is the persistence of deterministic assumptions in Operations Research approaches to Formula 1 race strategy, despite the inherently volatile nature of tyre degradation and pit stop execution on race day. To address this, the work developed a two-stage stochastic MILP formulation capable of treating key performance parameters as random variables rather than fixed point estimates, and further complemented it with a Monte Carlo simulation framework to assess strategy robustness under yellow flag disruption periods, an operational uncertainty not previously integrated into this class of optimisation models. In doing so, the thesis sought to answer two related but distinct questions raised throughout Chapters 1 to 3: first, what is the optimal strategy to plan for under normal, uninterrupted racing conditions, and second, how robust does that plan remain once the uncertainty inherent to racing, both in the model's own parameters and in external disruptions such as yellow flags, is taken into account.

7.1 Achievements

In the first part of this thesis, a deterministic MILP model was successfully developed to determine optimal tyre sequence, stint lengths, and pit stop timings for all 25 Grand Prix circuits. However, an initial application of literature-derived degradation and performance coefficients revealed a calibration bias, where the hard compound was never selected, which contradicts real Formula 1 racing patterns. Through a systematic sensitivity analysis, this discrepancy was traced to an excessively high pace penalty assigned to the hard compound. Therefore, by reducing this parameter by 25%, the parameter was identified as the threshold at which the model's output realigned with the tyre hierarchy observed between the 2022 and 2024 seasons. Ultimately, the deterministic results demonstrated that the choice between one-stop and two-stop strategies is never governed by any single variable in isolation, but rather emerges from the simultaneous interaction of the circuit's base lap time, its total race distance, the pit stop time loss, and Pirelli's specific compound allocation for that weekend, with different factors becoming dominant under varying circuit conditions.

Secondly, the deterministic model was extended into a two-stage stochastic programming framework, achieved by evaluating each fixed first-stage strategy across 20 discrete scenarios sampled from the empirical distributions of tyre degradation and pit stop duration. By applying the recalibration logic derived from the deterministic analysis, the stochastic model generated a tyre usage distribution that preserved the established compound hierarchy while successfully capturing the flexibility introduced by parameter uncertainty. The explicit modelling of uncertainty was shown to yield a consistent advantage, quantified by a negative Value of the Stochastic Solution across all circuits. These results confirm that strategies derived from full scenario-based evaluations systematically outperform those optimised against single, average case estimates. Furthermore, the validity of these stochastic strategies was assessed against two independent benchmarks: Pirelli's official pre-race recommendations, and the strategies actually executed by 2024 race winners. In instances where divergences occurred, these were not attributed to model flaws, but were instead traced to explicit and identifiable causes, such as the model's fixed tyre lifespan constraints, the simplifying assumption of a constant degradation multiplier across stints, and the inability of season-averaged beta parameters to capture the specific competitive dynamics, track conditions, or driver behaviour of an individual race.

Lastly, the final objective of assessing strategy robustness during yellow flag disruptions was addressed using a Monte Carlo simulation framework built upon the stochastic model's optimal strategies. This framework was run across 10000 replications for representative one-stop and two-stop circuits. The simulations confirmed that a yellow flag occurring within a pre-planned pit window consistently and measurably reduces total race time relative to an equivalent disruption falling outside that window, validating the strategic logic of anticipating a pit stop to exploit the reduced pit lane penalty. However, it was also observed that widening the pit window yields only a modest increase in the probability of capturing this benefit, since the window itself spans a small fraction of the total race distance. Furthermore, the two-stop strategies exhibited consistently higher variance in total race time outcomes than their one-stop counterparts. This was identified as a direct structural consequence of the simulation design, where the same sampled pit stop duration was applied to both stops within a single race replication.

Overall, these results confirm that the two dimensions of race strategy identified at the outset of this thesis, determining the best plan under normal conditions and assessing its robustness once uncertainty is introduced, were both addressed through a coherent and internally consistent modelling framework, with each phase of the analysis building directly on the calibration and validation established in the preceding one.

7.2 Future Work

Several of the proposed refinements target how tyre degradation and lifespan are parameterised within the model, since the current formulation relies on generic, literature-derived values that do not fully reflect the diversity of circuit conditions across a season. A first extension would be to calibrate the degradation parameters individually for each circuit and for each Pirelli compound grade, rather than relying on a single set of values pooled across all circuits, since factors such as track abrasiveness, lap length, and the proportion of high-speed straights versus technical corners vary substantially between circuits and jointly influence how a given compound degrades over the course of a race. A second extension concerns the way tyre lifespan is defined. Rather than imposing a fixed maximum stint length for each compound, lifespan could instead be modelled as an interval or range, reflecting the combined influence of lap length and circuit-specific durability effects, as well as the fact that this limit is not a strict cutoff, but rather a recommended threshold driven primarily by safety and performance considerations. Strategists occasionally choose to exceed it when doing so proves more advantageous overall. Modelling lifespan as a range rather than a fixed value would allow the optimisation to capture this flexibility, rather than excluding such strategies by construction.

Beyond calibration and lifespan, several additional refinements could further improve the realism of the degradation model. One such extension would be to incorporate driver-specific degradation coefficients, reflecting the fact that individual driving style and tyre management can meaningfully affect how quickly a given compound degrades, rather than assuming a single generic driver behaviour across all strategies. Another possible refinement involves the stint multiplier, currently held constant ($M_i = 1$) across all stints. Allowing this multiplier to vary would make it possible to capture the strategic value of stint order, once empirical data on fuel load evolution and track evolution over the course of a race becomes available. Finally, the pit stop time loss, currently modelled using a single global triangular distribution across all circuits, could instead be defined on a circuit-specific basis, provided sufficiently granular pit lane data becomes available, since factors such as pit lane length and layout vary meaningfully from one circuit to another.

