Faster than Hamilton! - Optimising F1 Strategies

MATH5003M - Assignment in Mathematics (30cr)
Supervisor: Cédric Beaume
University of Leeds

James Murkin: 201120564

May 7, 2021

UNIVERSITY OF LEEDS

School of Mathematics

Declaration of Academic Integrity
for Individual Pieces of Work

I am aware that the University defines plagiarism as presenting someone else's work as your own. Work means any intellectual output, and typically includes text, data, images, sound or performance.

I promise that in the attached submission I have not presented anyone else's work as my own and I have not colluded with others in the preparation of this work. Where I have taken advantage of the work of others, I have given full acknowledgement. I have read and understood the University's published rules on plagiarism and also any more detailed rules specified at School or module level. I know that if I commit plagiarism I can be expelled from the University and that it is my responsibility to be aware of the University's regulations on plagiarism and their importance.

I re-confirm my consent to the University copying and distributing any or all of my work in any form and using third parties (who may be based outside the EU/EEA) to monitor breaches of regulations, to verify whether my work contains plagiarised material, and for quality assurance purposes.

I confirm that I have declared all mitigating circumstances that may be relevant to the assessment of this piece of work and that I wish to have taken into account. I am aware of the School's policy on mitigation and procedures for the submission of statements and evidence of mitigation. I am aware of the penalties imposed for the late submission of coursework.

Student Signature JAMurkin Date 07/05/21

Student Name JAMES MURKIN Student Number 201120564

Please note.
When you become a registered student of the University at first and any subsequent registration you sign the following authorisation and declaration:
"I confirm that the information I have given on this form is correct. I agree to observe the provisions of the University's Charter, Statutes, Ordinances, Regulations and Codes of Practice for the time being in force. I know that it is my responsibility to be aware of their contents and that I can read them on the University web site. I acknowledge my obligation under the Payment of Fees Section in the Handbook to pay all charges to the University on demand.
I agree to the University processing my personal data (including sensitive data) in accordance with its Code of Practice on Data Protection http://www.leeds.ac.uk/dpa . I consent to the University making available to third parties (who may be based outside the European Economic Area) any of my work in any form for standards and monitoring purposes including verifying the absence of plagiarised material. I agree that third parties may retain copies of my work for these purposes on the understanding that the third party will not disclose my identity."

Contents

Chapter 1

Introduction

Formula 1 is a motorsport which consists of 20 drivers from 10 teams attempting to complete a set number of laps around a track quicker than all other drivers. Points are awarded to drivers based on their finishing position. After each race, the points scored by each driver is added to their total for the season, thus creating a Drivers' Championship. There is an equivalent Constructors' Championship, where each team is given the combined points scored by their drivers. After the season has been completed, which typically consists of 20 separate races, the driver and team with the most points wins their respective championship. The goal of each team is therefore a simple one: give themselves the best chance of securing the Constructors' and/or Drivers' Championship. A team will have the best chance of winning at least one of the championships if they can provide their drivers with both a good car and good race strategies.

The aim of this project is to study how we can model a car's lap time around a race track. This will involve modelling how tyres wear out over time and how much the fuel load affects the lap time. With the model we create, we will work towards finding the fastest possible way of completing a race of N laps. While it is possible for a car to complete a whole race without stopping, cars are able to make pit stops, where fuel is added and the tyres are changed at the cost of losing time in the pit lane. We will see what conditions are required so that it is faster to make a given number of pit stops in the races, rather than driving non-stop from start to finish.

Using our lap time model, we will "become" a Strategy Engineer for a Formula 1 team at a given race, where we will find the best strategy for this race according to the data we acquire.[1] This involves analysing lap time data from the race to estimate all the necessary values. Once we find the fastest strategy, we will compare it to the race winning strategy and also fit the race winning strategy using our model. We will discuss ways in which we can adapt or extend the simple model we create, as well as other factors that can influence the strategy decisions that are made in a real life situation.

The rest of this project will cover the F1 Strategy Competition run by Cédric Beaume, which involves creating a virtual team by allocating a fixed budget into 4 separate components, hiring drivers through a drafting process, and then putting our race strategies to the test against 12 other competitors. Using Python, we will use Monte Carlo simulations to give ourselves what we think is the best possible chance of winning the team championship. After the season is complete, we will discuss how we could have altered our initial budget allocation in order to give ourself the best chance of winning against the 12 other teams.

In Chapter 2, we will develop our lap time model and use it to formulate race strategies. Chapter 3 will be focused around a real life race, for which we will create an optimal race strategy using data from the race, followed by a discussion on areas our model can be improved. Chapter 4 will cover the preparation for the F1 Strategy Competition, discussing the budget allocation process and how we can give ourselves what we believe is the best chance of winning the championship. Chapter 5 will review the events of the competition, discussing what we did well, what did not go well, and what we could have done differently.

Chapter 2

Lap time modelling

In this chapter, we will start by introducing how we can model the time it takes for a driver to complete a lap, depending on a few simple factors. We will then use this model to determine what strategy can be used in order to complete a race in the quickest time possible.

2.1 The simple lap time model

When modelling the time taken for a driver to complete a lap of a track, there are two important variables: the amount of fuel onboard the car and the level of tyre wear. A car with more fuel will be a heavier car, and will therefore experience greater inertia when attempting to accelerate and decelerate which has the effect of slowing a car down. As tyres are used for more and more laps, they begin to wear out. A worn set of tyres will provide less grip than a new set, thus slowing the car down.

We can therefore model a lap time [2] by writing

$$t = t_b + p_t + p_f, \tag{2.1}$$

where the separate components, all measured in seconds, are defined:

t: The total time taken to complete a lap;
t_b: The base lap time around the given track, assuming minimal fuel and no tyre wear;
p_t: Time penalty due to tyre wear;
p_f: Time penalty due to fuel level.

During a Formula 1 race weekend, which consists of 3 practice sessions, a qualifying session and a race, strategy engineers for each team will be tasked with estimating these components. The base lap time will be a constant, whereas the fuel component will be a decreasing function of laps completed and the tyre wear component will be an increasing function of laps completed.

Example 2.1

Suppose that at a given race weekend, a strategy engineer estimates each of the above components with:

$$\begin{aligned} t_b &= 90; \\ p_t(l;T) &= Tl; \\ p_f(l;F,l_{\max}) &= F(l_{\max}-l). \end{aligned}$$

$$\implies t = 90 + Tl + F(l_{\max}-l), \qquad l = 0, 1, \ldots, l_{\max}-1.$$

Here, $T > 0$ is the additional time penalty per lap due to tyre wear, $F > 0$ is the additional time penalty per lap of fuel onboard the car, l is the number of laps completed in the current stint and $l_{\max}$ is the total number of laps in the stint. When $l = l_{\max}$, every lap in the stint has been completed, so there is no fuel in the car. At this point, either the race will have been completed, or we would make a pit stop to change tyres and acquire more fuel.

If we let $T = 0.05, F = 0.1, l_{\max} = 20$, then on the 11^{th} lap (i.e. $l = 10$ since only 10 laps have been completed), the lap time according to this model would be

$$\begin{aligned} t &= 90 + 0.05 \cdot 10 + 0.01 \cdot (20-10) \\ \implies t &= 91.5 \text{ seconds} \end{aligned}$$

2.2 Extending to a stint

The obvious extension to having a model for calculating a driver's lap time around a circuit is to consider the set of lap times across a stint, which is simply a set of successive laps. The notion of a stint has already been encoded by the use of the variable l in our lap time model, as the tyre wear and fuel load penalties are a function of the number of laps completed in the stint. We will see how the time to complete a stint of laps is of interest to a strategy engineer when multiple tyre compounds are available to choose from, although for now we will just consider methods for calculating the time taken to complete a stint.

To define a way of calculating the time taken to complete a stint, we will be using the following lap time model:

$$t_{l+1} = t_b + Tl + F(l_{\max}-l), \qquad l = 0, 1, \ldots, l_{\max}-1, \tag{2.2}$$

where $T, F > 0$ are the tyre wear and fuel level time penalties respectively, and t_b, l, and $l_{\max}$ are the same as defined in (2.1). Here, t_{l+1} is simply just the lap time on lap $l+1$, at which point only l laps have been completed in the stint. We assume a linear law for the tyre wear penalty for simplicity, although other laws, such as quadratic or exponential, may be appropriate. In reality, there are many factors that influence the wear rate of a tyre, which we will discuss later, but a linear law can occur in some circumstances. A linear law for the fuel load penalty is appropriate here because the weight of the car is significantly larger than the weight of the fuel that is burned across a single lap. As a result, the percentage change in weight of the car after a lap of fuel has been used will be close to a constant value, therefore the change in lap time will be effectively linear. If there were a motorsport with incredibly light vehicles, the percentage change in weight of the car after a lap of fuel is used would increase more noticeably, making a linear law for the fuel load penalty less appropriate.

Proposition 2.2

Suppose we have a stint of length $l_{\max}$, and let t_b be the base lap time, T the tyre wear time penalty and F the fuel load penalty. Then, the total time to complete the stint is

$$t_S = \sum_{l=0}^{l_{\max}-1} t_l = l_{\max} t_b + \frac{l_{\max}}{2} \left[T(l_{\max} - 1) + F(l_{\max} + 1) \right]. \tag{2.3}$$

Proof. We have

$$\begin{aligned}
t_s = \sum_{l=0}^{l_{\max}-1} t_l &= \sum_{l=0}^{l_{\max}-1} t_b + Tl + F(l_{\max} - l) \\
&= l_{\max} t_b + T(0 + 1 + \cdots + (l_{\max} - 1)) + F(l_{\max} + (l_{\max} - 1) + \cdots + 2 + 1) \\
&= l_{\max} t_b + T \frac{(l_{\max} - 1) l_{\max}}{2} + F \frac{l_{\max}(l_{\max} + 1)}{2} \\
\implies \sum_{l=0}^{l_{\max}-1} t_l &= l_{\max} t_b + \frac{l_{\max}}{2} \left[T(l_{\max} - 1) + F(l_{\max} + 1) \right].
\end{aligned}$$

□

Example 2.3

Let $T = 0.05, F = 0.1, t_b = 90$, and suppose we wish to complete a stint of $l_{\max} = 20$ laps. Then, the total time to complete this stint is

$$\begin{aligned}
\sum_{l=0}^{l_{\max}-1} t_l &= 20 \cdot 90 + \frac{20}{2} [0.05(20 - 1) + 0.1(20 + 1)] \\
&= 1830.5\text{s}.
\end{aligned}$$

2.3 Extending to a full race

According to Article 24.4k of the 2021 Formula 1 Sporting Regulations, in a dry race drivers a required to use at least 2 of the 3 available dry-weather tyre compounds else they would be disqualified from the race results.[3] As a result, at least one pit stop is required at some point during the race in order to change the tyres. We will later consider ways of modelling the time required to make a pit stop, as they are highly susceptible to any mistakes made by the pit crew or driver. For now, we will assume that there are no errors made during a pit stop, and that the only effect on the time taken is the amount of fuel that is being added to the car. When devising a strategy for a given race, a Strategy Engineer would need to consider the time taken to make a pit stop. For instance, making more pit stops throughout the race allows a driver to have lower fuel and freshers tyres on average, but they would lose a significant amount of time in the pit lane relative to someone making fewer stops throughout the race.

We will see how we can find the fastest possible race strategy and also how we can compare different strategies. For simplicity, we will also assume that there is only one tyre compound available to the drivers and so, we will also lift the requirement for a pit stop to be made.

Definition 2.4

Let $l_1, l_2, \ldots, l_r \in \mathbb{N}$ be the laps on which a pit stop is made, with $l_0 = 0 < l_1 < l_2 < \cdots < l_r$, and suppose the race has l_{r+1} laps, with $l_r < l_{r+1}$. Then, the intervals $[l_0+1, l_1], [l_1+1, l_2], \ldots, [l_r+1, l_{r+1}]$ are the $r+1$ *stints* that make this strategy. The $r+1$ stints define an *r-stop strategy.*

Example 2.5

Suppose we have a race of 40 laps. An example of a 3-stop strategy is one that stops on laps $l_1 = 10, l_2 = 20$, and $l_3 = 30$, with $l_4 = 40$ being the length of the race. The 4 stints in this strategy consist of laps 1-10, 11-20, 21-30, 31-40. An alternative 2-stop strategy could have $l_1 = 15, l_2 = 30$, and $l_3 = 40$, where the 3 stints are made of laps 1-15, 16-30, 31-40.

With a formalised definition of a race strategy for any number of pit stops, we can now work towards a method for calculating the total race time for a given strategy.

Proposition 2.6

Let $T, F > 0$ be the respective tyre wear and fuel load penalties at a given track, and suppose we wish to wish to use an r-stop strategy at a race of l_{r+1} laps where the base lap-time is given by t_b. Then, the total time required to complete the race, according to the simple model, is

$$\begin{aligned} t_R(l_1, \ldots, l_{r+1}; F, T, t_s, t_b) = l_{r+1} t_b + F\frac{l_1(l_1+1)}{2} + T\frac{l_1(l_1-1)}{2} \\ + \sum_{i=2}^{r+1} \Bigg[t_s + \frac{1}{2}(l_i - l_{i-1}) + F\frac{(l_i - l_{i-1})(l_i - l_{i-1} + 1)}{2} \\ + T\frac{(l_i - l_{i-1})(l_i - l_{i-1} - 1)}{2} \Bigg]. \end{aligned} \tag{2.4}$$

Here, we have assumed that the time required to make a pit stop is given by

$$t_p = t_s + \frac{1}{2}(l_i - l_{i-1}), \tag{2.5}$$

where t_s is some constant, in seconds, and we add half a second to the stop time for every lap of fuel that is added to the car. We interpret t_s as the time required to drive through the pit lane without making a pit stop.

We first notice that the base lap time t_b is counted on every lap, and so we can simply add it to the total sum l_{r+1} times. The other terms outside the summation are simply the total tyre wear and fuel load penalties for a stint of length l_1. The summation is used to calculate the individual stint times for the remaining r stints in the race after the first stint has been completed where the tyre wear and fuel load penalties are equivalent to those found in (2.3). We have also included the time required to make the necessary pit stop in between each stint, as given by (2.5).

Example 2.7

Let $T = 0.05, F = 0.1, t_b = 90, t_s = 15$, and suppose we wish to complete a 60 lap race with the following 2-stop strategy: $l_1 = 20, l_2 = 40$, hence $l_3 = 60$. Then, the total race time for this strategy is

$$
\begin{aligned}
t_R &= 60 \cdot 90 + 0.1\frac{20 \cdot 21}{2} + 0.05\frac{20 \cdot 19}{2} \\
&\qquad + \left(15 + \frac{1}{2}(40-20) + 0.1\frac{(40-20)(40-20+1)}{2} + 0.05\frac{(40-20)(40-20-1)}{2}\right) \\
&\qquad + \left(15 + \frac{1}{2}(60-40) + 0.1\frac{(60-40)(60-40+1)}{2} + 0.05\frac{(60-40)(60-40-1)}{2}\right) \\
\implies t_R &= 5400 + 21 + \frac{19}{2} + \left(15 + 10 + 21 + \frac{19}{2}\right) + \left(15 + 10 + 21 + \frac{19}{2}\right) \\
&= 5541.5\text{s} \\
&= 1\text{hr } 32\text{m } 20.5\text{s}.
\end{aligned}
$$

Therefore, this driver's race would take just over 1 hour and 32 minutes according to this 2-stop strategy.

We can compare this to a 3-stop strategy. Let $l_1 = 15, l_2 = 30, l_3 = 45$, and hence $l_4 = 60$. This strategy would give a total race time of 5536.5 seconds, which is a 5 second improvement on the 2-stop strategy!

Being able to calculate the time it would take for a driver to complete a race given a race strategy is one thing, but a Strategy Engineer would want to know how to find the best possible strategy. As such, we would want to minimise the value of t_R with respect to the laps on which pit stops are made.

Theorem 2.8

Let $l_1, \dots, l_r$ define an r-stop strategy in a race of l_{r+1} laps, with $l_r < l_{r+1}$. Let F be the time penalty per lap of fuel onboard the car, and T be the time penalty per lap completed on the tyres. Let t_s be the time required for the car to drive through the pit lane and suppose it takes half a second to add a lap of fuel in a pit stop.

Then, starting with the fixed race length l_{r+1}, the optimal laps on which to make a pit stop can be found recursively with the following formula:

$$
l_i = \frac{i}{i+1} l_{i+1} + \frac{1}{2(i+1)(T+F)}, \qquad \forall i = 1, \dots, r \tag{2.6}
$$

Proof. Since we aim to minimise the fuel penalty across the race, we want to find the l_i that minimises (1.7). Starting with l_1, we take the partial derivative of t_R with respect to l_1, and set to 0.

$$
\begin{aligned}
\frac{\partial t_R}{\partial l_1} &= \frac{F}{2}l_1 + \frac{F}{2}(l_1+1) + \frac{T}{2}l_1 + \frac{T}{2}(l_1-1) - \frac{1}{2} - \frac{F}{2}(l_2-l_1+1) - \frac{F}{2}(l_2-l_1) \\
&\quad - \frac{T}{2}(l_2-l_1-1) - \frac{T}{2}(l_2-l_1) = 0 \\
\implies 0 &= 4l_1(T+F) - 2l_2(T+F) - 1 \\
\implies l_1 &= \frac{l_2}{2} + \frac{1}{4(T+F)}.
\end{aligned}
$$

For the remaining l_i, we expand the sum so as to only consider terms with l_i. Again, we take the partial derivative of t_R with respect to l_i, and set to 0.

$$
\begin{aligned}
\frac{\partial t_R}{\partial l_i} = \frac{\partial}{\partial l_i}\Bigg[&\left(t_s + \frac{1}{2}(l_i - l_{i-1}) + F\frac{l_i - l_{i-1}}{2}(l_i - l_{i-1} + 1) + T\frac{l_i - l_{i-1}}{2}(l_i - l_{i-1} - 1)\right) \\
&+ \left(t_s + \frac{1}{2}(l_{i+1} - l_i) + F\frac{l_{i+1} - l_i}{2}(l_{i+1} - l_i + 1) + T\frac{l_{i+1} - l_i}{2}(l_{i+1} - l_i - 1)\right)\Bigg]
\end{aligned}
$$

$$
\begin{aligned}
\implies \frac{\partial t_R}{\partial l_i} &= \frac{T}{2}(2l_i - 2l_{i-1} - 1 - 2l_{i+1} + 2l_i + 1) + \frac{F}{2}(2l_i - 2l_{i-1} + 1 - 2l_{i+1} + 2l_i - 1) = 0 \\
\implies 0 &= (T+F)(4l_i - 2l_{i-1} - 2l_{i+1}) \\
\implies 2l_i &= l_{i-1} + l_{i+1}. \qquad (\star)
\end{aligned}
$$

Using this relationship, we can find all the l_i. We use the fact that $l_1 = \frac{l_2}{2} + \frac{1}{4(T+F)}$ and that the race length, l_{r+1}, is fixed.