A further set of extensions concerns the uncertain events that could occur in a race, such as the treatment of weather conditions and race disruptions, both of which were deliberately left outside the scope of this thesis, due to the lack of reliable data on these events. The current model could be extended to rain affected races by incorporating intermediate and full wet tyre compounds, which follow distinct degradation and performance profiles from the slick compounds considered here. Similarly, the treatment of yellow flag disruptions could be made more sophisticated in two ways. First, rather than only assessing the robustness of a given strategy to yellow flags after the fact through Monte Carlo simulation, their occurrence could instead be endogenised directly into the stochastic optimisation itself, allowing the model to actively plan around this uncertainty rather than merely evaluate it. Second, the Monte Carlo simulation could be extended to allow multiple safety car or virtual safety car events within a single race replication, relaxing the single-occurrence assumption adopted in this thesis, and to model safety car and virtual safety car periods as separate events with distinct probabilities and time-loss profiles, rather than aggregating them into a single yellow flag category.

Another direction for future work concerns the competitive and strategic realism of the model. The current formulation optimises a single driver's strategy in isolation, without accounting for the behaviour of competitors on track. A more online, reactive version of the model could instead account for rival strategic behaviour, adjusting a driver's own pit timing and tyre age in response to the decisions made by competitors, for instance by pitting earlier to remain synchronised with the field rather than optimising purely against a fixed, pre-race plan. The optimisation model itself could be further perfected by incorporating game-theoretic elements, relaxing the single-driver, isolated strategy assumption altogether and treating race strategy not as a single optimisation problem but as an interactive one shaped by the anticipated responses of other competitors.

Finally, a natural direction for future work is to extend the model to the regulatory framework introduced in 2026. Rather than the aerodynamic changes themselves, the more relevant extension concerns the new energy management dimension, since drivers can now harvest and deploy battery power strategically over a lap, adding a further layer of decision-making alongside tyre choice and pit stop timing. Incorporating energy management into the optimisation framework would allow the model to better reflect the strategic landscape under the new regulations, though this was not feasible within the scope of this thesis due to the absence of historical race data at the time of modelling. Once sufficient data becomes available, the model could be extended to jointly optimise tyre and energy management strategy.

Bibliography

[1] C. L. W. Choo. Real-Time Decision Making in Motorsports: Analytics for Improving Professional Car Race Strategy. 2015.

[2] F. I. de l'Automobile. 2025 FORMULA ONE SPORTING REGULATIONS, apr 2025. URL https://www.fia.com/regulation/category/110. Issue 5.

[3] J. Thorne. What Is Grid Position in Racing? F1 Grid & Pole Position, may 2026. URL https://worldofspeed.org/f1/what-is-grid-position-in-racing/. Accessed: 2026-06-12.

[4] M. Mills. How many laps does each Formula 1 race have?, jan 2025. URL https://motorsporttickets.com/blog/how-many-laps-does-each-formula-1-race-have/. Accessed: 2026-06-12.

[5] Pirelli. FORMULA 1® AND PIRELLI. URL https://www.pirelli.com/tyres/en-ww/motorsport/car/formula-1. Accessed: 2026-06-10.

[6] The F1 Explained. Why F1 Tires Are Designed to Fail [C1 to C6], aug 2025. URL https://www.youtube.com/watch?v=c8IJQ0vMMgw.

[7] The beginner's guide to F1 tyres. URL https://www.formula1.com/en/latest/article/the-beginners-guide-to-formula-1-tyres.61SvF0Kfg29UR2SPhakDqd. Accessed:2026-06-12.

[8] M. P. Q. García. Optimization of Pit Stop Strategies in Formula 1 Racing: A Data-Driven Approach. 2024.

[9] F. Aguad and C. Thraves. Optimizing pit stop strategies in Formula 1 with dynamic programming and game theory. *European Journal of Operational Research*, 319(3):908–919, Dec. 2024.

[10] O. F. C. Heine and C. Thraves. On the optimization of pit stop strategies via dynamic programming. *Central European Journal of Operations Research*, 31(1):239–268, 2023.

[11] E. Spencer. F1 Explained: Why Has the Pitlane Speed Limit Been Changed for the Dutch Grand Prix?, aug 2025. URL https://www.autoevolution.com/news/f1-why-has-the-pitlane-speed-limit-been-changed-for-the-dutch-grand-prix-256727.html. Section: Auto Motorsport; Accessed: 2026-06-10.

[12] F1 strategy explained: What's an undercut, overcut, a DRS train and more, jun 2026. URL https://www.motorsport.com/f1/news/f1-strategy-explained-whats-an-undercut-overcut-a-drs-train-and-more/10500434/. Accessed: 2026-06-12.

[13] A. Heilmeier, M. Graf, and M. Lienkamp. A Race Simulation for Strategy Decisions in Circuit Motorsports: 21st IEEE International Conference on Intelligent Transportation Systems, ITSC 2018. *2018 IEEE Intelligent Transportation Systems Conference, ITSC 2018*, pages 2986–2993, dec 2018.

[14] F1 2026: come funzionerà il pulsante ≪Override≫. Sorpassi più semplici a velocità altissime. URL https://www.formula1.it/news/27700/1/f1-2026-come-funzionera-il-pulsante-override-sorpassi-piu-semplici-a-velocita-altissime. Accessed: 2026-06-12.

[15] J. T. Westbrook. Here's How Virtual Safety Cars Work In Formula One, may 2018. URL https://www.jalopnik.com/heres-how-virtual-safety-cars-work-1826237975/. Accessed: 2026-06-12.

[16] A. Heilmeier, A. Thomaser, M. Graf, and J. Betz. Virtual Strategy Engineer: Using Artificial Neural Networks for Making Race Strategy Decisions in Circuit Motorsport. *Applied Sciences*, 10(21): 7805, jan 2020.

[17] A. Bonomi, E. Turri, and G. Iacca. Evolutionary F1 Race Strategy. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation*, pages 1925–1932, Lisbon, Portugal, July 2023. ACM. Accessed: 2026-04-22.

[18] J. Bekker and W. Lotz. Planning Formula One race strategies using discrete-event simulation. *Journal of the Operational Research Society*, 60(7):952–961, 2009.

[19] Formula 1 World Championship (1950 - 2024), . URL https://www.kaggle.com/datasets/rohanrao/formula-1-world-championship-1950-2020.

[20] F. I. de l'Automobile. 2022 FORMULA ONE Technical REGULATIONS, feb 2021. URL https://www.fia.com/regulation/category/110. Issue 3.

[21] F. I. de l'Automobile. 2021 FORMULA ONE Technical REGULATIONS, jun 2021. URL https://www.fia.com/regulation/category/110. Issue 10.

[22] F. I. de l'Automobile. 2023 FORMULA ONE Technical REGULATIONS, dec 2022. URL https://www.fia.com/regulation/category/110. Issue 4.

[23] F. I. de l'Automobile. 2024 FORMULA ONE Technical REGULATIONS, oct 2024. URL https://www.fia.com/regulation/category/110. Issue 8.

[24] M. Schofield. F1 tire rules for the 2025 season, explained, feb 2025. URL https://www.sbnation.com/formula-one/2025/2/17/24366939/f1-tire-rules-2025-season. Accessed: 2026-06-12.

[25] F1 Tyres Explained 2024 | Blackcircles.com, 2024. URL https://www.blackcircles.com/news/f1-tyres-explained. Accessed: 2026-06-12.

[26] Pole Position Experts. How Long Do F1 Hard Compound Tires Last?, dec 2025. URL https://www.youtube.com/watch?v=gczL9yZ_1kk. Accessed: 2026-06-12.

[27] How Long Do F1 Tires Last? (Full Explanation) - FLOW RACERS, aug 2023. URL https://flowracers.com/blog/how-long-do-f1-tires-last/. Accessed: 2026-06-12.

[28] L. Pretorius. How Many Laps Do F1 Tires Last?, oct 2021. URL https://onestopracing.com/how-many-laps-do-f1-tires-last/. Accessed: 2026-06-12.

[29] All compounds on track to end the season, 2024. URL https://press.pirelli.com/all-compounds-on-track-to-end-the-season/.

[30] D. Carter. GP Italy 2024 - Possible race strategies: risk of graining opens up to different scenarios at Monza, Sept. 2024. URL https://scuderiafans.com/gp-italy-2024-possible-race-strategies-risk-of-graining-opens-up-to-different-scenarios-at-monz Accessed: 2026-06-22.

[31] D. Grangier. Quelles stratégies possibles pour le Grand Prix de Grande-Bretagne 2024 ?, July 2024. URL https://www.autohebdo.fr/actualites/f1/quelles-strategies-possibles-pour-le-gp-de-grande-bretagne-2024.html. Accessed: 2026-06-22.

[32] D. Carter. 2024 F1 Dutch GP - Possible race strategies at Zandvoort: watch out for the overcut, Aug. 2024. URL https://scuderiafans.com/2024-f1-dutch-gp-possible-race-strategies-at-zandvoort-watch-out-for-the-overcut/. Accessed: 2026-06-22.

[33] LC1597. [Pirelli] Possible race strategies for Hungarian GP 2024, July 2024. URL https://www.reddit.com/r/formula1/comments/1e8h2hi/pirelli_possible_race_strategies_for_hungarian_gp/. Accessed: 2026-06-22.

[34] S. Fans. Pirelli race strategies for F1 in Monaco: which compound to start with? | 2024 Monte Carlo GP, May 2024. URL https://scuderiafans.com/pirelli-race-strategies-for-f1-in-monaco-which-compound-to-start-with-2024-monte-carlo-gp/. Accessed: 2026-06-22.

[35] S. Fans. How many pit stops in 2024 F1 Saudi Arabian GP? Pirelli details strategies for Jeddah race, Mar. 2024. URL https://scuderiafans.com/how-many-pit-stops-in-2024-f1-saudi-arabian-gp-pirelli-details-strategies-for-jeddah-race/. Accessed: 2026-06-22.

[36] S. Fans. How many pit stops in F1 Miami GP? Here are the possible race strategies, May 2024. URL https://scuderiafans.com/

how-many-pit-stops-in-f1-miami-gp-here-are-the-possible-race-strategies/. Accessed: 2026-06-22.

[37] F1 Stint Data with Aggression Scores, . URL https://www.kaggle.com/datasets/akashrane2609/f1-stint-data-with-aggression-scores.

[38] Norris fights back against Verstappen to win Dutch GP, 2024. URL https://www.formula1.com/en/latest/article/norris-fights-back-against-verstappen-to-end-home-heros-run-of-dutch-gp-wins.3IPKJnhPnBEjWGAfrGeZQo. Accessed: 2026-06-22.

[39] Leclerc thrills the Tifosi as he triumphs at Monza, 2024. URL https://www.formula1.com/en/latest/article/leclerc-thrills-the-tifosi-to-triumph-at-monza-ahead-of-piastri-and-norris.1aiYZF3rWZp2Q9yQtcuvqV. Accessed: 2026-06-22.

[40] 2024 Spanish Grand Prix race report and highlights: Verstappen just holds off Norris to win super-tight Spanish Grand Prix as Hamilton claims first podium of 2024 | Formula 1®, 2024. URL https://www.formula1.com/en/latest/article/verstappen-holds-of-norris-challenge-to-seal-victory-at-the-spanish-grand.7a5UhXMAtkhbslSDNiPVes. Accessed: 2026-06-22.

[41] Verstappen crowned champion as Russell heads Mercedes 1-2, 2024. URL https://www.formula1.com/en/latest/article/verstappen-crowned-champion-as-russell-heads-mercedes-1-2-in-las-vegas.3TZdhZEthj8sADUPU8p6Ni. Accessed: 2026-06-22.