Substituting l_1 into $(\star)$, and setting $i = 2$, we find that

$$
\begin{aligned}
2l_2 &= \frac{l_2}{2} + \frac{1}{4(T+F)} + l_3 \\
\implies l_2 &= \frac{2}{3}l_3 + \frac{1}{6(T+F)}.
\end{aligned}
$$

Now, for $i = 3$, we find that

$$
\begin{aligned}
2l_3 &= \frac{2}{3}l_3 + \frac{1}{6(T+F)} + l_4 \\
\implies l_3 &= \frac{3}{4}l_4 + \frac{1}{8(T+F)}.
\end{aligned}
$$

As we increase i up to r, we find that this pattern holds, thus giving the recurrence formula

$$
l_i = \frac{i}{i+1}l_{i+1} + \frac{1}{2(i+1)(T+F)}. \qquad \square
$$

It should be noted that the optimal stopping laps are not affected by the time taken to complete the pit stop, as we are effectively minimising the total fuel load and tyre penalty *on track*. When comparing strategies with the same number of stops, t_s is a constant penalty, regardless of when the pit stops are made.

Example 2.9

Suppose we want to find the fastest 2-stop strategy at a race of 25 laps. Let $F = 0.1, T = 0.1, t_s = 15, t_b = 90$. By setting $l_3 = 25$, we find that

$$\begin{aligned}
l_2 &= \frac{2}{3} \cdot 25 + \frac{1}{2(3)(0.1+0.1)} \\
&= 17.5 \\
\implies l_1 &= \frac{1}{2} \cdot 17.5 + \frac{1}{2(2)(0.1+0.1)} \\
&= 10.
\end{aligned}$$

Therefore, the optimal laps on which to make the pit stops for this 2-stop strategy are laps $l_1 = 10$ and $l_2 = 17.5$. Of course, we cannot stop halfway through the lap, so we will use (2.8) to determine whether pitting on lap 17 or 18 is the better strategy.

$$\begin{aligned}
l_2 = 17: \qquad t_R &= 2250 + \frac{11}{2} + \frac{9}{2} + \left(15 + \frac{1}{2}(17-10) + 0.1\frac{7 \cdot 8}{2} + 0.1\frac{7 \cdot 6}{2}\right) \\
&\quad + \left(15 + \frac{1}{2}(25-17) + 0.1\frac{8 \cdot 9}{2} + 0.1\frac{8 \cdot 7}{2}\right) \\
&= 2308.8; \\
l_2 = 18: \qquad t_R &= 2250 + \frac{11}{2} + \frac{9}{2} + \left(15 + \frac{1}{2}(18-10) + 0.1\frac{8 \cdot 9}{2} + 0.1\frac{8 \cdot 7}{2}\right) \\
&\quad + \left(15 + \frac{1}{2}(25-18) + 0.1\frac{7 \cdot 8}{2} + 0.1\frac{7 \cdot 6}{2}\right) \\
&= 2308.8.
\end{aligned}$$

Therefore, it does not matter whether the second pit stop is made on lap 17 or 18 as they both have the same overall race time! In fact, in cases where the recursive formula yields an exact half-lap as the optimal stopping lap, it does not matter whether we round up or down when selecting the strategy.

Example 2.10

In the same race of 25 laps, suppose we want to find the fastest 1-stop strategy for the given tyre wear and fuel level penalties. With l_2=25, we find that

$$\begin{aligned}
l_1 &= \frac{1}{2} \cdot 25 + \frac{1}{2(2)(0.1+0.1)} \\
&= 13.75.
\end{aligned}$$

Therefore, stopping on lap 14 would provide the optimal 1-stop strategy for this race. To show this, we will compare the total race time when stopping on lap 14 with the race time when stopping on lap 13.

$$l_1 = 13: \qquad t_R = 2250 + 9.1 + 7.8 + \left(15 + \frac{1}{2}(25-13) + 0.1\frac{12 \cdot 13}{2} + 0.1\frac{12 \cdot 11}{2}\right)$$
$$= 2302.3;$$
$$l_1 = 14: \qquad t_R = 2250 + 10.5 + 9.1 + \left(15 + \frac{1}{2}(25-14) + 0.1\frac{11 \cdot 12}{2} + 0.1\frac{11 \cdot 10}{2}\right)$$
$$= 2302.2.$$

Therefore, stopping on lap 14 results in a race time that is 0.1s quicker than stopping on lap 13. The optimal 1-stop strategy is also 6.6s faster than the optimal 2-stop strategy! When deciding which strategy to use in a race of a given length, the best approach is to find the optimal r-stop strategy for each r (up to some maximum value) and then compare the total race times for each of these optimal strategies.

Note

Our method for finding the optimal laps to stop essentially splits the race into two separate components: the first stint and the remaining r stints. When we choose $r \geq 2$, our method finds the optimal lap on which to make the first pit stop, and then divides the remaining laps into r stints of equal length. We can see this by considering a simple rearrangement of ($\star$) from the proof of Theorem 2.8:

$$2l_i = l_{i-1} + l_{i+1} \implies l_i = \frac{l_{i-1} + l_{i+1}}{2}.$$

Clearly, the optimal stop laps are the halfway point between the preceding and successive stops (or race length, if $i = r$).

For instance, suppose we have a race of 60 laps with $T = 0.05, F = 0.1$ and we wish to complete a 3-stop strategy. According to (2.6), we would make our pit stops on laps 18, 32, and 46. Hence, the first stint consists of 18 laps, while the 3 remaining stints are 14 laps each.

2.4 Modelling a pit stop

In the previous section, we assumed a simple model for the time taken to complete a pit stop, which depended on only 2 variables:

$$t_p = t_s + \frac{1}{2}(l_i - l_{i-1}), \tag{2.7}$$

where t_s is the base pit lane time, and every lap of fuel added to the car takes an additional 0.5 seconds. As mentioned previously in Proposition 2.6, we interpret t_s as the time required to drive through the pit lane without making a pit stop, which depends only on the track.

In this section, we will consider three possible modifications in order to offer a more "realistic" model for the time taken to complete a pit stop.

Relationship between fuel flow and track length

Article 5.3 of the 2021 Formula 1 Sporting Regulations state that the race length (in terms of laps) is the minimum number of laps to ensure at least 305km has been driven.[3] For instance, the Silverstone Circuit in Great Britain is 5.891km[4] in length. Each race at Silverstone consists of 52 laps, resulting in the total race distance being 306.198km, which is just over the required

305km. Note that this is slightly less than the distance covered by 52 laps of the circuit (306.332km) as the start line is around 130m away from the timing line where the race ends. By comparison, the Red Bull Ring in Spielberg, Austria, is only 4.318km[5] in length, resulting in a race of 71 laps. The rules also state that cars are required to use no more than 110kg of fuel during a race. Therefore, assuming teams fill up the cars entirely at the start of a race, and then finish the race with no fuel, each lap of the Silverstone Circuit uses fuel at a rate of 2.115kg/lap, whereas at the Red Bull Ring fuel is used at a rate of 1.549kg/lap. In general, longer laps have a higher rate of fuel usage per lap.

When devising strategies for different tracks, one thing that a Strategy Engineer can assume to be constant is that rate at which fuel is added to the car (in kg/s) during a pit stop. By assuming this to be constant, the time taken to add a single lap of fuel is dependent on the length of the track or, equivalently, the number of laps in the race. In the 2021 Formula 1 Season, there are scheduled to be 1394 racing laps at 23 circuits, giving an average of ≈ 60.61 laps/race. For simplicity, we will assume that the average number of laps in a race is 60 laps, and that the amount of time required to add a single lap of fuel during a pit stop at a race of 60 laps is 0.5 seconds, as in the original pit stop model. Using this, we can make a simple modification to the pit stop model:

$$t_p = t_s + \frac{1}{2} \cdot \frac{60}{l_{r+1}} (l_i - l_{i-1}), \tag{2.8}$$

where l_{r+1} is the number of laps in the race. With this model, tracks with fewer than 60 laps will see the time required to add a single lap of fuel increased above 0.5s, and vice versa for races with more laps.

Influence of changing tyres

Historically, when refuelling was permitted during Formula 1 races, pit crews were able to change the tyres at the same time as refuelling the car. This is also the case in many forms of motorsport, for which our simple model can be adapted if necessary. Because F1 tyres are only connected to the car with one lug nut, according to Article 12.8.2 of the 2021 Technical Regulations [6], tyre changes are completed long before the fuel is added in its entirety. In other motorsports, however, this is not necessarily the case, so there may be occasions where the time to change the tyres is greater than the time to add the fuel to the car. As such, our model can be simply adapted to the following:

$$t_p = t_s + \min\left(t_t, \frac{1}{2}(l_i - l_{i-1})\right), \tag{2.9}$$

where t_t is the time taken to change the tyres, which for now we assume to be some constant. We have also assumed that fuel is added to the car at a rate of 0.5s/lap, in contrast to the modification in (2.8).

The case when refuelling is banned

According to Article 30.1c in the 2021 Formula 1 Sporting Regulations, refuelling is not permitted during the race[3], so at pit stops only tyres need to be changed. As such, teams aim to change the tyres on the car as quick as possible. For instance, the fastest pit stop in Formula 1 history was completed at the Brazilian Grand Prix in 2019, where Red Bull Racing changed the tyres on Max Verstappen's car in 1.82 seconds.[7] In order to model this, we need to consider the time taken to drive through the pit lane t_s and the time taken to change the tyres. The latter can be described by a random variable. If we assume the time taken to change the tyres in a pit stop is normally distributed, then we could write our pit stop model as

$$t_p = t_s + X_T, \tag{2.10}$$

where t_s is the same as in (2.7), and $X_T \sim \mathcal{N}\left(\mu, \sigma^2\right)$ is a random variable that describes the time taken to change the tyres in a pit stop, which we have assumed to be normally distributed with mean μ and variance σ^2. A Strategy Engineer would need to estimate the values of μ and σ^2, although a simple estimate could be $\hat{\mu} = 2.5$ and $\hat{\sigma}^2 = 0.25^2$. With these estimates, we would expect 95% of all pit stops to have the tyre change completed within 2 to 3 seconds.

2.5 Example: 10-lap race

Suppose we have a race of $l_{r+1} = 10$ laps around some circuit. A team's Strategy Engineer estimates that the tyre wear penalty is $T = 0.15$ seconds per lap completed on the tyres, and that each lap of fuel onboard the car adds an additional $F = 0.2$ seconds to the lap time. Suppose the base lap time is $t_b = 90$ seconds, and that it takes $t_s = 5$ seconds to drive through the pit lane. We will assume the simple pit stop model (2.7). We wish to find the fastest strategy for this race.

In a race of 10 laps, we can make up to 9 pit stops (we cannot pit on lap 10), so we will find the optimal laps to pit for strategies with up to 9 stops. Using (2.6), we find that the following strategies are optimal for the given number of stops:

$$\begin{aligned}
\text{1 stop:} &\quad l_1 = 6\\
\text{2 stops:} &\quad l_1 = 4, l_2 = 7\\
\text{3 stops:} &\quad l_1 = 4, l_2 = 6, l_3 = 8\\
\text{4 stops:} &\quad l_1 = 3, l_2 = 5, l_3 = 7, l_4 = 8\\
\text{5 stops:} &\quad l_1 = 3, l_2 = 4, l_3 = 6, l_4 = 7, l_5 = 9\\
\text{6 stops:} &\quad l_1 = 3, l_2 = 4, l_3 = 5, l_4 = 6, l_5 = 8, l_6 = 9\\
\text{7 stops:} &\quad l_1 = 2, l_2 = 4, l_3 = 5, l_4 = 6, l_5 = 7, l_6 = 8, l_7 = 9\\
\text{8 stops:} &\quad l_1 = 2, l_2 = 3, l_3 = 4, l_4 = 5, l_5 = 6, l_6 = 7, l_7 = 8, l_8 = 9\\
\text{9 stops:} &\quad l_1 = 1, l_2 = 2, l_3 = 3, l_4 = 4, l_5 = 5, l_6 = 6, l_7 = 7, l_8 = 8, l_9 = 9
\end{aligned}$$

Now that we have optimal strategies for each number of stops, we can calculate the total race times for each of these using (2.8), and then compare them to see which is the fastest strategy. We will also be considering the 0-stop strategy whereby the race is simply treated as one 10-lap stint.

Figure 2.1 shows a plot of the total race times for each of the strategies described above. The fastest of these strategies is simply the one that minimises the total race-time, hence for this track a 1-stop strategy is the fastest with a total race time of 916.35 seconds. We can also see that a 0-stop strategy is only slightly slower at 917.75 seconds. We can clearly see that completing more than 2 pit stops results in a significantly slower race time.

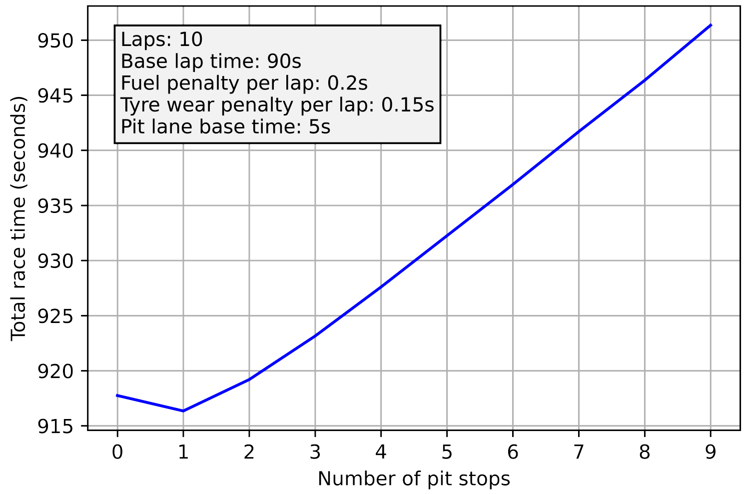

Figure 2.1: (2.5 - Example) Total race times of optimal r-stop strategies for $r \leq 9$. Race consists of 10 laps with $t_b = 90, F = 0.2, T = 0.15, t_s = 5$.

2.6 Sensitivity to parameters

We have shown how to find the optimal laps on which to make pit stops for an r-stop strategy, as well as how we to many the best number of pit stops to make during a race. However, a Strategy Engineer may be interested in knowing how the optimal strategy is affected by the various parameters we have considered so far.

We have four variables that affect the optimal strategy for a given race: $l_{r+1}, T, F,$ and t_s. Some simple intuition tells us we would expect changing these variables to have the following effects:

l_{r+1} : Increasing l_{r+1} is likely to increase the optimal number of stops. This is because a longer race will require the car to carry more fuel and have worse tyres on average, which slows the car down. Making more stops reduces the amount of tyre wear and fuel in the car on average, thus reducing the overall fuel load and tyre wear penalty.

T : Increasing T means there is a greater penalty for using a given set of tyres for more laps in the race. A large T would make it desirable to have fresh tyres more often during the race, meaning more pit stops may have to be made.

F : Increasing F means there is a greater penalty for carrying large amounts of fuel. As a result, it is desirable to carry as little fuel as possible, meaning the optimal number of stops is likely to increase.

t_s : Increasing t_s would mean pit stops would take longer to complete. As a result, it is desirable to make fewer stops so as to waste as little time in the pits as possible. We would expect the optimal number of stops to decrease as t_s increases.

For all of the above, the opposite indeed holds. For instance, a race with fewer laps would likely see the optimal strategy having fewer stops.

To demonstrate the effects of changing these parameters, we will make some modifications to the race defined in the previous example.

Increasing the race length

We will double the length of the race to 20 laps. As said previously, we would expect more stops to be necessary in a longer race.

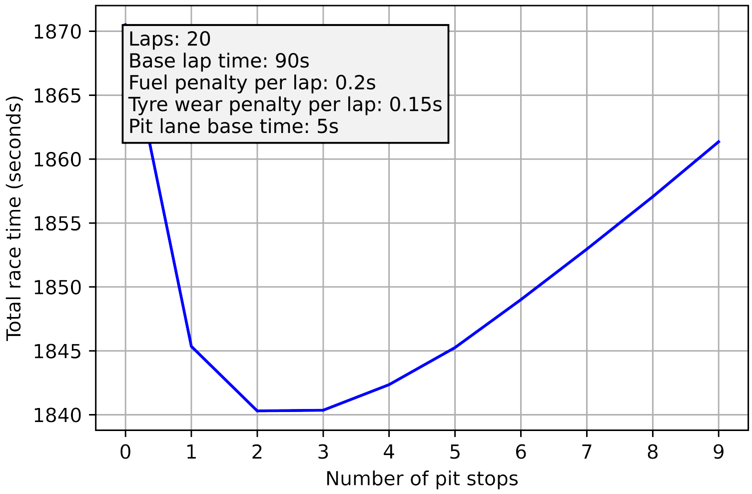

Figure 2.2: Race length has been doubled to 20 laps.

Figure 2.2 shows that completing either 2 or 3 stops will result in the fastest total race time for this extended race of 20 laps. It turns out that the optimal 2-stop strategy has a total race time of 1840.3 seconds, whereas the optimal 3-stop strategy is only 0.05s slower at 1840.35 seconds! The optimal 1-stop strategy now has a total race time of 1845.35 seconds, which is just over 5 seconds slower than the fastest possible strategy at this track.

Decreasing tyre wear penalty

We will decrease the effect of tyre wear by reducing the penalty to $T = 0.05$. Because there is less of a penalty for running longer stints, we would expect the optimal number of stops in the race to decrease as a result of this change.