[42] 2024 Monaco Grand Prix race report and highlights: Leclerc clinches long-awaited home win in Monaco ahead of Piastri and Sainz after early drama | Formula 1®, 2024. URL https://www.formula1.com/en/latest/article/leclerc-clinches-long-awaited-home-win-in-monaco-ahead-of-piastri-and-sainz.4EmeQeb2P27aS8yGIVMOEO. Accessed: 2026-06-22.

Appendix A

Pirelli Pre-Race Strategy Infographics

The figures presented in this appendix correspond to the official pre-race strategy infographics provided by Pirelli for the 2024 Formula 1 season, for the seasons analysed on 6.3.1. These figures outline the recommended tire compound combinations and estimated pit stop windows for each circuit, as well the c-grade allocation t the race. The Tables 6.15, 6.16, 6.17, 6.18, 6.19, 6.20, and 6.21 summarise the strategic data extracted from these visual resources.

Italian Grand Prix

Figure A.1: Pirelli's pre-race strategy recommendations for the Italian Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [30]

British Grand Prix

Figure A.2: Pirelli's pre-race strategy recommendations for the British Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [31]

Dutch Grand Prix

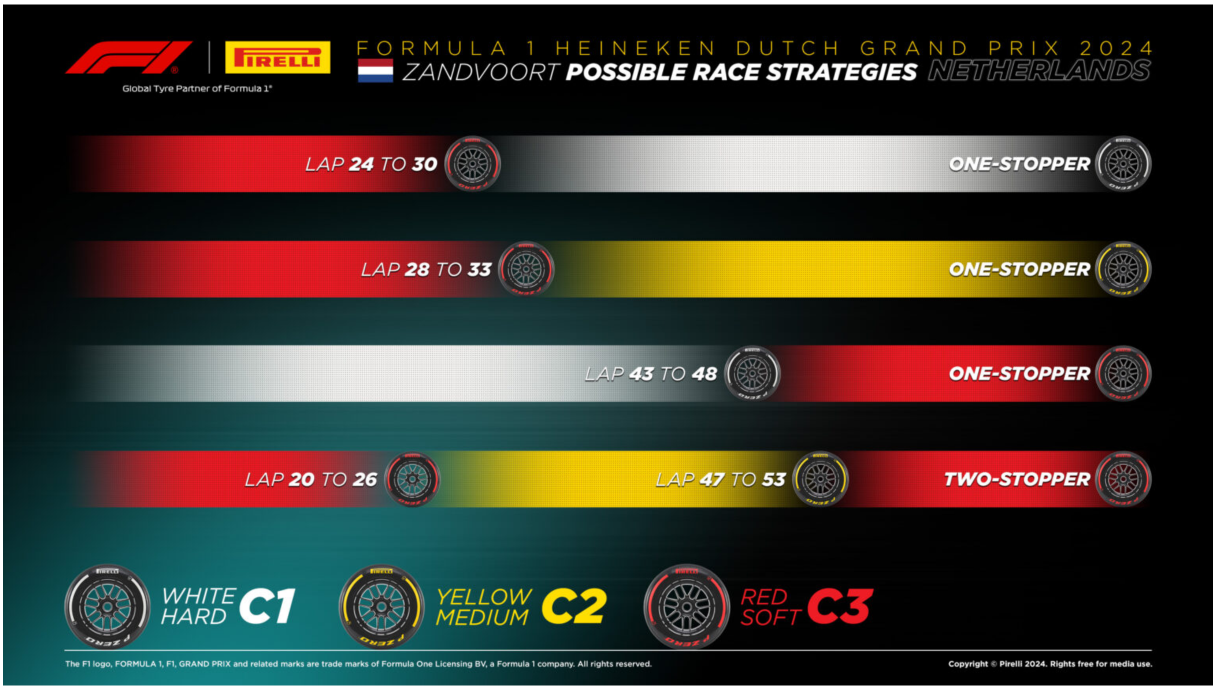

Figure A.3: Pirelli's pre-race strategy recommendations for the Dutch Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [32]

Hungarian Grand Prix

Figure A.4: Pirelli's pre-race strategy recommendations for the Hungarian Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [33]

Monaco Grand Prix

Figure A.5: Pirelli's pre-race strategy recommendations for the Monaco Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [34]

Saudi Arabia Grand Prix

Figure A.6: Pirelli's pre-race strategy recommendations for the Saudi Arabia Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [35]

Miami Grand Prix

Figure A.7: Pirelli's pre-race strategy recommendations for the Miami Grand Prix 2024, illustrating the proposed pit windows and compound sequences. Source: [36]

Appendix B

Monte Carlo Simulation Results

B.1 One-Stop Strategy

The Table B.1 and Figure B.1 provide the statistical summary and the distribution of the total race time (T_{total}) and its deviation (ΔT) for the one-stop strategy (Japanese Grand Prix) simulation across all window widths. Additionally, the scatter plots in Figures B.2 and B.3 illustrate the sensitivity of the total race time to the yellow flag start lap (j) for the extreme window widths of $w = 3$ and $w = 7$.

w	Group	Variable	Mean	Std	Min	Max	Q1	Q2	Q3	P5	P95
$w = 3$	Overall	T_{total}	4897.45	81.04	4754.28	5020.87	4812.85	4918.11	4963.03	4768.42	5003.37
	Overall	ΔT	149.63	81.04	-9.20	257.40	49.38	154.64	199.56	4.95	239.90
	No YF	T_{total}	4773.84	6.45	4754.28	4792.61	4769.34	4773.51	4778.33	4763.73	4784.93
	No YF	ΔT	10.37	6.45	-9.20	29.14	5.86	10.03	14.86	0.25	21.45
	YF in win.	T_{total}	4930.77	41.72	4863.46	4996.75	4906.73	4918.07	4957.01	4870.56	4990.67
	YF in win.	ΔT	167.30	41.72	99.98	233.28	143.25	154.60	193.53	107.09	227.19
	YF out of win.	T_{total}	4937.34	47.65	4792.83	5020.87	4893.38	4931.94	4971.91	4875.54	5005.27
	YF out of win.	ΔT	173.87	47.65	34.35	257.40	129.91	168.47	208.44	112.07	241.80
$w = 5$	Overall	T_{total}	4898.15	80.01	4755.04	5020.13	4842.49	4917.43	4962.19	4768.22	5003.78
	Overall	ΔT	134.68	80.01	-8.44	256.66	79.01	153.96	198.71	4.74	240.30
	No YF	T_{total}	4773.70	6.48	4755.04	4794.28	4769.01	4773.42	4778.02	4763.63	4784.85
	No YF	ΔT	10.23	6.48	-8.44	30.80	5.53	9.95	14.55	0.15	21.38
	YF in win.	T_{total}	4929.38	42.45	4865.37	4999.48	4880.80	4918.81	4955.23	4870.89	4991.49
	YF in win.	ΔT	165.91	42.45	101.89	236.01	117.33	155.33	191.76	107.42	228.02
	YF out of win.	T_{total}	4936.41	47.95	4797.03	5020.13	4892.60	4930.39	4970.39	4875.51	5005.83
	YF out of win.	ΔT	172.94	47.95	33.55	256.66	129.13	166.92	206.92	112.03	242.36
$w = 7$	Overall	T_{total}	4900.26	79.65	4755.92	5020.45	4853.57	4918.86	4963.23	4768.61	5002.99
	Overall	ΔT	136.79	79.65	-7.56	256.98	90.10	155.39	199.75	5.13	239.51
	No YF	T_{total}	4773.58	6.30	4755.92	4792.97	4769.15	4773.42	4777.81	4763.39	4784.50
	No YF	ΔT	10.11	6.30	-7.56	29.50	5.67	9.95	14.33	-0.08	21.03
	YF in win.	T_{total}	4931.89	42.90	4865.22	4998.49	4882.68	4920.57	4981.96	4872.24	4992.54
	YF in win.	ΔT	168.41	42.90	101.75	235.01	119.21	157.09	218.48	108.77	229.06
	YF out of win.	T_{total}	4937.39	47.90	4795.60	5020.45	4893.60	4932.73	4970.92	4873.69	5005.72
	YF out of win.	ΔT	173.91	47.90	32.12	256.98	130.13	169.25	207.45	110.22	242.24

Table B.1: Statistical summary of variables across different window widths for the Suzuka circuit. Q2 also represents the median.

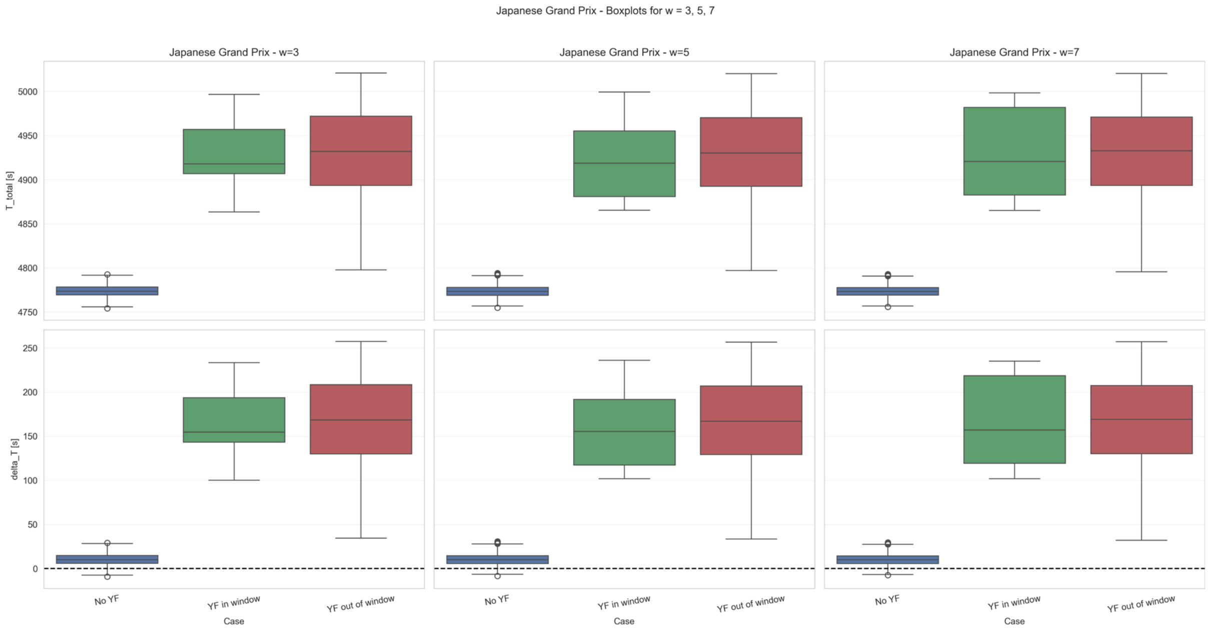

Figure B.1: Statistical distribution of race outcomes for the Suzuka circuit, under all window widths. The top panels show total race time (T_{total}), and the bottom panels show the deviation from the optimal strategy (ΔT) across the three yellow flag scenarios.