After reducing the tyre wear penalty, we can see on Figure 2.3 (blue) that the fastest strategy for this race would be a 0-stop strategy. The total race time for the 0-stop strategy is 913.25s, which is 4.5s faster than the 0-stop strategy in the original example. The optimal 1-stop strategy, which would involve making a pit stop on lap 6, has a total race time of 914.25s which is 1 second slower than the 0-stop strategy. With the reduced tyre wear, this 1-stop strategy is quicker than the optimal 1-stop strategy in the original example, as we would expect. We can also see that as we increase the number of pit stops, the race times for the original and reduced tyre wear races converge. This is because the tyre wear penalty is not applied on the first lap of the stint, so on the 9-stop strategy where every lap is completed on fresh tyres, the lap times for this race and the original are the same.

Increasing fuel penalty

We will now increase the fuel penalty to $F = 0.5$. If we interpret this as the cars requiring more fuel to complete a lap, then across the race the cars will be carrying more fuel on average. We would therefore expect more stops to be made.

Figure 2.3 (orange) shows that for this race, the optimal strategy is still the 1-stop strategy, but we can see that it is only just quicker than the 2-stop strategy. Conversely, the 0-stop strategy

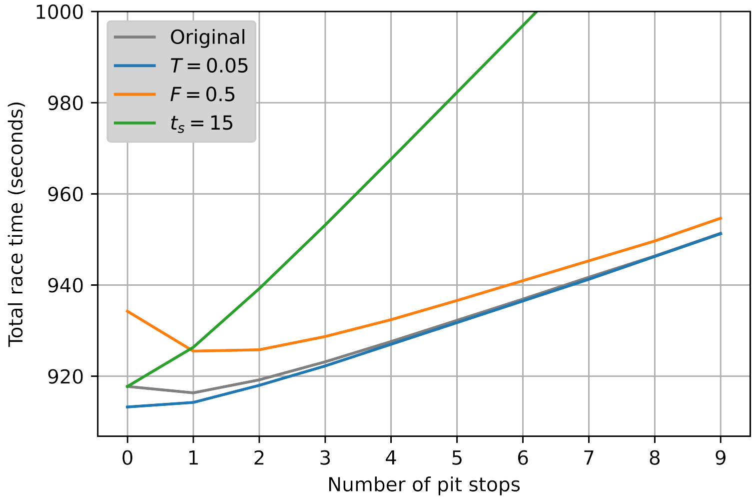

Figure 2.3: Combined plot of all changes that have been made separately. Original example (grey); tyre wear penalty has been reduced to $T = 0.05$ (blue); fuel load penalty has been increased to $F = 0.5$ (orange); pit stop base time has been increased to $t_s = 15$ (green).

is now significantly slower than making 1 or 2 stops, making it even worse than completing a 4-stop strategy! With the increased fuel penalty, the optimal 1-stop strategy now involves making a pit stop on lap 5, rather than lap 6, and the total race time is increased to 925.5s. The optimal 2-stop strategy has a total race time of 925.8s, which is 0.3s slower than the optimal 1-stop strategy. Similarly for the race with reduced tyre wear, the race times for the race with the higher fuel penalty begin approach the original race times as the number of pit stops increases. However, because the fuel penalty is considered on every lap in the stint, the race times are never able to converge. In the 9-stop strategy, where every stint is exactly one lap, each lap in the race with the additional fuel penalty is 0.3s slower than the original race.

Increasing pit stop time

By increasing the base pit lane time to $t_s = 15$ seconds, we would expect fewer pit stops to be made as there is now a greater penalty for making a pit stop.

Figure 2.3 (green) shows us that the 0-stop strategy is now the fastest strategy after the pit lane time had been increased. Because this change just adds 10 seconds to the total race time in the original example for each pit stop that is made, the 0-stop strategy avoids this additional penalty altogether. As such, the race time for the 0-stop strategy is 917.75 seconds, as in the original example, whereas the race time for the 1-stop strategy has been increased by 10 seconds to 926.35 seconds. We can also clearly see how much slower the races are with multiple stops when compared to the original race.

2.7 Example: N-lap race

Having seen the method for finding the best strategy for a short race, we can very easily extend it to find the best strategy for a race of a more typical length.

Suppose we wish to complete a race of 63 laps of a given track, where a Strategy Engineer has estimated the following parameters: $T = 0.05, F = 0.13, t_s = 17, t_b = 75$. Considering strategies with at most 8 pit stops, and assuming the simple pit stop model (2.7), we wish to find the fastest strategy for this race.

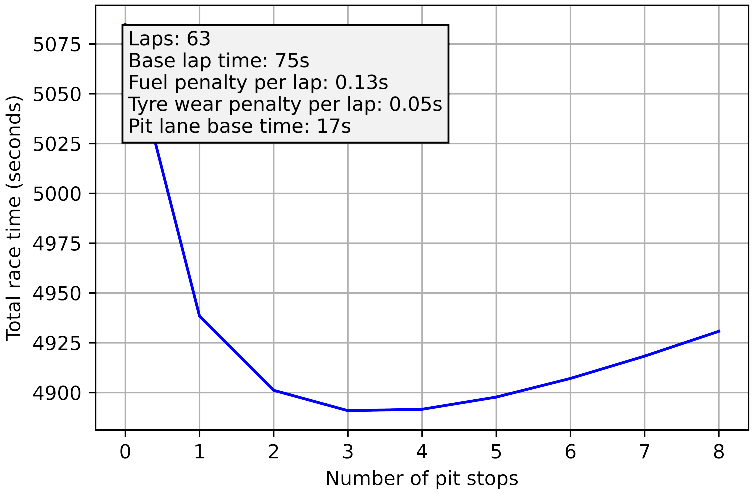

Figure 2.4: Total race times of optimal strategies for given number of pit stops.

Figure 2.6 shows us that the optimal 3-stop strategy is the fastest strategy for this race. This strategy involves stopping on laps 18, 33, and 48 (notice that each stint is 15 laps, apart from the first), and the total race time is 4890.93s. We can see that the optimal 4-stop strategy (pitting on laps 15,27,39, and 51) is only slightly slower, as its race time is 4891.61s. Completing either 2 or 3 pit stops in this race is clearly the best option.

2.8 Summary

- We can model a car's lap time around a track by using

$$t = t_b + Tl + F(l_{\text{max}} - l), \qquad l = 0, 1, \dots, l_{\text{max}} - 1,$$

  where t is the lap time, t_b is the base lap time with minimal fuel and fresh tyres, T is the tyre wear penalty per lap completed on the tyres, F is the fuel load penalty per lap of fuel in the car, l_{max} is the number of laps in a stint, and l is the number of laps completed.

- The time taken to complete a stint of l_{max} laps is given by

$$t_S = \sum_{l=0}^{l_{\text{max}}-1} t_l = l_{\text{max}} t_b + \frac{l_{\text{max}}}{2} \left[T(l_{\text{max}} - 1) + F(l_{\text{max}} + 1) \right], \qquad l = 0, 1, \dots, l_{\text{max}} - 1,$$

  where t_l is the lap time on lap l, at which point only $l - 1$ laps have been completed in the in the stint.

- A race of l_{r+1} laps can be split into $r + 1$ stints: $[l_0 + 1, l_1], [l_1 + 1, l_2], \dots, [l_r + 1, l_{r+1}]$, where pit stops are made on laps $l_1, l_2, \dots l_r$. The laps on which pit stops are made define an r-stop strategy.

- The time taken to complete a race of l_{r+1} laps using an r-stop strategy is given by

$$\begin{aligned} t_R(l_1,\dots,l_{r+1};F,T,t_s,t_b) = l_{r+1}t_b + F\frac{l_1(l_1+1)}{2} + T\frac{l_1(l_1-1)}{2} \\ + \sum_{i=2}^{r+1}\Bigg[t_s + \frac{1}{2}(l_i - l_{i-1}) + F\frac{(l_i-l_{i-1})(l_i-l_{i-1}+1)}{2} \\ + T\frac{(l_i-l_{i-1})(l_i-l_{i-1}-1)}{2}\Bigg]. \end{aligned}$$

- Given a race of l_{r+1} laps, if we wish to complete an r-stop strategy, the optimal laps on which to pit can be found recursively using

$$l_i = \frac{i}{i+1}l_{i+1} + \frac{1}{2(i+1)(T+F)}, \qquad \forall i = 1,\dots,r.$$

- The time taken to complete a pit stop is

$$t_p = t_s + \frac{1}{2}(l_i - l_{i-1}),$$

  where t_p is the time taken to complete the pit stop, t_s is the time taken to drive through the pit lane without changing tyres or adding fuel, and it takes 0.5s to add a single lap of fuel to the car in the pit stop.

- The optimal race strategy is affected by the race length l_{r+1}, the tyre wear penalty T, the fuel load penalty F, and the pit lane base time t_s.

Chapter 3

Applying the simple model

In this chapter, our aim is to apply our simple to a real life Formula 1, as if we were a strategy engineer for a given team. We will first estimate the necessary parameters in the model and then devise a number of strategies for the race. These strategies will then be compared to the winning strategy for the race.

In order to apply our model, we need to choose a Formula 1 race that meets the following criteria:

1. The race must have been dry throughout, so no wet tyre compounds have been used.

   Currently, we have not modelled the effect of the weather on lap times as there is great variability in the effect of rain on the lap time around a track.

2. The race must not have featured either the Safety Car or Virtual Safety Car.

   In the event of a crash or some other scenario that may endanger drivers or marshals on the track, either the SC or VSC can be deployed. The SC groups all cars together to the same point on track at a reduced speed, whereas the VSC essentially puts a speed limit on all cars for a given amount of time. If either of these occur at some point during the race, it gives teams the opportunity to make a pit stop without losing as much time as they would normally. This is because the pit lane speed limit remains the same but cars are travelling slower around the track. For instance, it may take 25 seconds for a car to pit and change tyres, and this is the same for normal racing conditions and when the SC/VSC is deployed, but the lap time may increase from 80s to 120s because of the SC or VSC. As a result, pit stops are relatively quicker under the VSC or SC, thus making new strategies viable. These events are highly situational, hence it is very difficult to create a model that adequately captures the effects of either the SC or VSC on race strategies.

3. There must not have been a red flag.

   When red flags are deployed on a race track, all cars must return to the pits immediately. This typically happens after a large crash or torrential rain. During a red flag, F1 teams are allowed to change the tyres on the car without penalty, as there is no time loss when doing so. As such, red flags can greatly alter the strategies used by teams in a race.

With these criteria, our race of choice will be the 2020 70$^{\text{th}}$ Anniversary Grand Prix, held at Silverstone Circuit in Great Britain. This race consisted of 52 laps and saw Max Verstappen take victory in the Red Bull Racing car ahead of Lewis Hamilton and Valtteri Bottas, both of whom drove for Mercedes.[8] During the race, refuelling was not permitted, and drivers were required to use at least 2 of the available 3 dry tyre compounds at some point in the race. This race was notable due to the generally high levels of tyre wear that were expected during the race as Pirelli, the tyre suppliers for Formula 1, opted to use softer tyre compounds than what would normally be expected for a track such as Silverstone. The tyre compounds that were chosen by Pirelli were the C2 (hard), C3 (medium), and C4 (soft) compounds.

3.1 Extraction of data

To estimate the tyre wear and fuel load penalties, we will be considering the lap times[9] from various drivers throughout the race where possible. We will first find an estimate for F so all lap times can be 'fuel-adjusted'.

Finding F

In order to find a good estimate for F using data from the race, we will look for stints of a similar length on the same tyre compound. Formula 1 drivers regularly adapt their driving style throughout the race depending on various circumstances.[10] If they have been asked to complete a long stint on a set of tyres, they may then have to drive more conservatively in order to ensure the tyres last as long as necessary. By driving conservatively, drivers will brake earlier, carry less speed through corners, and be more careful when accelerating out of a corner which combine to give a slower lap time than what may be possible. Conversely, if a driver has been asked to perform a shorter stint on a set of tyres, they would likely drive more aggressively. As a result, by finding stints of similar length on equivalent tyre compounds, we can assume that the drivers have employed a similar driving style across the stint, thus making the stints directly comparable (or as close to directly comparable as possible).

An additional consideration when choosing stints to estimate the value of F is the amount of traffic a car is in. The speed of a Formula 1 car is highly dependent on its aerodynamic performance, and when a car follows another car closely (typically within 2 seconds), turbulence from the car ahead negatively affects the airflow for the following car, which slows the following car through corners. As a result, by attempting to find drivers who were in clear air for the majority of a stint, their lap times will be more directly influenced by the fuel load and tyre wear.

We are also looking for sets of lap times that appear to follow the same trend. That is, if we find a stint where the lap times are generally getting slower, then we should aim to compare these laps to another set of laps that also appear to be getting slower at a similar. This will help to find sets of lap times where the driver is employing a similar driving style, allowing the laps to be more directly comparable.

In our chosen race, Valtteri Bottas completed two separate stints on the hard compound tyres, with the stints being 19 and 20 laps respectively. To find our estimate of F, we will find the average difference between lap times in each of these stints, and then divide this difference by the number of laps between the first lap in each stint. We will then check this for other drivers in the race who also completed multiple stints on the same compound with the same length. Because of the nature of Formula 1 circuits, we will be disregarding the in-laps (laps on which the driver entered the pit lane for a pit stop) and out-laps (laps on which the driver exited the pit lane after a pit stop).

A plot of Bottas' lap times across these stints is given in Figure 3.1, although we have discarded the last 3 laps of Bottas' third stint as they were both noticeably slower than the other laps in the stint, and also the first 2 laps and last 3 laps of Bottas' second stint as they were noticeably faster than the preceding laps.

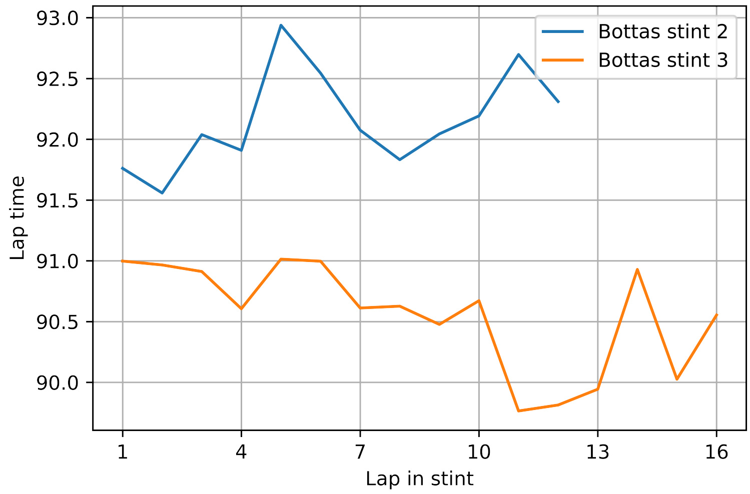

Figure 3.1: Lap times for Valtteri Bottas at the 2020 70$^{\text{th}}$ Anniversary Grand Prix. Both stints were completed on the C2 (hard) tyre compound, with stint 2 (blue) being completed before stint 3 (orange). First lap of stint 2 was lap 17, while first lap of stint 3 was lap 34.

Our estimate of F is given by

$$\hat{F} = \frac{\frac{1}{L_1}\sum_{i=1}^{L_1} l_{1,i} - \frac{1}{L_2}\sum_{j=1}^{L_2} l_{2,j}}{L_{\text{diff}}} = \frac{\bar{l}_1 - \bar{l}_2}{L_{\text{diff}}}, \tag{3.1}$$

where $l_{1,i}$ is the i^{th} lap time from stint 1, $l_{2,j}$ is the j^{th} lap time from stint 2, L_1 is the number of laps in stint 1 (excluding in-laps, out-laps, and other discarded laps), L_2 is the number of laps in stint 2, and L_{diff} is the number of laps in the race between the start of stint 1 and stint 2.

Using Bottas' lap times, we find the following estimate for F

$$\begin{aligned} \hat{F} &= \frac{92.1581\ldots - 90.5570\ldots}{17} \\ &= 0.09418\ldots \\ \implies \hat{F} &= 0.0942 \qquad \text{(4d.p.)}. \end{aligned}$$

Therefore, we would expect lap of fuel onboard the car to increase the lap time around the track by 0.0942s. To check this, we will check perform the same calculation for a number of other stints from other drivers.

Driver	Tyres	Stint 1 average	Stint 2 average	L_{diff}	$\hat{F}$
Bottas	Hard	92.1581	90.5570	17	0.0942
Verstappen	Hard	91.7294	90.1294	22	0.0724
Hamilton	Hard	92.0592	89.468	27	0.0960
Hulkenberg	Medium	92.1593	91.2494	15	0.0607
Albon	Hard	92.7129	90.3248	24	0.0995
Ricciardo	Medium	93.1727	92.3204	14	0.0609
Kvyat	Hard	93.7914	92.1758	19	0.0850
Norris	Hard	92.6939	91.4611	17	0.0725
Sainz	Hard	93.3286	91.1454	34	0.0642

Table 3.1: Estimates of $\hat{F}$ from various drivers. Note that the column names 'Stint 1' and 'Stint 2' do not necessarily mean these were the driver's first and second respective stints in the race, and also note that some lap times have been excluded from the full set of times to give us more comparable lap times.

If we take the mean of the estimates of $\hat{F}$, we get $\overline{\hat{F}} = 0.0784$. Some estimates are noticeably smaller than others, such as Hulkenberg's, Ricciardo's, and Sainz's which lowers the average. If we were to discount these estimates, we would get $\overline{\hat{F}} = 0.0866$, from which we could use the estimate $\hat{F} = 0.09$. However, according to analysis from some F1 teams, the time penalty per lap of fuel is around $F = 0.112$ seconds![11] There may be a number of reasons for this difference. For instance, our calculation relies on drivers being clear of traffic ahead for a long portion of a stint, which is a reasonably rare occurrence. Of course, it is a race, so if a driver knows there is another car not too far ahead, they will likely try to overtake the other driver (if they have good reason to do so), thus forcing them to follow closely for a number of laps before being able to make an overtake. Additionally, towards the end of races, where cars tend to be more spread out, a team may instruct a driver to ease off a little in order to ensure the car makes it to the end of the race. This means the lap times in the final stint are potentially slower than what they could be if the driver where to drive were to drive how they did earlier in the race. The slower lap times means there is a smaller difference in the average lap time between stints, which in turn lowers the estimate of F.