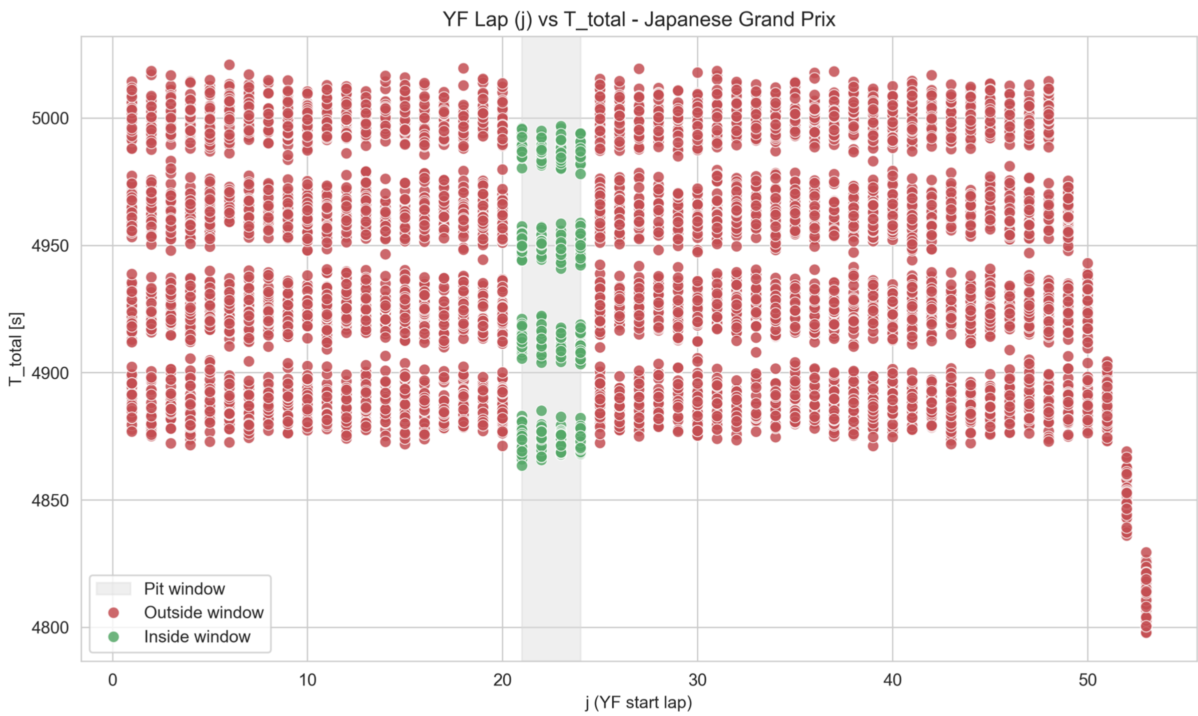

Figure B.2: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Suzuka circuit when $w = 3$.

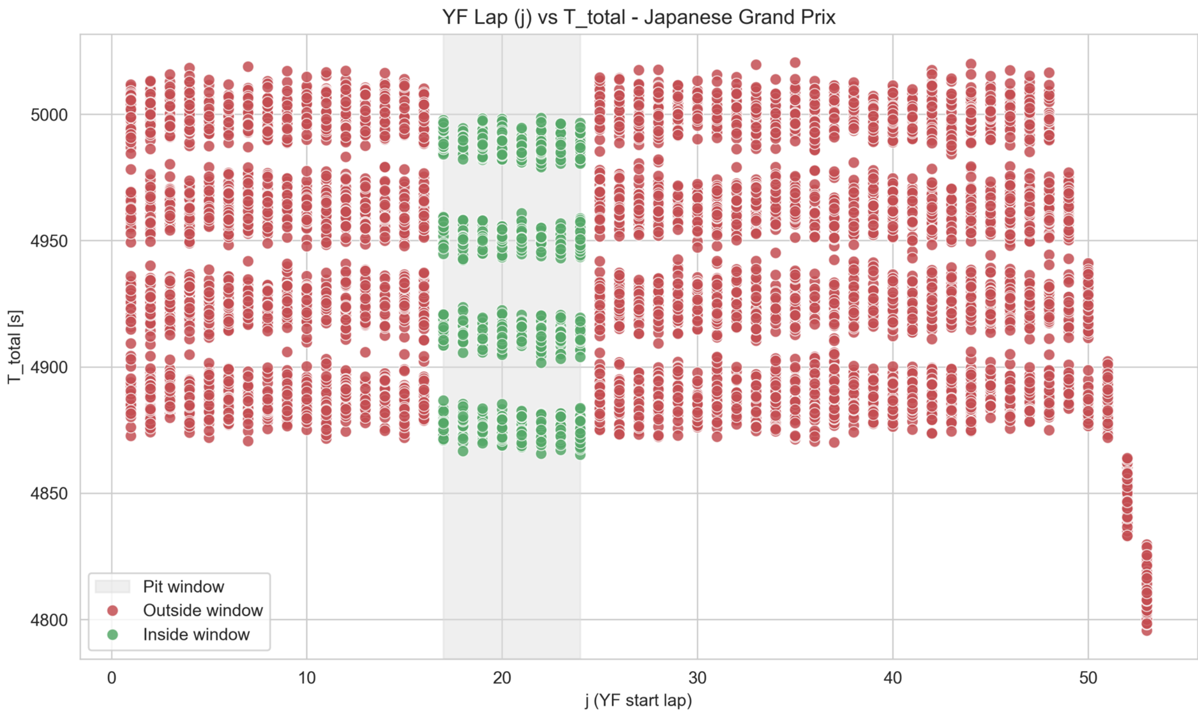

Figure B.3: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Suzuka circuit when $w = 7$.

B.2 Two-Stop Strategy

The Table B.2 and Figure B.4 provide the statistical summary and the distribution of the total race time (T_{total}) and its deviation (ΔT) for the one-stop strategy (Japanese Grand Prix) simulation across all window widths. Additionally, the scatter plots in Figures B.5 and B.6 illustrate the sensitivity of the total race time to the yellow flag start lap (j) for the extreme window widths of $w = 3$ and $w = 7$.

w	Group	Variable	Mean	Std	Min	Max	Q1	Q2	Q3	P5	P95
w=3	Overall	T_{total}	5260.84	64.10	5132.73	5376.05	5225.77	5272.25	5311.75	5151.78	5347.09
	Overall	ΔT	128.55	64.10	0.44	243.76	93.48	139.96	179.46	19.49	214.80
	No YF	T_{total}	5161.58	11.95	5132.73	5199.54	5152.66	5160.55	5169.38	5143.27	5182.70
	No YF	ΔT	29.29	11.95	0.44	67.26	20.37	28.26	37.09	10.98	50.41
	YF in win.	T_{total}	5281.92	33.91	5218.44	5348.36	5252.33	5283.04	5306.26	5231.35	5335.54
	YF in win.	ΔT	149.63	33.91	86.15	216.07	120.04	150.75	173.97	99.06	203.25
	YF out of win.	T_{total}	5291.66	38.54	5165.75	5376.05	5261.31	5291.82	5324.45	5235.32	5351.42
	YF out of win.	ΔT	159.37	38.54	33.46	243.76	129.02	159.53	192.16	103.03	219.13
w=5	Overall	T_{total}	5260.26	63.61	5129.74	5373.96	5226.48	5271.85	5310.87	5150.90	5344.69
	Overall	ΔT	127.97	63.61	-2.55	241.67	94.19	139.56	178.58	18.61	212.40
	No YF	T_{total}	5161.44	12.40	5129.74	5200.18	5152.16	5160.09	5170.42	5143.00	5183.21
	No YF	ΔT	29.15	12.40	-2.55	67.89	19.87	27.80	38.13	10.71	50.92
	YF in win.	T_{total}	5284.33	34.85	5221.48	5353.19	5252.77	5286.27	5315.89	5231.25	5335.51
	YF in win.	ΔT	152.04	34.85	89.19	220.90	120.48	153.98	183.60	98.96	203.22
	YF out of win.	T_{total}	5290.78	38.33	5165.66	5373.96	5261.29	5290.75	5322.71	5234.10	5350.31
	YF out of win.	ΔT	158.49	38.33	33.37	241.67	129.00	158.46	190.42	101.81	218.02
w=7	Overall	T_{total}	5260.55	63.80	5126.94	5375.70	5226.50	5272.30	5311.28	5151.08	5344.11
	Overall	ΔT	128.26	63.80	-5.35	243.41	94.21	140.01	178.99	18.79	211.82
	No YF	T_{total}	5161.19	11.93	5126.94	5200.70	5152.56	5160.56	5169.19	5142.45	5181.70
	No YF	ΔT	28.90	11.93	-5.35	68.41	20.27	28.27	36.90	10.16	49.41
	YF in win.	T_{total}	5285.17	34.38	5220.49	5353.42	5255.20	5288.01	5316.55	5233.31	5335.80
	YF in win.	ΔT	152.88	34.38	88.20	221.13	122.91	155.72	184.26	101.02	203.51
	YF out of win.	T_{total}	5291.57	38.48	5166.56	5375.70	5262.19	5293.00	5323.79	5234.09	5349.61
	YF out of win.	ΔT	159.28	38.48	34.27	243.41	129.90	160.71	191.50	101.80	217.32

Table B.2: Statistical summary of variables across different window widths for the Zandvoort circuit. Q2 also represents the median.

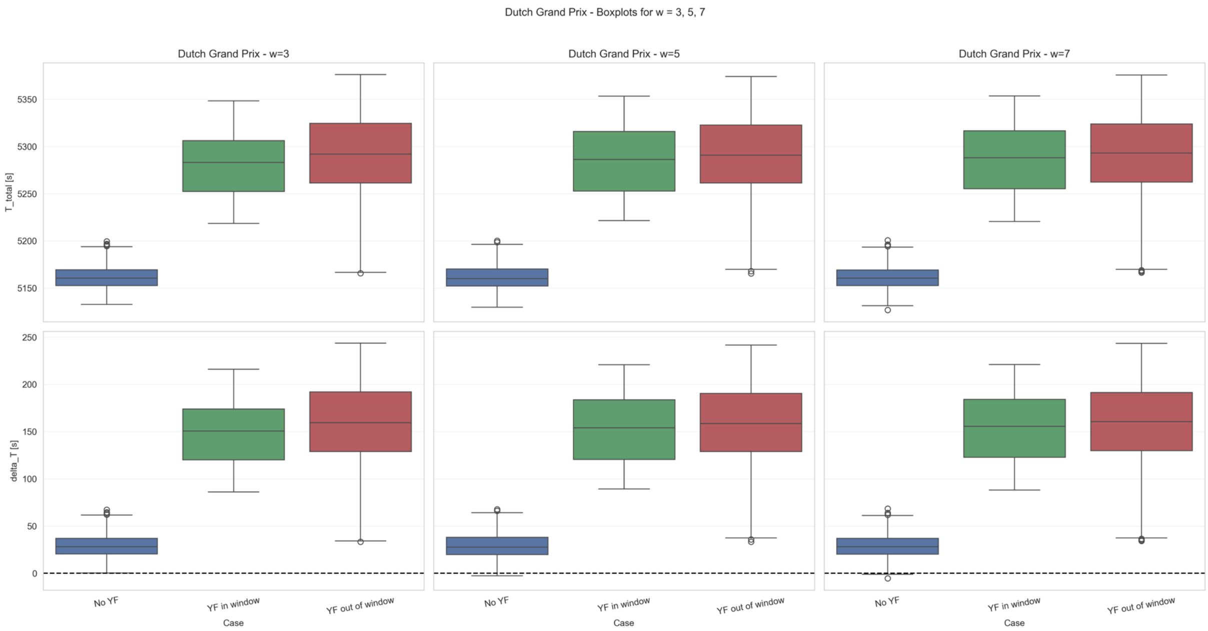

Figure B.4: Statistical distribution of race outcomes for the Zandvoort circuit, under all window widths. The top panels show total race time (T_{total}), and the bottom panels show the deviation from the optimal strategy (ΔT) across the three yellow flag scenarios.