Considering the estimate of F from the table above and that this is likely to be an underestimate, we will use $\underline{\hat{F} = 0.1}$ as our estimated value of the fuel load penalty.

Finding T

In the model described in the first chapter, we only had one value of T to consider, as we assumed there was only one tyre compound available. However, as mentioned previously, F1 teams can choose between 3 different tyre compounds which each have their own unique characteristics. The three compounds at each track are referred to as 'hard', 'medium', and 'soft' tyres. The hard tyres are the most durable, but are initially the slowest of the 3 compounds. The soft tyres are the fastest outright, but quickly wear out, eventually becoming slower than the other two compounds. The medium tyres provide a middle ground between the hard and soft tyres, and typically offer the best compromise between speed and durability.

One way of modelling the differences between the 3 compounds is a simple adjustment to (2.2):

$$t_l = t_b + T_c l + F(l_{\max} - l) + p_c, \tag{3.2}$$

where T_c is the per lap tyre penalty for tyre compound c, and p_c is the constant time penalty for tyre compound c. Using this model, we would set $p_S \equiv 0$, meaning the soft tyre compound has no constant lap time penalty. We would also assume that $p_S \leq p_M \leq p_H$, meaning the soft tyres are quicker than the medium tyres, which are quicker than the hard tyres, at the start of a stint (i.e. $l = 0$). We would also expect $T_H \leq T_M \leq T_S$ to ensure the hard tyres have the lowest tyre wear rate while the soft tyres have the highest.

To begin finding an estimate for T_c and p_c for each compound c, we need to find some candidate stints to find the data. As before when estimating F, we will ideally use stints in which the driver was not affected by traffic ahead, or only minimally so. After finding the candidate stints, we will then "fuel-adjust" them, whereby we subtract

$$\hat{F}(l_{r+1} - L)$$

from each lap time, where $\hat{F} = 0.1$ is our previous estimate of F, l_{r+1} is the number of laps in the race, and L is the number of laps completed in the race. After this, we will fit a linear regression model to the lap times to find T_c, which is given by the coefficient β_1 from $y = \beta_0 + \beta_1 x$, where x is the number of laps completed in the stint. To find p_c, we will first need to find an estimate for the base lap time, t_b.

Hard tyres (C2)

Our estimate of T_H will come from Verstappen's first stint in the race, where we consider laps 14 to 25.

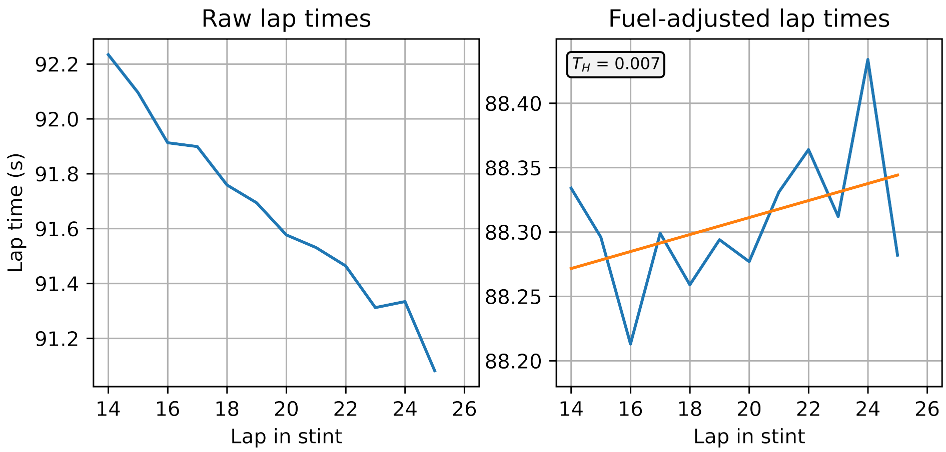

Figure 3.2: Plots of Verstappen's raw (left) and fuel-adjusted (right) lap times from his second stint on the medium tyres. Orange line shows fitted values of linear regression model.

Figure 3.2 shows that the fuel-adjusted lap times on the hard tyres tend to increase over time. For this fitted regression model, we find that $\hat{T_H} = 0.0065979$. This can be rounded to $\hat{T_H} = 0.007$ which suggests that the hard tyres get slower by 0.007s for every lap completed on them.

Medium tyres (C3)

Our estimate of T_M will come from Sainz's second stint in the race, where we consider laps 24 to 33.

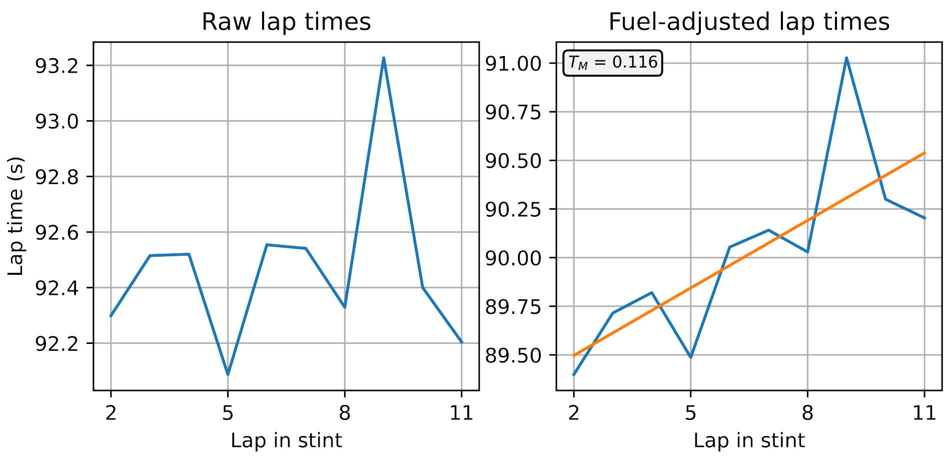

Figure 3.3: Plots of Sainz's raw (left) and fuel-adjusted (right) lap times from his first stint on the hard tyres. Orange line shows fitted values of linear regression model.

Much like with the hard tyres, we can see from Figure 3.3 that the fuel-adjusted lap times tend to increase for the medium tyres. We find that $\hat{T}_M = 0.11568485$ which can be rounded to $\hat{T}_M = 0.116$.

Soft tyres (C4)

Unfortunately for this race, the soft tyres were only used by one driver for one stint. If we were to use the method as for the other tyre compounds, we would find that $\hat{T}_S \approx 0.13$, which is only slightly larger than our estimate of T_M. These laps were completed at the end of the race when the driver's position was essentially fixed for the rest of the race, as the cars ahead were too far ahead to catch up, and he was under no threat from behind. As such, the driver was likely being conservative to ensure the car made it to the end of the race, meaning the tyres were not wearing down as quick as they would normally, thus reducing the estimate of T_S.

However, we are able to use the lap times from the free practice sessions[12],[13] (which is actually the where these values would be obtained from) to acquire further estimates of T_S. Of course, we do not know exactly how much fuel is onboard each car in the practice sessions, but our fuel-adjustment method will alter the lap times so that they were completed with the same amount of fuel (not necessarily no fuel), which enables us to estimate T_S as before.

For our calculations, we will be using Sebastian Vettel's lap times from the second practice session, specifically laps 15 to 21 from his session, as he had the longest clean stint on soft tyres throughout the session.

Figure 3.4 shows that the soft tyres quickly lose pace, with our estimate being $\hat{T}_S = 0.39732143$ which we could round to $\hat{T}_S = 0.4$. This is much greater than $\hat{T}_M = 0.116$ and is in agreement with our knowledge that the tyres chosen for this race were arguably too soft.

Overall, we have found that $\underline{\hat{T}_H = 0.007, \hat{T}_M = 0.116, \hat{T}_S = 0.4}$.

Finding t_b and p_c

To find t_b, we simply take a driver's lap times and subtract the fuel and tyre penalties using the values we have calculated. However, each car will have a different base lap time due to different engines, aerodynamics, etc., so we will need to estimate the average pace difference between each of the different cars and drivers. To do this, we will take the gap from the leader to the

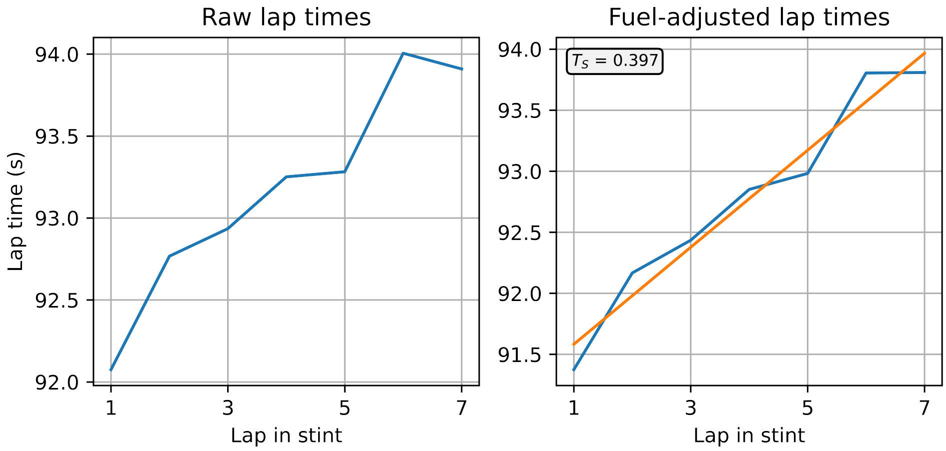

Figure 3.4: Plots of Vettel's raw (left) and fuel-adjusted (right) lap times from his FP2 stint on the soft tyres. Orange line shows fitted values of linear regression model.

lead driver of a given team at the end of the race and divide by the number of laps the driver completed. For instance, Hamilton was the lead Mercedes car in the race and finished 11.326s behind Verstappen in the Red Bull, meaning he lost ≈ 0.218s per lap on average.

We will find a number of base lap times for each tyre compound (where possible) to find both t_b and p_c for $c = S, M, H$. We will take the average lap time from a tyre wear and fuel-adjusted stint from a number of drivers.

Driver	Team adjust	Tyre	Laps	Average adjusted lap time	$\hat{t}_b$
Verstappen	0	Hard	14-25	88.1784	88.1784
Leclerc	0.5634	Hard	20-44	88.8490	88.2856
Norris	1.2605	Hard	32-40	89.8538	88.5933
Hamilton	0.2178	Hard	43-51	88.8330	88.6152
Kvyat	1.340	Hard	2-18	89.4284	88.0884
Leclerc	0.5634	Medium	2-17	87.7176	87.1542
Sainz	1.2605	Medium	24-33	89.3796	87.9551
Ricciardo	1.2456	Medium	2-13	87.8687	86.6231
Norris	1.2605	Medium	2-11	88.1137	86.8532
Stroll	0.8180	Medium	2-17	87.5105	86.6925
Hulkenberg	0.8180	Soft	46-52	88.7703 (excl. lap 51)	87.9523

Table 3.2: Estimates of the base lap time t_b for each tyre compound. Average lap time is found for each driver in the given lap range, which is then adjusted according to the average amount of time they lost to the leader per lap in the race.

If we take the average base time for each tyre compound from Table 3.2, we get $\hat{t}_b = 88.3522$ from the hard tyres, $\hat{t}_b = 87.0557$, while the estimate from the soft tyres, $\hat{t}_b = 87.9523$, is not very useful considering there is only one value. From this, we could say that the medium tyres are about 1.3 seconds quicker than the hard tyres when brand new. According to data from Pirelli, the hard tyres are around 0.7s slower than the medium tyres, which are a further 0.6s slower than the soft tyres in qualifying, although this would be different under race conditions.[14]

Given the level of variability within the data, it is very difficult to accurately say what t_b and p_c are, although we can make some very rough estimates. If we say that $\hat{t}_b = 86.5$, then

$\hat{p}_S = 0, \hat{p}_M = 0.6, \hat{p}_H = 1.8$, then we would argue, according to estimates of T_c, that after just 2 laps the soft tyres would be worse than the medium tyres, and after around 5 laps they would be worse than the hard tyres. Meanwhile, it would take the medium tyres around 12 laps before they become slower than the hard tyres.

Finding t_s

To find an estimate for t_s, we need to consider both the in-lap and out-lap. Currently, our model assumes that the time taken to complete a pit stop is simply added to the lap on which the pit stop is made. However, because the pit entry is before the timing line and the pit exit is after, time is lost on 2 laps whenever a pit stop is made. To account for this, we will simply estimate the lap time on both the in-lap and out-lap, and then find the difference between these and the lap times recorded in the race. We will then add these differences together to find the total time lost due to a pit stop.

To make calculations easier, we can make use of the fact that our tyre wear and fuel load laws are linear, and so we can add $T - F$ to the previous lap time to find an estimate for the in-lap time. For instance, we have $\hat{F} = 0.1$ and $\hat{T}_M = 0.116$. Therefore, every lap completed on the medium tyres should be about 0.016s slower than the previous lap, as the fuel load decreases but the tyre wear penalty increase. For the hards, we hard expect each lap to be about 0.093s quicker than the previous, and for the softs we would expect each lap to be about 0.3s slower than the previous lap.

Driver	In-lap		Out-lap		Total time lost
	Actual	Estimated	Actual	Estimated	
Raikkonen	92.973	93.692	113.308	93.126	19.463
Grosjean	92.097	92.500	116.257	92.310	23.544
Verstappen	90.444	90.989	111.315	91.582	19.188
Hulkenberg	91.426	91.270	110.849	90.593	20.412
Bottas	90.071	90.344	110.959	91.091	19.617
Norris	92.315	92.843	112.180	92.440	19.212
Albon	90.440	91.040	111.242	91.138	19.504
Russell	92.492	93.346	111.343	92.068	18.421

If we take the average of the final column, we get $\hat{t}_s = 19.920\ldots$, so we could choose our estimate of the time taken to complete a pit stop as $\underline{\hat{t}_s = 20}$ seconds.

Other considerations

One additional aspect we need to consider for our strategy calculations is the additional time taken to complete the first lap of the race. Because Formula 1 races begin with a standing start, the cars have to accelerate to their normal speeds while under racing conditions, so they will lose time when compared to a regular racing lap. To estimate this, we will consider the additional lap 1 time penalty for the leading car by estimating their first lap if they were to start the race at full speed. Valterri Bottas was in the lead at the end of the first lap with a lap time of 92.486s and his second lap was 91.384s. Because he started on the medium tyres, we would expect his first lap to be 0.016s quicker than his second lap, which would be 91.400s. Therefore, we estimate that he lost about 1.1 seconds on the first lap because of the grid start. When finding the race time of the fastest possible strategy, we will add $\hat{p}_{\text{lap 1}} = 1.1$s onto the first lap time.

3.2 Finding the optimal strategy

To find the optimal race strategy, we will use the following values:

- $\hat{t}_b = 86.5, \hat{t}_s = 20, l_{r+1} = 52$;
- $\hat{F} = 0.1$;
- $\hat{T}_S = 0.4, \hat{T}_M = 0.116, \hat{T}_H = 0.007$;
- $\hat{p}_S = 0, \hat{p}_M = 0.6, \hat{p}_H = 1.8$;
- $\hat{p}_{\text{lap 1}} = 1.1$.

Because we have multiple tyre compounds available to us, we cannot use (2.6) to find the optimal strategy. Instead, we will calculate the total stint time for all possible stints, and then, using Python, find the combinations of stints that minimise the total race time. We should note that because refuelling is not allowed, the time taken to complete a pit stop will be exactly $\hat{t}_s = 20$ seconds, ignoring any random error time as discussed in Section 2.4. We also need to ensure that the strategy we select uses at least 2 of the available tyre compounds, as per the current rules in F1.

After considering all possible strategies that feature at least 2 of the available tyre compounds, we find that the following strategy is the quickest: **M13 H52**. Because refuelling is not allowed, the strategy H39 M52 would be equally fast according to our calculations. After considering the additional lap 1 penalty of 1.1 seconds, the total race time for this strategy would be $t_R = 4749.135$s, which is equal to 1hr 19mins 9.135s. If we compare this to the race winning time of 1hr 19mins 41.993s, we can see that our optimal strategy is just over 30s quicker than the race winning strategy!

If we were to consider the fastest 2-stop strategy, we would pick M12 M24 H52 (or any rearrangement). The total race time for this strategy would be 1hr 19mins 16.7s, which is quicker than the race winning strategy by around 25 seconds.

We can also calculate the race time of the race winning strategy (H26 M32 H56) using our estimated values, 1hr 19mins 28.645s. This is just 13.348s off the true value: which is an error of only 0.28%! This suggests to us that this model fits this race well.

3.3 Goodness of fit

To see how well our model fits, we will compare the race winning strategy to our optimal and fitted strategies in more detail.

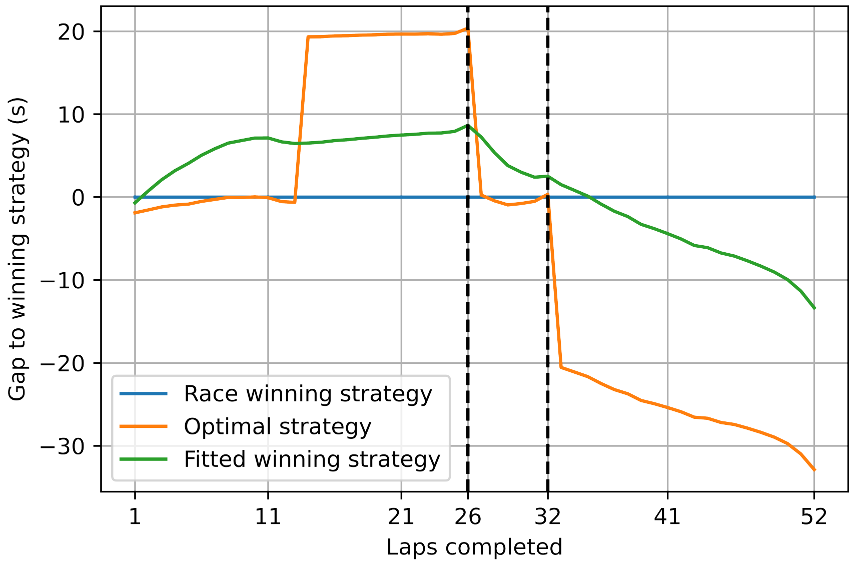

Figure 3.5: Race winning strategy compared to fitted and optimal strategies. Plot shows lap by lap gap to the winning strategy (seconds) for the other strategies. A negative delta means a given strategy is ahead of the winning strategy on that lap.

Figure 3.5 shows where the fitted and optimal strategies deviate from the race winning strategy. We can clearly see how far ahead the optimal strategy is compared to the winning strategy, while the fitted winning strategy is less far ahead. Considering only the fitted strategy, in the first stint of the race our model falls behind Verstappen across the first 11 laps, but loses very little until the first pit stop is made. When on the medium tyres, our model predicts that Verstappen would drive much faster than he did, so our fitted model catches up until the final pit stop is made. After the final pit stop, our fitted model overtakes Verstappen and then gradually pulls away towards the end of the race. At this point in the real race, Verstappen was comfortably in the lead, and so had little reason to push to the extent we predicted the tyres would allow.

From Figure 3.5, we see that our model is unable to capture some features in Verstappen's race, such the points in the race where he is pushing (laps 1-11) and also when he is backing off (laps 32-52). These sets of laps that deviate from our predicted lap times can be explained by the presence/absence of other cars close to Verstappen. At the start of the race, Verstappen was closely following Hamilton and Bottas who were both on the medium tyres, and so he may have felt the need to push to make sure he could keep up with them. His laps in this part of the race closely match the optimal strategy, which used medium tyres for the first stint, demonstrating the extent to which he was pushing. After making his first stop, Verstappen had come out ahead of both Hamilton and Bottas, so he only needed to maintain his lead over them both. Bottas and Verstappen both made their final stop at the same time, leaving Hamilton out in the lead. If Hamilton were to pit, he would have ended up behind Verstappen, so Verstappen only needed to match the pace of Hamilton to guarantee himself track position once Hamilton pitted. After Hamilton did pit, he emerged so far behind Verstappen that he was not able to threaten Verstappen for the lead, so Verstappen was able to ease off for the rest of the race.

The drop off in Verstappen's lap times at the end of the race can be summarised with the following quote from Alain Prost: *"It is ideal to win the race at the lowest speed possible."*[15]

3.4 Extending the model

Now that we have applied our model to a real life scenario and also discussed where it fails to capture certain features in a driver's lap times, we can consider ways of extending this model to improve how well it fits.

Relation between fuel load and tyre wear

When a car is carrying more fuel, it experiences greater inertia when accelerating, decelerating, and changing direction which in turn demands more grip from the tyres. As a result, the tyres on a car will wear out quicker when the car is heavier. We can model this with a simple adjustment to (2.2):

$$t_l = t_b + Tl + F(l_{\max} - l) + \alpha T \sum_{n=1}^{l} (l_{\max} - n), \tag{3.3}$$

where α is some positive constant. This works by adding a cumulative sum for every lap that is completed on the tyres that is based on the number of laps of fuel remaining in the car. As the stint progresses, the additional term converges to $\alpha T \frac{(l_{\max} - 1)l_{\max}}{2}$. If a team were to use this model, they would become discouraged from running excessively long stints on a single set of tyres, as the additional penalty per lap is an increasing function of the number of laps in the stint.

Driving style

Should the situation necessitate it, a Formula 1 driver may choose to driver more aggressively or more cautiously at some stages in a race. By driving more aggressively, faster lap times can be achieved but tyre wear may be increased which will slow the car down in the long run. Conversely, by driving more cautiously, lap times are slowed in the short term but tyre can be reduced, enabling a driver to prolong a stint. We can model this with the following

$$t_l = t_b + d_P Tl + F(l_{\max} - l) + \beta(1 - d_P), \tag{3.4}$$

where d_P is the "driver-push factor" that can take values between, say, 0.8 and 1.2, and β is an adjustment constant for the additional lap penalty $1 - d_P$ which we can interpret as the driver aiming for lap times that are some constant amount slower or faster than normal. This model assumes that a driver will push at the same rate throughout a stint, although it would be possible to adapt this to the driver pushing on a subset of laps within a stint. With this model, if a driver were to drive more aggressively, meaning $d_P > 1$, they would be faster at the start of the stint, but as the stint progresses the additional tyre wear begins to dominate, resulting in ultimately slower lap times and vice versa for $d_P < 1$.

Variable fuel consumption

Because refuelling is banned in Formula 1 races, drivers need to manage their fuel consumption during the race. At some stages of the race, such as when chasing a car ahead, it may be worth using more fuel than normal. This comes at the cost of needing to save fuel later in the race. One reason for using a fuel at a decreased rate could be to save some time in the pits as less fuel needs to be added. Of course, lap times will be slower as the engine will be turned down, but in some cases the immediate track position may be worth it. We can apply this to our simple model with the following modification:

$$t_l = t_b + Tl + F\left(l_{\max} - l\frac{l_{\max}}{l_{\max}+D}\right) + \frac{\alpha D}{l_{\max}+D}, \tag{3.5}$$

where $D \in \mathbb{Z}$ is the stint length adjustment constant. If $D > 0$, then we would extend the length of the stint by D laps by using less fuel in each lap, and vice versa for $D < 0$. Instead of using 1 lap of fuel per lap, we use $\frac{l_{\max}}{l_{\max}+D}$ laps of fuel which ensures all fuel is used by the end of the stint. The additional constant term measures the effect on lap time from having the engine turned up or down, and it is tuned by the constant α. This constant is proportional to the ratio between D and the length of the stint, as a larger ratio suggests more extreme engine settings would need to be used to either save or use extra fuel, in turn affecting the lap time more significantly. We assume that the engine mode is fixed throughout the stint, and so the lap time effect is fixed across the stint.

3.5 Real life strategies

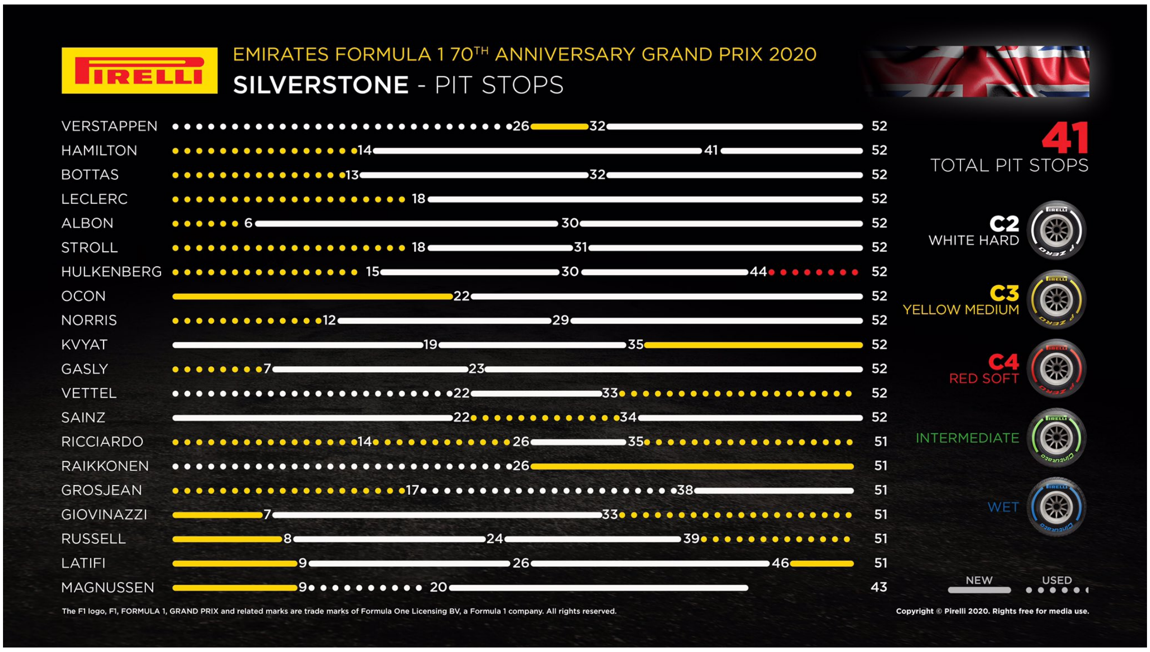

Figure 3.6: Strategies used at the 2020 70th Anniversary Grand Prix.[16]

Figure 3.6 shows a summary of the strategies used at the 2020 70th Anniversary Grand Prix. We can see that Verstappen completed 26 laps on a used set of hard tyres, then did 6 laps on mediums before finishing the race with a 20 lap stint on hard tyres. Conversely, the Mercedes drivers of Hamilton and Bottas started on used mediums before completing 2 stints on the hard tyres. We can see from the chart that most teams completed the majority of the race on hard tyres, as these were generally considered to be the best race tyre at this race. Only Nico Hulkenberg used the soft tyres during this race, as mosts teams found them to have overly high wear rates, thus making them largely unsuitable for the race.

Focusing on Verstappen, we see that he completed a very short stint on the medium tyres, whereas most other drivers completed significantly more laps on this compound. As the race was progressing, the Red Bull team realised that the hard tyres offered the best performance in the race, and thus used the medium tyres for a short a stint as they could, with this stint involving Verstappen driving aggressively to get the most out the tyres before pitting again.

Track position is incredibly important in F1 races, as cars in free air are able to both drive quicker and also keep their car and tyres cool. We will discuss how strategists acted both reactively and proactively throughout various stages of this race.

3.5.1 Undercuts and overcuts

When one driver is closely following another, the team has two options: pit before the car ahead, or pit after. Pitting before the car ahead, typically by a lap or two, is referred to as an "undercut".[17] The motive for doing this is that by pitting sooner, a car has fresh tyres while the car they were following is still on worn tyres. Then, when the car ahead does pit, the car behind should have used their fresh tyres to "overtake" the other car through the strategy, as the car that was ahead would come out of the pits behind the car that made the undercut. This strategy is very common in Formula 1 under the current rules, as it is often worth risking being slightly slower towards the end of the race in order to be ahead of the car in front after pit stops are made.

The opposite of this is the "overcut". This involves staying out a few laps longer than the car ahead. While this is less common in current F1, it is often used in cases where there is a risk of coming out in traffic after a pit stop, or if they would come out behind a car they were battling before the leading car made a pit stop. Often, when two drivers are battling and the car behind attempts an undercut, the car ahead may delay making a pit stop for several laps. This is because if they pit the lap after the car behind them did, they may end up behind them, leaving them in traffic and thus slowing them down. By delaying the pit stop, they open up other strategy options (possibly making one less stop during the race) and also keep themselves in clean air, thus maximising on track performance for the remainder of the stint. This strategy is less common in the current F1 rules, but in previous years when refuelling was allowed during the race, it was a much more common strategy. When two cars are battling and one makes a pit stop, they are then carrying significantly more fuel than the car that stayed out. This results in a large pace difference between the two cars, so when the second car makes a pit stop a number of laps later, they may come out of the pits ahead of the car that pitted earlier!

In Figure 3.2, we can see that Giovinazzi makes a pit stop on lap 7, changing from the medium tyres to the hard tyres. At the start of lap 7, Magnussen led Giovinazzi, Russell, and Latifi. Giovinazzi came into the pits first on lap 7, then Russell pitted on lap 8, followed by Magnussen and Latifi on lap 9. After all four drivers had made their pit stop, Giovinazzi was ahead of Russell, Magnussen and Latifi. By pitting earlier, both Giovinazzi and Russell gained positions because of the undercut!

In the first stint of the race, Bottas was in the lead ahead of Hamilton and Verstappen. Bottas then made the first pit stop on lap 13, followed by Hamilton on lap 14. Mercedes may have made these pit stops in order to prevent Verstappen from using the undercut. However, Verstappen decided to stay out for another 12 laps. By staying out longer, Verstappen used the clear air to open up enough of a lead ahead of Bottas and Hamilton so that when he did make his first stop, he actually came out ahead of them both! This successful overcut allowed Verstappen to eventually win the race.

3.5.2 Other potential factors to model

Tyre temperature

Throughout the course of a race, there are many variables that can affect the performance of a car along with the tyre wear and fuel load. One such variable is the temperature during the race, specifically the track temperature. When track temperatures are low, tyres are less able to generate temperature and thus provide less grip, resulting in slower lap times. Conversely, when track temperatures are high, tyres are more likely to overheat which also slows the cars down. There is a temperature window in which the tyres perform at their best, resulting in the quickest lap times. The trouble that teams face is that the 3 compounds at each race often have different operating windows and so when the temperature changes, a tyre that was performing well may find itself being either too hot or too cold, thus slowing down. In general, harder tyre compounds perform better when it is hot, while softer tyre compounds perform better in cooler conditions. Teams need to pay attention to significant temperature changes throughout the race, as being on the right tyre at the right time is vital to performing well in a race.

We can imagine that the effect of tyre temperature on lap time takes the form of $\exp\left(-(\tau - \tau_{opt})^2\right)$, with τ being the tyre temperature and τ_{opt} being the optimal tyre temperature. There is therefore a clear optimal peak in performance with declining performance on either side when the temperature is too high/low.

Track grip

As a race track is used more and more, the rubber from the tyres is laid down onto the track, resulting in the track being "rubbered-in". This results in the track providing more grip for the cars, allowing for faster lap times. When a track is "green", which is the opposite of a rubbered-in track, teams may opt to use softer tyre compounds as the tyres will be less able to generate temperature due to the reduced grip. As the race progresses, harder tyre compounds become increasingly more viable as the track provides more grip allowing the tyres to get up to the required temperature more easily. This effect is most apparent on street circuits, as they are very rarely used for racing meaning the track surface is significantly less grippy at the start of a weekend.

Similarly, new track surfaces tend to provide very little grip, as the bitumen and oils can seep from the tarmac causing the track to be more slippery. As the surface is used more and more, these oils and bitumen are gradually removed, resulting in the track providing more grip.[18] This was a significant issue at the 2020 Turkish Grand Prix, where a 2-week old track surface and cold conditions resulted in the track being incredibly slippery at first, but lap times managed to decrease by 6.5 seconds within just a few hours of usage.

The level of grip on a track surface can take the form of $1 - \exp(-x)$, where $x > 0$ is some measure of track usage (e.g. laps completed). Here, grip increases very quickly when the track is first being used, but over time the improvement in grip becomes negligible.

Chapter 4

F1 Strategy Competition - Preparation

4.1 Introduction

In the F1 strategy competition, run by Cédric Beaume, a number of participants will create a virtual F1 team in which they will need to allocate a budget of $8i$ on various components.[19] The teams will then be assigned drivers through a drafting process, and then the season will begin. 10 races will take place across the season, with each team needing to determine a race strategy for their two drivers. At the end of each race, finishing drivers will be assigned points based on their finishing position. These points will be totalled for each team at the end of the season to determine the winner.

The goal of this competition is clear: win the team championship.

While it is possible to allocate the budget and determine race strategies using intuition alone, we will make use of various techniques to find the budgets and strategies that give the best chance of winning the championship.

4.2 The budget problem

Each team has a budget of $8i$ that they can spend in 4 areas: Reliability, Marketing, Chassis, and Engine. A team's reliability budget affects the probability that a driver finishes a race, with an increased budget resulting in a decreasing likelihood of retirement. The marketing budget is used to assign drivers. A team with a higher marketing budget, relative to the other teams, is more likely to hire the best drivers. The chassis and engine budgets have a guaranteed influence on performance, as they contribute directly to the performance of the car at each race.

A team can invest their budget to 1 decimal place (e.g. $1.1i$ into reliability, $3.0i$ into marketing, $3.1i$ into engine, and the remaining $0.8i$ into chassis). The goal is to find the best way to allocate the budget into the given areas to give the best chance of winning the competition.

4.2.1 Reliability

The reliability budget will determine the probability of a driver retiring in any given race. A greater investment into the car's reliability will result in it having a lower probability of retirement. The probability of a car retiring in any race is defined as follows:

$$p_{\text{DNF}}(R) = \left(\frac{1 - \text{erf}(1.5R - 1.8)}{2}\right)^4, \tag{4.1}$$

where R is a team's reliability budget, and $\text{erf}(z)$ is the error function.[20]

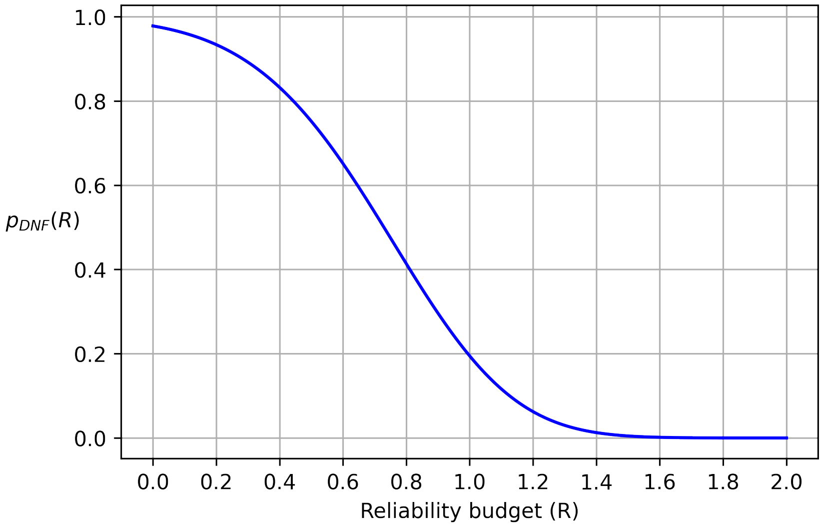

Figure 4.1: Retirement probability p_{DNF} as a function of R.