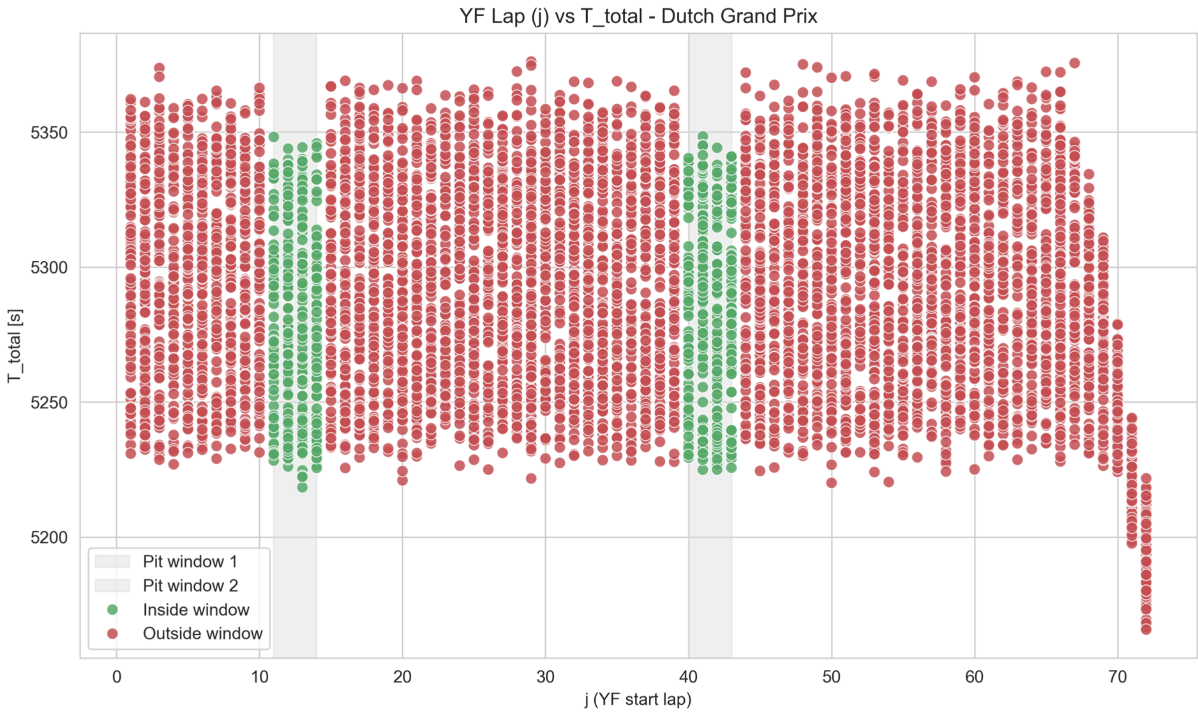

Figure B.5: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Zandvoort circuit when $w = 3$.

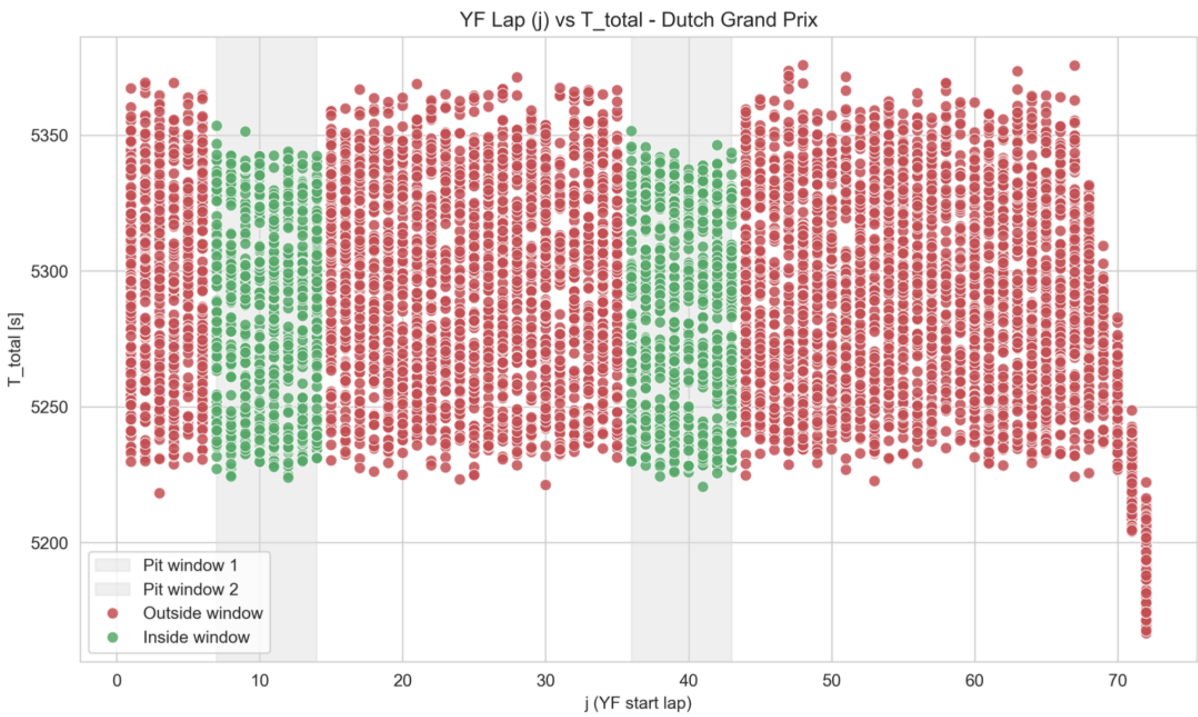

Figure B.6: Scatter plot of total race time (T_{total} in seconds) compared to the yellow flag start lap (j), for the Zandvoort circuit when $w = 7$.