As shown in Figure 4.1, we can see that the probability of retirement quickly approaches 0 as the reliability budget increases. In each race, a driver either finishes the race or they do not. Since we have a retirement probability p_{DNF} for each team, and each team has 20 opportunities to retire (10 races with 2 cars), we can define a binomial distribution. We can therefore calculate the expected number of retirements a team will experience in a single season, as well as the cumulative probabilities for amounts of retirements for each team:

Let N_{ret} be the number of retirements for a team across a season. Then,

$$\mathbb{E}[N_{\text{ret}}|R] = 20p_{\text{DNF}}(R); \tag{4.2}$$

$$\mathcal{P}[N_{\text{ret}} \leq N_{\max}|R, N_{\max} \leq 20] = \sum_{n=0}^{N_{\max}} \binom{20}{n} p_{\text{DNF}}(R)^n (1 - p_{\text{DNF}}(R))^{20-n}. \tag{4.3}$$

Example 4.1

Suppose a team invests $R = 0.9i$ into reliability. Then, for each of their drivers in each race, the probability of retiring is $p_{\text{DNF}}(0.9) \approx 0.296$. Hence,

$$\begin{aligned} \mathbb{E}[N_{\text{ret}}|R = 0.9] &= 20p_{\text{DNF}}(0.9) \\ &\approx 6. \end{aligned}$$

Therefore, we would expect the team experience around 6 retirements across the 10 race season. Equivalently, we would expect each driver in the team to retire from around 3 races across the 10 race season. A plot of the cumulative probabilities $\mathcal{P}[N_{\text{ret}} \leq N_{\max}|r = 0.9, N_{\max} \leq 20]$ is shown in Figure 4.2.

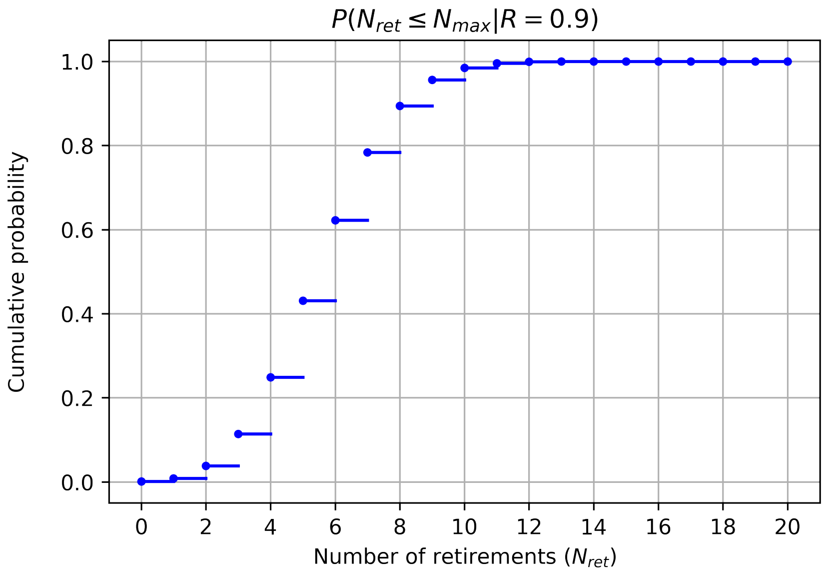

Figure 4.2: Cumulative probability plot of number of retirements with $R = 0.9$.

While it may be desirable for a team to not experience any retirements across the season, this requires a greater investment into reliability which would take away from the car's performance. A trade-off exists in the sense that a more reliable car is an ultimately slower car. Therefore, it may be a better approach to "allow" a certain number of retirements across the season in order to maximise the performance of the car. Of course, investing too little into reliability would result in too many retirements across the season. At the moment, it is difficult to determine what is the optimal value of R for a team, and we will need to consider the other aspects of the budget before determining how much to allocate towards reliability.

4.2.2 Marketing

The process in which the drivers are assigned is not deterministic, as there is a random element. Once each team has determined their marketing budget, they will be assigned a probability of hiring the next driver, with the drivers being assigned in decreasing order of performance. This probability is given by

$$p_i = \frac{m_i^4}{\sum\limits_{j=1}^{N} m_j^4}, \tag{4.4}$$

where m_i is the marketing budget of team i and N is the number of teams with at least one empty seat at that stage in the draft.

The probabilities for each team will be summed cumulatively, thus partitioning the interval $(0, 1)$ into N intervals, $M_t \subseteq (0, 1)$, with $t = 1, \ldots, N$, where N is the number of teams with at least one empty seat. A uniformly distributed number, x_{RNG}, in the interval $(0, 1)$ will be chosen. If $x_{\text{RNG}} \in M_t$, then team t will be assigned the next driver. Once a team has been assigned two drivers, they will be removed from the drafting process, and the probabilities for each team will be recalculated. This process is repeated until all drivers have been assigned.

Example 4.2

Suppose we have 3 teams who choose the following marketing budgets: $M_A = 2.2$, $M_B = 4.5$, $M_C = 0.3$. The 6 drivers that are to be assigned have the following performance ratings:

$5.0, 4.0, 3.0, 2.0, 1.0, 0.0$. At the first stage in the drafting process, where all teams are present, the teams have the following driver assignment probabilities according to (4.4)

$$p_A = \frac{2.2^4}{2.2^4 + 4.5^4 + 0.3^4} = 0.05403\ldots$$
$$p_B = \frac{4.5^4}{2.2^4 + 4.5^4 + 0.3^4} = 0.94594\ldots$$
$$p_C = \frac{0.3^4}{2.2^4 + 4.5^4 + 0.3^4} = 0.00001\ldots.$$

Clearly, Team B has the best chance of being the first and second drivers. The first random number is $x_{\text{RNG}} = 0.465\ldots$ which is in the interval $(p_A, p_A + p_B) = (0.05403\ldots, 0.99998\ldots)$. Therefore, the first driver with 5.0 performance is assigned to Team B. The next random number is $x_{\text{RNG}} = 0.091\ldots$, which is the same interval as before, meaning the second driver (4.0) is also assigned to Team B, which has now got two drivers, so Team B are removed from the drafting process. We now recalculate the driver assignment probabilities.

$$p_A = \frac{2.2^4}{2.2^4 + 0.3^4} = 0.99965\ldots$$
$$p_C = \frac{0.3^4}{2.2^4 + 0.3^4} = 0.00034\ldots.$$

Team A is now highly likely to be assigned the next 2 drivers. The next random number we draw is $x_{\text{RNG}} = 0.784\ldots$ which is the interval $(0, p_A) = (0, 0.99965\ldots)$, so the driver with performance 3.0 is assigned to Team A. The fourth random number we draw is $x_{\text{RNG}} = 0.351\ldots$ which is also in the same interval as before, so the driver with performance 2.0 is assigned to Team A. There are now only 2 drivers remaining in the draft, and Team C has 2 seats available, so both these drivers are automatically assigned to Team C.

In summary, Team A was assigned the drivers with performance 3.0 and 2.0, Team B was assigned the drivers with performance 5.0 and 4.0, and Team C was assigned the drivers with performance 1.0 and 0.0.

Given the random nature of this process, an analytical solution cannot be found. Therefore, Monte Carlo simulations will be used in order to find the optimal solution to this problem. This will involve assigning marketing budgets to 13 teams, at random, according to some distribution, and then performing the drafting process to assign drivers to the teams. The drivers that are assigned to each team will be recorded, and we will track both the average driver performance and the top driver's performance for each team, where each driver is assigned a performance value between 0 and 5, to 1 decimal place. For each budget distribution we consider, we will repeat the budget assignment and drafting process $n = 1,000,000$ times.

An important consideration when choosing the marketing budget is "return on investment". A team that spends more of their budget on marketing has less budget to spend on the chassis, engine, or reliability of their car, meaning a trade-off exists. As such, when performing these Monte Carlo simulations, we are interested in the "expected performance" of the car which is the sum of the mean top/average driver performance and the remaining budget not spent towards marketing. It should be noted that maximising return on investment does not necessarily mean a team has the best chance of winning the championship, as relative return on investment to the other teams would determine which teams gained.

Example 4.3

Suppose a team allocates $M = 2.4$ of their budget towards marketing and is assigned two drivers, A and B, with performance 4.2 and 3.8 respectively, the expected performance for these drivers would be

$$\mathbb{E}(P_A) = 4.2 + (8 - 2.4 - R) = 9.8 - R;$$
$$\mathbb{E}(P_B) = 3.8 + (8 - 2.4 - R) = 9.4 - R,$$

where R is the budget allocated towards reliability.

Clearly, a team would want the performance of their drivers to be greater than the budget they allocated towards marketing, else they could have simply used this budget in a way that guarantees additional performance.

Assumption of budgets

The first step in the Monte Carlo process is to consider how much of a team's budget may be allocated to marketing. A naïve assumption may be that each team chooses a uniform random value between 0 and 8, to 1 decimal place.

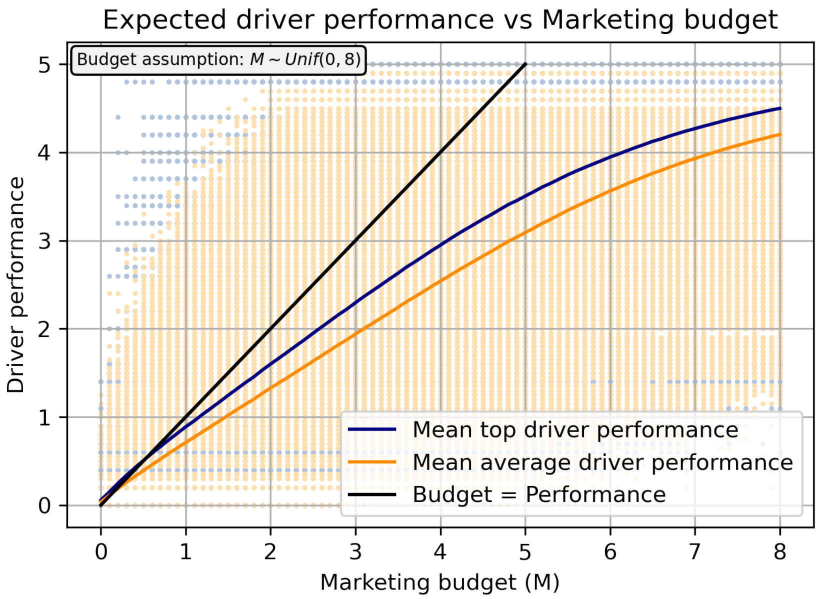

Figure 4.3: Raw results of drafting process assuming Unif(0,8) distribution.

Figures 4.3 and 4.4 shows the results of the Monte Carlo simulation for this budget assumption. We can see from the Figure 4.4 that expected performance is maximised when the marketing budget is close to 0. Therefore, should we assume this to be the true assumption, we would expect a low marketing budget to offer the best "return on investment" for a team.

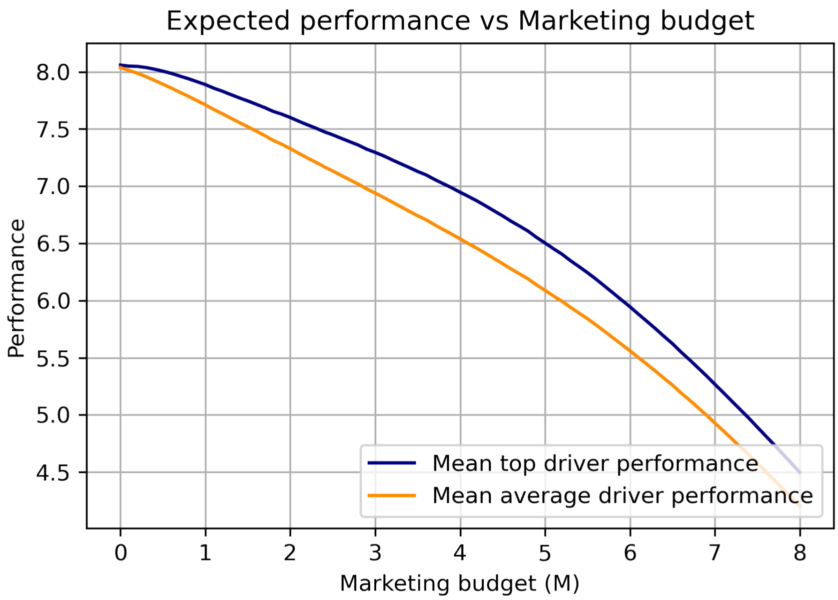

Figure 4.4: Expected performance for average and top drivers as a function of the marketing budget assigned assuming Unif(0,8) distribution.

However, it is easy to see that a uniform distribution is unlikely to be the true distribution. Firstly, it implies that teams have a 37.5% chance of choosing a marketing budget above $5i$. No team would realistically choose such a marketing budget as it would be an immediate waste of the available budget, as it is impossible to have a driver performance above 5. Secondly, we are able to study the marketing budgets that were chosen in previous seasons of this competition, as seen in Table 4.1.

Season 1	Season 2
2.6	0.2
2.5	2.8
2.3	2.8
3.2	2.5
2.6	2.7
2.6	2.7
3.0	2.6
1.6	1.8
5.0	2.0
N/A	0.0
N/A	0.0

Table 4.1: Marketing budgets for teams in previous seasons[21],[22]

While we only have 20 marketing budgets to consider, we can see that most teams seem to focus their marketing budget around 2.5, with only 3 teams having a marketing budget above 3. In previous seasons, each team had a total budget of $7i$ compared to $8i$ this season, but car reliability was fixed in previous seasons. We have already seen how a reliability budget of $r = 0.9$ seems to provide a reasonable level of reliability across a season, so the choice of marketing budget is largely unaffected by the overall increase in available budget.

With the data from previous seasons, we can find the mean budget and the variance so we can fit a normal distribution. We find that $\hat{\mu} = 2.275, \hat{\sigma}^2 = 1.282875$. However, when picking a random number according to a $\mathcal{N}(2.275, 1.282875)$ distribution, we have a chance of picking a number either less than 0 or greater than 5. In this case, we simply select either 0 or 5 respectively as the marketing budget. As before, we will pick $n = 1,000,000$ sets of 13 marketing budgets to see if a pattern emerges.

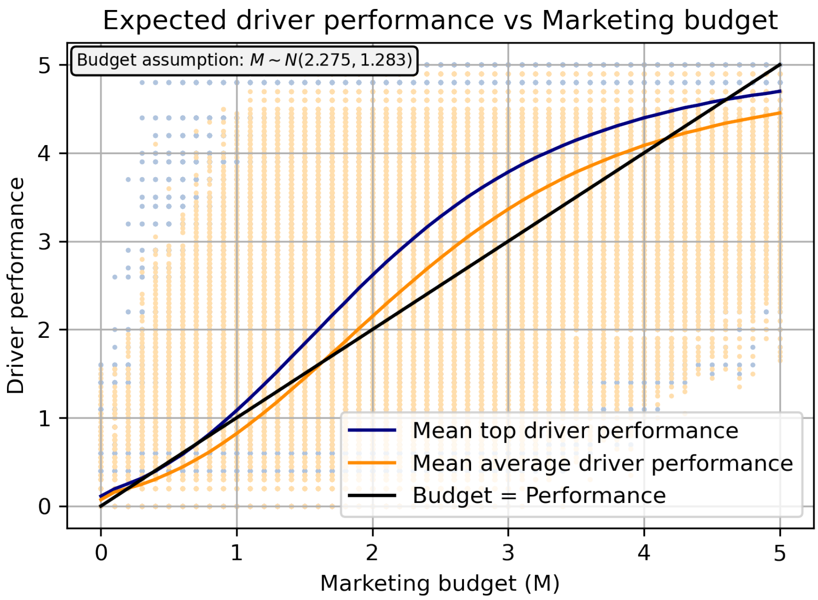

Figure 4.5: Raw results of drafting process assuming $\mathcal{N}(2.275, 1.282875)$ distribution.

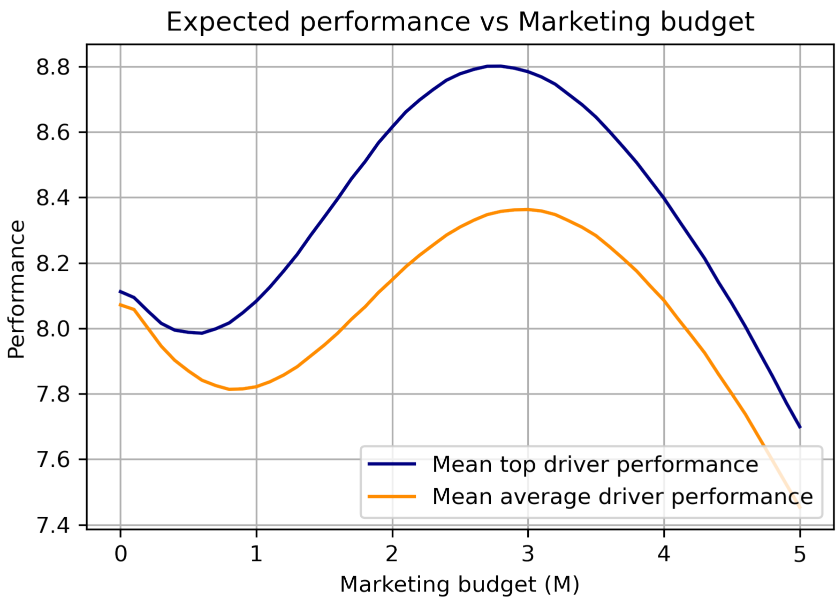

Figure 4.6: Expected performance for marketing budgets assigned assuming a $\mathcal{N}(2.275, 1.282875)$ distribution.

Figures 4.5 and 4.6 shows that there appears to be two distinct peaks in the average driver

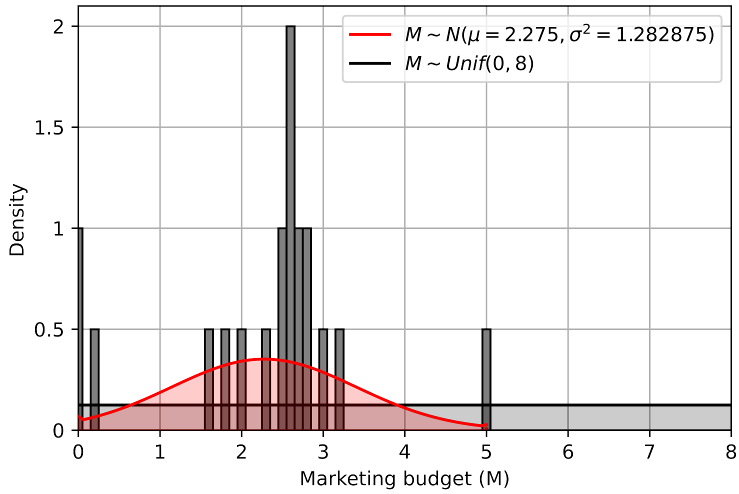

Figure 4.7: Uniform distribution and normal distribution shown on top of a histogram of the true data.

performance depending on the chosen marketing budget. One peak appears to be close to 0 while the other is around 2.8. The peak where the marketing budget is above the mean is expected, as we would expect teams that opt for an above average marketing budget to be allocated better drivers. The decline for high marketing budgets is likely due to the cap on driver performance. With our assumed distribution, teams have a ∼14% chance of choosing a budget over 3.5 which, given the method for calculating the driver assignment, means a team bidding over 3.5 is very likely to be assigned at least one of the first few drivers. As a result, investing too much into marketing will have diminishing returns, which explains the apparent reduction in the return on investment.

As for the smaller peak at the lower marketing budgets, this may be because of the drivers assigned at the end of the process. The 6 lowest driver performances available are: 0.9, 0.7, 0.6, 0.4, 0, 0. When these drivers are being assigned, it is likely that only 3-5 teams remain in the process, which increases each team's probability of being assigned the next driver. As a result, teams with a low marketing budget have a reasonable chance of being assigned drivers with a performance greater than their marketing budget, resulting in a positive return on investment. The dip between this peak and the greater peak may be because teams choosing these budgets fall into a "no man's land". These teams are very unlikely to be assigned the best available drivers, and are unlikely to benefit from occasions where few teams remain in the process, as is the case for very low marketing budgets.

Assuming a normal distribution for the marketing budgets appears to provide a more realistic distribution, although we should check to see how it actually compares the true data.

Figure 4.7 clearly shows the improvement of the normal distribution over the original uniform distribution. However, we can also see where it fails to fully capture the marketing budgets from the previous seasons. For instance, the normal distribution places too much weight on marketing budgets between 0.3 and 1.5, as well as between 3.2 and 5. It also fails to capture the peaks near 0 and 2.6.

To obtain a more accurate distribution, we need to define our own distribution that is based on the data available to us. We first split the interval [0,5] into a number of sub-intervals, each with a given probability of being selected.

$$\mathcal{P}(a \le M \le b) = \begin{cases} 0.1, & a = 0, b = 0.3 \\ 0.05, & a = 0.4, b = 1.3 \\ 0.1, & a = 1.4, b = 2 \\ 0.6, & a = 2.1, b = 3 \\ 0.1, & a = 3.3, b = 4 \\ 0.05, & a = 4.1, b = 5 \\ 0, & \text{otherwise.} \end{cases} \tag{4.5}$$

Then, each possible value within each subinterval has an equal chance of being chosen, giving the probability distribution

$$\mathcal{P}(M = m) = \mathbf{I}_{m \in [a,b]} \frac{\mathcal{P}(a \le M \le b)}{10(b - a + 0.1)} \qquad \forall m \in [0, 5], \tag{4.6}$$

where m is given to 1 decimal place, and

$$\mathbf{I}_{m \in [a,b]} = \begin{cases} 1, & \text{if } m \in [a, b] \\ 0, & \text{otherwise.} \end{cases}$$

For instance, the probability of randomly selecting a marketing budget of $M = 2.6$ is

$$\mathcal{P}(M = 2.6) = \frac{\mathcal{P}(2.1 \le M \le 3)}{10(3 - 2.1 + 0.1)} = 0.06,$$

where we have simply divided the total probability in the interval [2.1,3] by the number of possible values within the interval (i.e. 10).

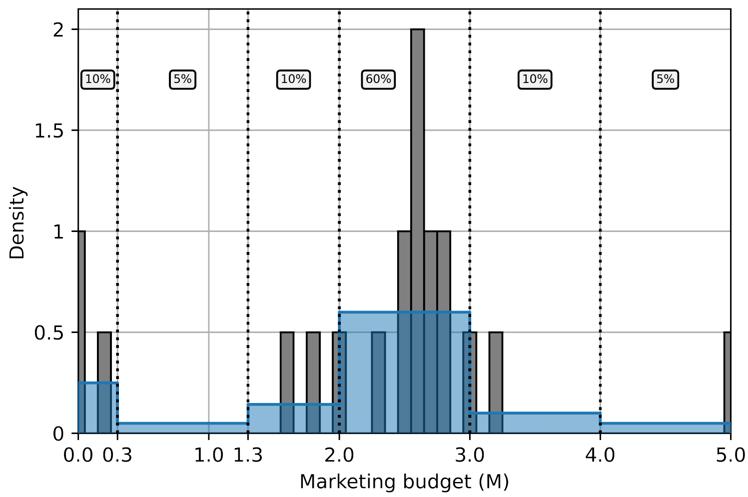

Figure 4.8: Custom distribution shown on top of a histogram of the true data. Probability of choosing a budget within each shaded area is indicated.

Figure 4.8 shows how this distribution fits to the true data. We can see how 60% of marketing budgets will be between 2.1 and 3, and that the density reduces as the marketing budget increases/decreases. There is also a slight peak at very low marketing budgets which reflects data from past seasons. Of course, this is again unlikely to be the true distribution (especially since this data is from a slightly different rule-set), but the usage of this data gives us some basis for our further study. In reality, the choice of marketing budget is more closely linked to game theory, as we have a number of rational individuals taking part in a non-cooperative game. Each player is aiming to predict what the others are doing, and thus predict what choice they can make to give themselves the best outcome. This is beyond the scope of this project, but it justifies the custom distribution we have defined.

For instance, a team may predict that one other team will submit a low marketing bid, say 0.1. This then makes submitting a low marketing bid (such as 0.2) more worthwhile, as they are more likely to see a positive return on investment, as explained previously. However, if every team thinks in this way, every team could end up submitting a low marketing bid. A team could then predict this would happen and realise they could get away with submitting a marketing bid of around 1.2 and still have a very good chance of being assigned the best drivers. This process of teams trying to outsmart others can continue indefinitely.

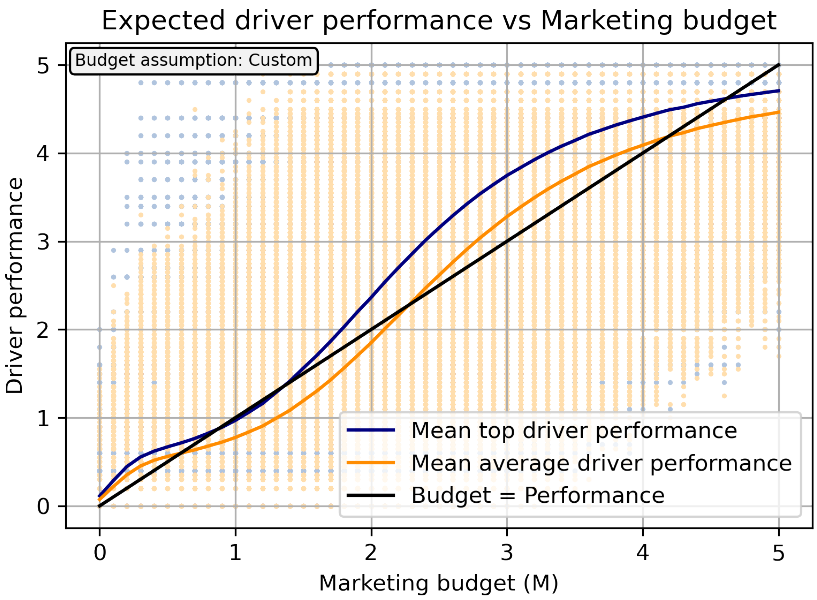

Figure 4.9: Raw results of drafting process assuming the custom distribution described in (4.5) and (4.6).

Figures 4.9 and 4.10 shows that, assuming the custom distribution described in (4.5) and (4.6), we see a pattern similar to the results for the normal distribution. The two distinct peaks appear to be more accentuated when assuming this distribution, and we can see that a team would expect the best return on investment if they were to choose $M \in [2.7, 3.2]$ as their marketing budget. The peak close to 0 also suggests there may be a valid reason to choose a very low marketing budget.

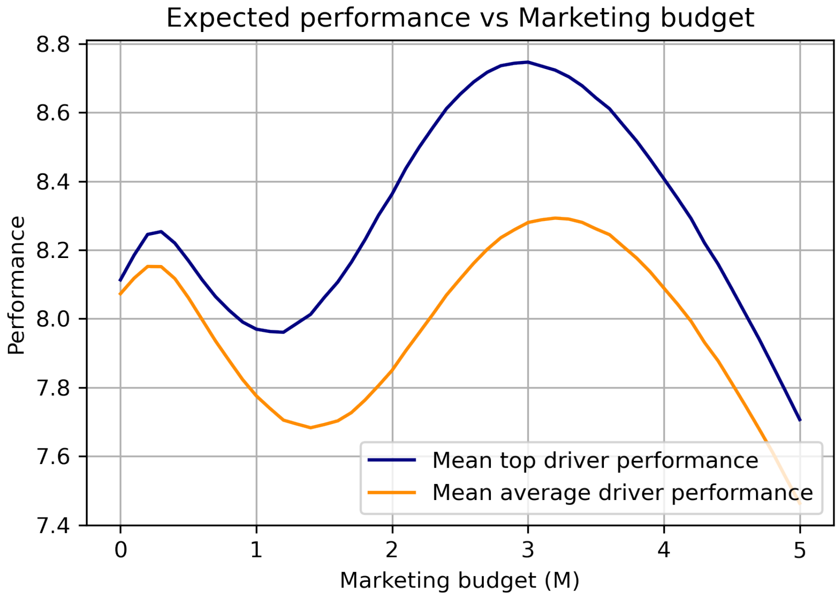

Figure 4.10: Expected performance for marketing budgets assigned assuming the custom distribution described in (4.5) and (4.6).

From this points onwards, we will be using the custom distribution described in (4.5) and (4.6) in future calculations/simulations that involve assigning the marketing budgets for teams.

4.2.3 Chassis and engine

The remaining 2 aspects of the budget allocation are closely linked, so we will discuss them together. In this competition, each track has a different "balance" that influences the performance of a car and driver around that track. For instance, a track may favour cars with a strong chassis, or it may be completely balanced thus valuing each aspect of the car's performance equally. Each track is assigned 3 track balance coefficients c_d, c_c, c_e which relate to the Driver, Chassis, and Engine respectively. The track performance of a driver from a given team is given by the following formula

$$P_{\text{driver}} = \frac{3(c_d D + c_c C + c_e E)}{c_d + c_c + c_e}, \tag{4.7}$$

where D is that driver's performance rating (out of 5), C is the team's chassis budget, and E is the team's engine budget.

Track	c_d	c_c	c_e
Spain	1	1	1
France	1	1	1
Monaco	2	1	1
Belgium	1	2	2
Italy	1	1	2
Hungary	1	2	1
USA	1	1	2
Japan	2	2	1
Canada	1	1	1
Australia	1	1	1

Table 4.2: Track balance coefficients for each track.

Table 4.2 shows the track balance coefficients c_d, c_c, c_e for each of the 10 tracks in the competition. For instance, Spain is a balanced track, Monaco favours the driver performance, and Belgium favours the chassis and engine performance.

Example 4.4

Suppose a team's drivers (A and B) have a performance rating of 4.2 and 2.5 respectively, and that they invested $C = 2.4$ and $E = 3.0$ into their chassis and engine. The track performance of each driver, according to (4.7), at Hungary is therefore

$$P_A = \frac{3(4.2 + 2 \cdot 2.4 + 3.0)}{1 + 2 + 1} = 9;$$
$$P_B = \frac{3(2.5 + 2 \cdot 2.4 + 3.0)}{1 + 2 + 1} = 7.725.$$

By comparison, the track performance of these drivers at Monaco is

$$P_A = \frac{3(2 \cdot 4.2 + 2.4 + 3.0)}{2 + 1 + 1} = 10.35;$$
$$P_B = \frac{3(2 \cdot 2.5 + 2.4 + 3.0)}{2 + 1 + 1} = 7.8.$$

Driver A has a track performance that is 1.275 greater than driver B's at Hungary, but at Monaco, which favours driver performance, this track performance difference increases to 2.55.

When allocating their budget, a team would want to know how split their budget between the chassis and engine which is influenced directly by the track balance coefficients in Table 4.2. To determine whether it is worth investing more into either the engine or chassis, we will consider the weighting of each component at each track.

Track	c_d weight	c_c weight	c_e weight
Spain	1/3	1/3	1/3
France	1/3	1/3	1/3
Monaco	1/2	1/4	1/4
Belgium	1/5	2/5	2/5
Italy	1/4	1/4	1/2
Hungary	1/4	1/2	1/4
USA	1/4	1/4	1/2
Japan	2/5	2/5	1/5
Canada	1/3	1/3	1/3
Australia	1/3	1/3	1/3
Average	0.3183...	0.3383...	0.343...

Table 4.3: Weighted track balance coefficients for each track.

By calculating the average weight for each of the performance components, we can see from Table 4.3 that the engine coefficient is slightly larger than the other 2. This may suggest to a team that investing more of their budget into the engine would result in the best performance across the season, but maybe only marginally so.

The obvious question is how much more of their budget should a team invest into the engine compared to the chassis. For instance, if a team invests 1.1 into reliability and 2.6 into marketing, how should they split the remaining 4.3 between the chassis and the engine? Should they put it all into the engine, or should it be closer to a 50:50 split? While it is difficult to predict exactly what is better, we can at least discuss the advantages of doing one or the other. In the next section, we will use Monte Carlo simulations to see what split between engine and chassis yields the best results.

If a team were to invest heavily into their engine, thus investing very little into their chassis, they would likely see very strong performance on any track that gives extra weight to the engine, but they would lose out on tracks that place extra weight on the chassis. It is important to note that on balanced tracks, how a team splits the budget between the engine and chassis will have no effect on the performance of the car. Across the season, there are two tracks that see the engine have a higher weight than the chassis (Italy and USA), while there are two tracks that see the chassis have a higher weight than the engine (Hungary and Japan). Conversely, if a team were to split their budget between the chassis and engine more evenly, they would have more consistent performances across the season. If they were to do a 60:40 split in favour of the engine (or some similar split), they would not gain much at engine-favoured tracks, but they would also not lose out too much at chassis-favoured tracks.

Whether it is worth prioritising consistent results over more variable results is dependent on the points system.

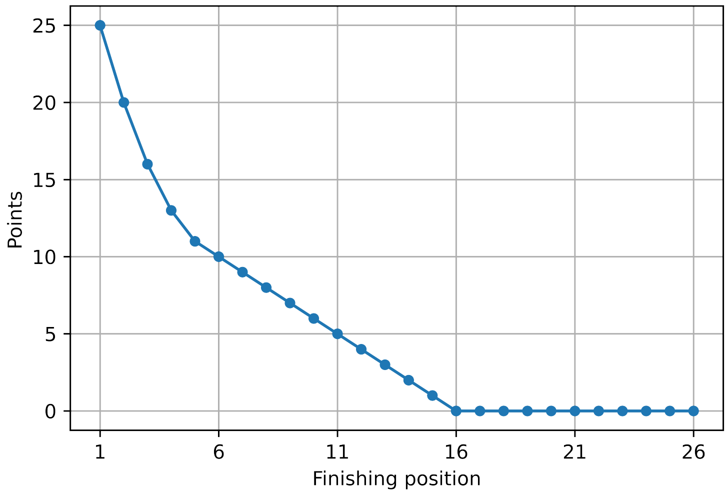

Figure 4.11: Points distribution for the competition with 26 drivers.

Figure 4.11 shows that the points distribution in this competition is weighted towards the better finishing positions. For instance, the points difference between finishing in 1$^{\text{st}}$ and 2$^{\text{nd}}$ is 5 points, whereas between 11$^{\text{th}}$ and 12$^{\text{th}}$, it is only 1 point. As a result, one could argue that it is worthwhile for a team to allocate their budgets in a way that gives them a chance of getting closer to the better positions in some, even if it sacrifices positions in other races. As mentioned previously, a team that heavily focuses on their engine would likely have a better chance of getting the best finishing positions in some races, but would lose out in other races. It is difficult to assess whether this is a worthwhile decision at this stage.

Another thing for a team to consider is how other teams may split their budget between the chassis and the engine. This is closely linked to game theory, much like the marketing budget. While it may be optimal to invest more into the engine budget, if every team does so, the benefit from doing so is reduced. It may even be worthwhile for a team to invest into the chassis instead, as they would have a very good chance of getting a strong result at the two tracks that emphasise the chassis performance above the engine. Again, it is difficult to determine what a team should do in the event of all other teams allocating their budgets in the supposedly optimal way.

4.3 Solving the budget problem

Having discussed potential effects of allocating the budget between the 4 components in various ways, we will now use Monte Carlo simulations to find the optimal budget allocation. This will involve a) generating random team budgets according to predetermined distributions; b) assigning drivers to the teams using the drafting process; c) simulating each race in a season to find the championship winner.

We will create a population of 20,000 team budgets and, at each stage, randomly select 13 of these budgets to compete in a championship. At the end of each championship, we will record which team won the championship. For each set of 13 teams, we will repeat the championship 10 times to help account for anomalistic results. We will repeat this until each team has taken part in around 10,000 seasons. At this point, we will sort the team budgets by the proportion of championships they have won to see if there is any similarity between teams that won the greatest proportion of championships.

Allocation of budgets

When randomly choosing a budget for each team, it is done in the following steps:

1) Reliability budget: R.

   We choose R according to the following distribution

$$R \sim \mathcal{N}(1, 0.25^2).$$

   We have already found that a team choosing a reliability budget of $R = 0.9$ would expect to see around 6 retirements across a season, and this reduces to around 4 for a reliability budget of $R = 1$. We would therefore expect teams to choose a reliability budget around 1 with some variance. By choosing a random number according to this distribution and then rounding to 1 decimal place, we give each team their reliability budget, although we restrict this budget to be between 0.5 and 1.5 by rounding up or down where necessary.

2) Marketing budget: M.

   We choose M according to the distributions described by (4.5) and (4.6).

3) Chassis and engine budgets: C and E.

   The remaining budget of $8 - M - R$ is split uniformly between C and E. While we discussed how prioritising E may provide the best results, we do not yet know if this is actually the case.

Simulating races and seasons

Once we have assigned drivers to each team, we can start simulating the 10 races in the season. This is done by calculating the track performance of each driver at each track. In Example 4.3, we saw how the track performance of drivers is calculated given the track balance coefficients, driver, engine, and chassis rating. We do this for each driver in each team, and then sort the drivers according to their track performance. The driver with the highest track performance rating is declared as the winner, with the second highest finishing second, and so on. The points are given to each driver according to the points distribution shown in Figure 4.8. These points are then summed across the season of 10 races, with the driver and team with the most points being declared the winner! We are interested in the team that is declared the winner as, after all, this is a team-based competition. Each season is then repeated 10 times, with each team being assigned (hopefully) different drivers each time.

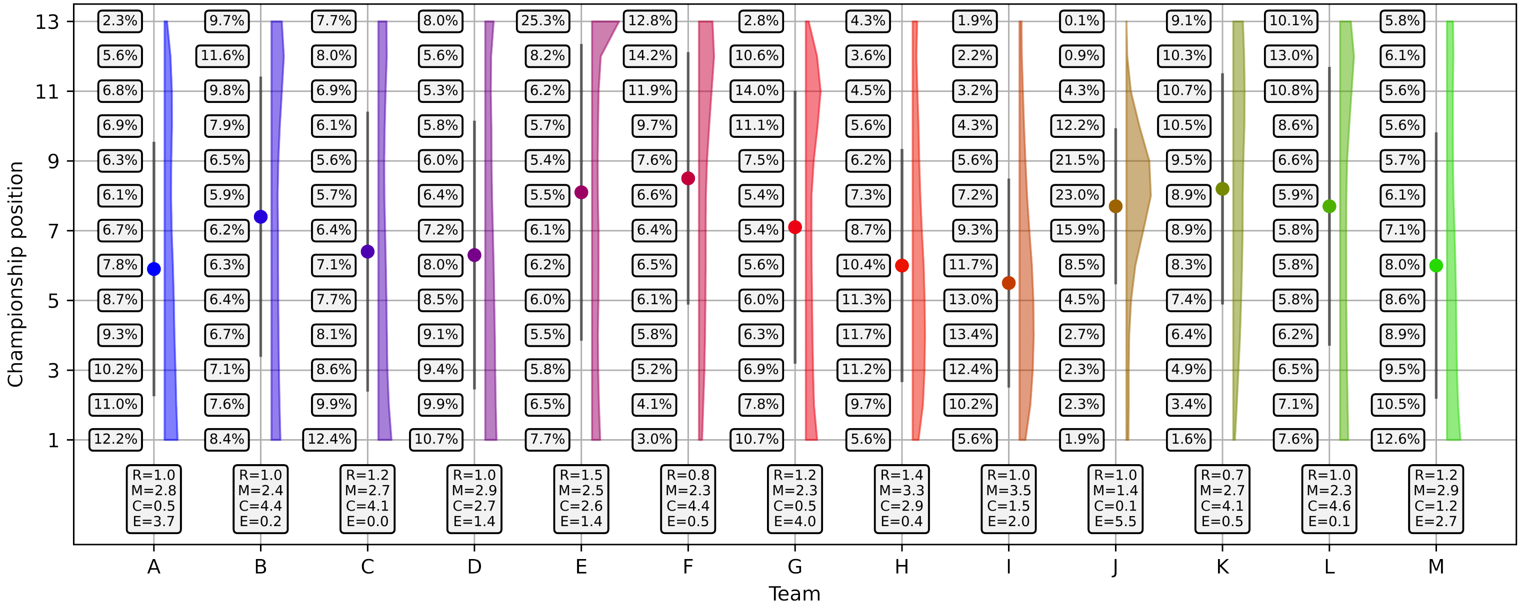

Figure 4.12: Summary of 100,000 simulated seasons with 13 fixed teams. Each team's finishing position distribution is shown as a percentage and as a plot. Team budgets are shown at the bottom.

Example 4.5

Figure 4.12 shows the summary of 100,000 seasons with 13 fixed teams. Here, we have performed the driver allocation process at the start of each season, and recorded each team's finishing position at the end of the season. Out of these 13 teams, we would determine that teams A, C, and M are the "best", as they have the highest proportion of championship victories. Teams F, J, and K had the fewest number of championship victories. One thing of interest to us is the shape of the distribution curves for each team. Considering the marketing budgets for each team, we see that teams with a high marketing budget tend to have relatively flat curves. Conversely, teams with lower marketing budgets tend to have a noticeable bulge. Team J is one such team, although its marketing budget is not that much lower than the other teams. As for the reliability budgets, team E has a very high reliability budget and sees themselves finish in last 25.3% of the time! The point marked between the percentages and distribution plot shows the teams average finishing position in the simulated seasons. Interestingly, team I has the best average finishing position but they only win the championship on in 5.6% of the seasons.

While this example is not exhaustive of every possible effect on a team's finishing position, it demonstrates some of our findings from the previous section.

Results of simulations

Having sorted the 20,000 teams in order of the proportion of championships they won, we can look for any similarities between teams near the top.

Rank	R	M	C	E	Champs. won	Seasons entered	**Win rate**
1	1.1	2.9	0.2	3.9	1819	10380	**0.175**
2	1.1	2.8	0.0	4.1	1844	10580	**0.174**
3	1.1	2.8	0.1	4.0	1734	10030	**0.173**
4	1.2	2.7	0.0	4.1	1804	10460	**0.172**
5	1.1	2.7	0.2	4.0	1854	10760	**0.172**
6	1.2	2.9	0.1	3.8	1803	10570	**0.171**
7	1.1	2.8	0.0	4.1	1765	10360	**0.170**
8	1.1	2.8	0.1	4.0	1788	10500	**0.170**
9	1.1	2.9	0.1	3.9	1753	10330	**0.170**
10	1.1	2.7	0.0	4.2	1816	10730	**0.169**
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Table 4.4: Top 10 teams according to championship win rate after Monte Carlo simulations.

Table 4.4 shows the top 10 teams from the Monte Carlo simulations, ranked according to their championship win rate. We can see a pattern in how the budgets have been allocated among the best teams: reliability around 1.1; marketing around 2.8; chassis around 0.1; and engine around 4.0. These results are in agreement with our previous discussions, although the reliability budget may be a little larger than we expected. Our study of the marketing budgets with the custom distribution suggested that a marketing budget in the interval [2.7,3.2] appears to result in the greatest return on investment for a team, which may give them an advantage over other teams in the competition, thus giving them a better chance of winning overall. We can also see that teams who invest heavily in the engine instead of the chassis are more likely to win.

To further our analysis of these results, we will consider each aspect of the budget individually.

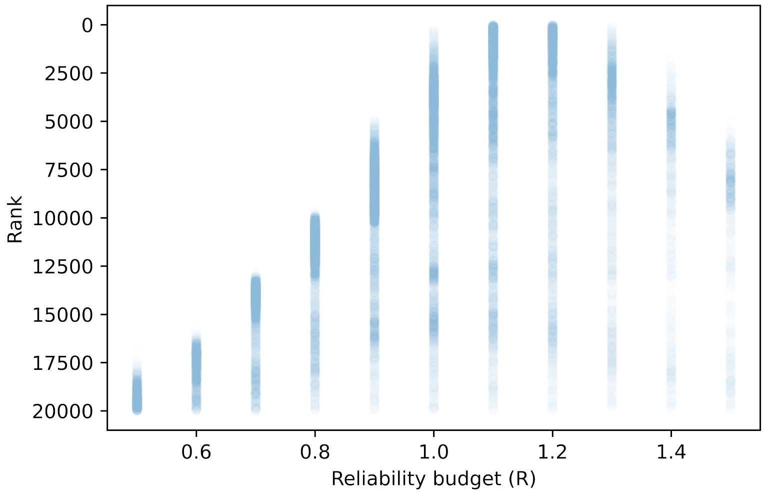

Figure 4.13: Ranking of teams plotted against their reliability budgets. Ranking closer to 0 is better. Points are slightly transparent, so darker areas indicate more densely packed points.

Figure 4.13 compares the reliability budgets assigned to teams and their finishing positions. By looking at the darker areas for each R, we can see where teams finished when using the given reliability budget. We can see that for $R \in [1, 1.3]$, these teams are closely packed towards the top of the plot, which indicates better performance. Conversely, teams with $R < 0.9$ tended to finish towards the back, as they have a worse ranking. The best performing reliability budgets appear to be $R = 1.1, 1.2$. $R = 1.0$ may also be valid, albeit a little more risky. As discussed previously, it may be worthwhile to risk a slightly worse reliability record in order to have a slightly better performing car.

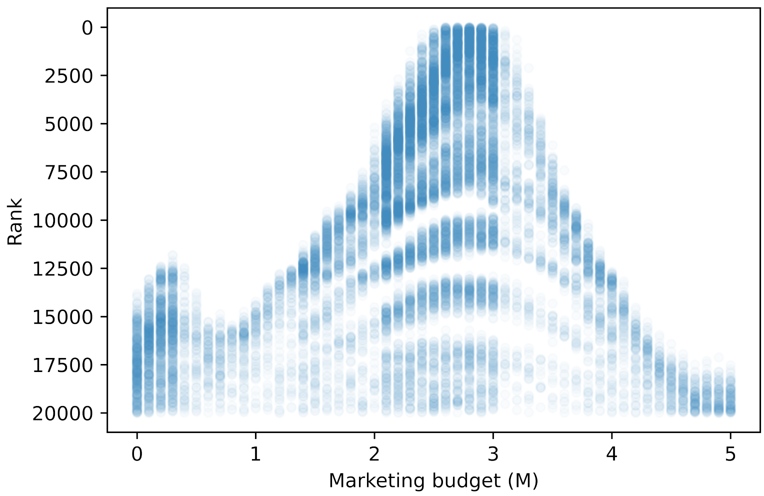

Figure 4.14: Ranking of teams plotted against their marketing budgets.

Figure 4.14 shows the ranking of teams according to their marketing budget. As before, darker colours towards the top of the plot suggests better performance. We can see that the best performing teams have $M \in [2.6, 3]$ which is what we expected following our previous analysis.

We can also see how teams that how a marketing budget close to 0 also perform better than those with $M \in [0.4, 1.3]$, but we can see that they generally have a low chance of winning championships. Also, teams who invest very highly in marketing perform very poorly, which is likely a result of them having little budget remaining for their chassis and engine. There are also some interesting gaps in the middle of the plot, which can be explained by the high density of teams with low marketing budgets that are ranked in the lower 50% of all teams.

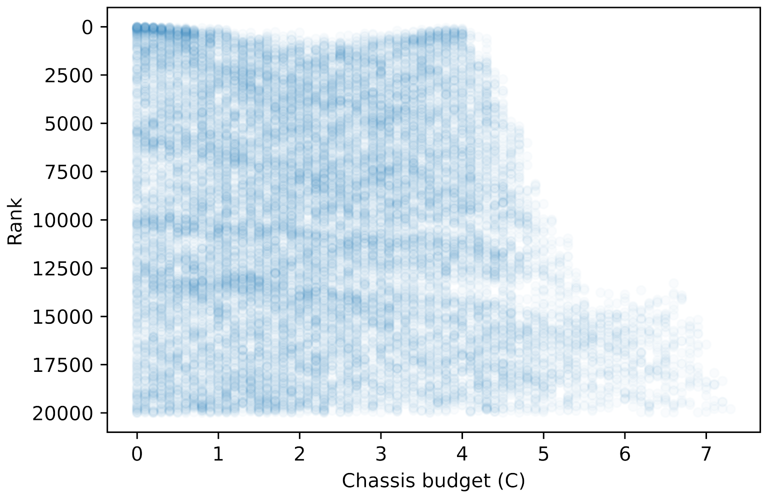

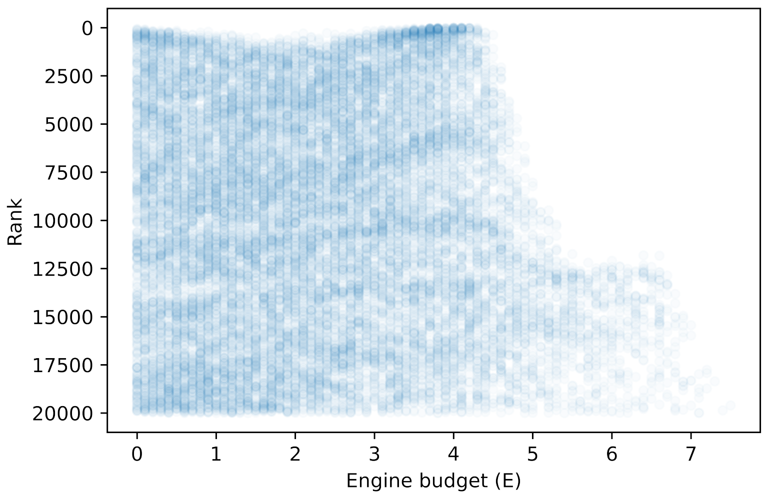

Figure 4.15: Ranking of teams plotted against their chassis (top) and engine (bottom) budgets.

Figure 4.15 shows the ranking of teams according to their chassis budget and engine budget. While there appear to be fewer clusters of teams, we can see two small clusters at the top of the plot. Considering the chassis plot, the first cluster is in the interval $C \in [0, 0.7]$, with the other in the interval $C \in [3.4, 4]$. The first interval is cases where the teams decide to prioritise their engine budget over their chassis budget (the inverse can be seen on the engine plot). We can also see that the highest ranking occurs when the engine is prioritised, although the difference is only marginal. This tells us that teams with low chassis/high engine budgets are more likely to win championships. These plots show noticeably less clustering in the data, which is indicative of a greater variance in the results, suggesting that the choice of chassis/engine budget is less influential on the win rate than the choice of marketing and reliability budgets.

It is also important to note that there is a slight dip in both plots for both $C, E \in [1.2, 2.8]$ relative to the 2 apparent peaks on each plot. This tells us opting for a more even split between the engine and chassis reduces the chance of winning the championship.

Example 4.6

Returning to the previous example, we have changed the budget of Team A to see the effects on the overall standings. Figure 4.16 shows the same summary plot as before, although Team A has had its marketing budget decreased to $M = 0.2$ and its engine budget increased to $E = 6.3$. We can clearly see the bulge on the density plot, which suggests that team A now finishes more consistently. We can also see from the marked point that team A has the best average finishing position of all teams. However, it has the lowest number of championship wins! Clearly, reducing the marketing budget has resulted in this team being more consistent, but ultimately less successful than before. Considering the other teams, team C is now the most successful while team J has performed slightly worse. Interestingly, team C has the best chance of winning even though they favoured the chassis over the engine.

We have learnt that a team with a low marketing budget may have a smaller chance of winning championships, but may instead finish more consistently around a certain position, dependent on the other teams.

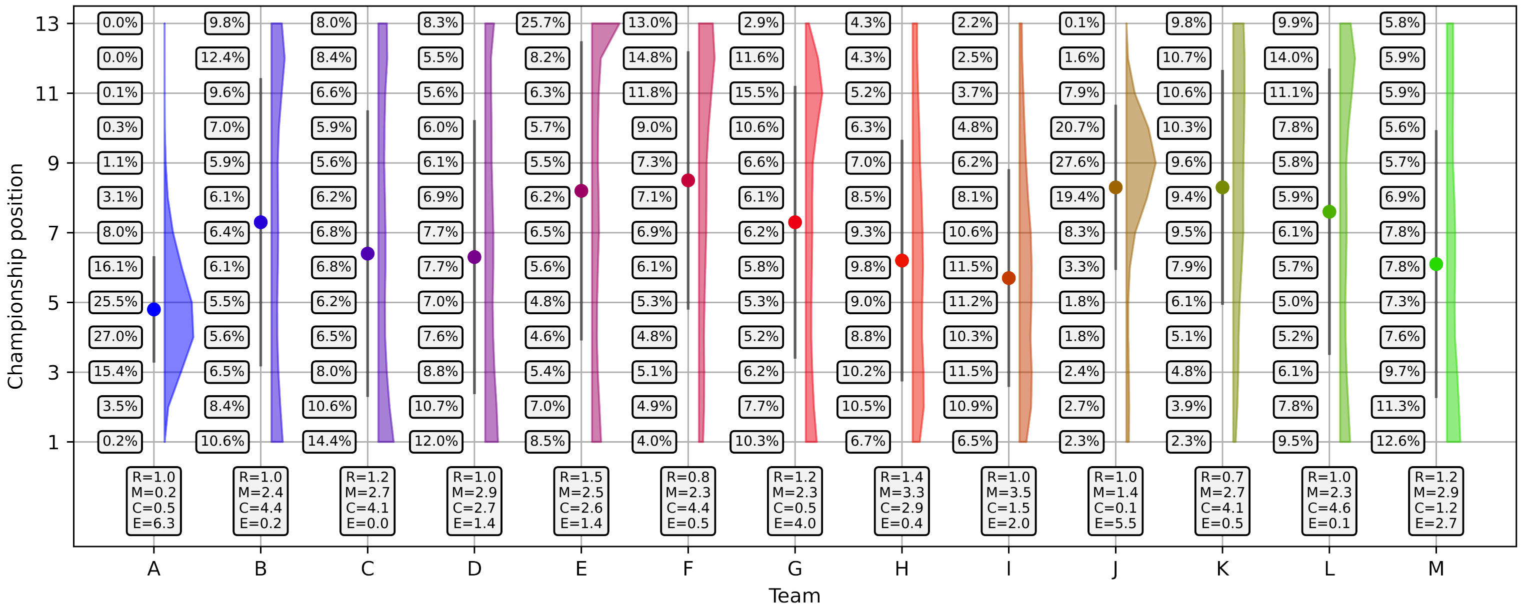

Figure 4.16: Summary of 100,000 simulated seasons with 13 fixed teams. Team B-M are the same as in Example 4.5, while Team A has had its marketing budget reduced and its engine budget increased.

4.4 Final budget

Having considered the results from our Monte Carlo simulations, the budget we choose will be based around the following observations:

- **Reliability**

  We have seen that choosing $R \approx 1.1$ is likely to give the best chance at winning a championship, although there was some justification for choosing $R = 1$. Choosing $R = 1$ would either result in a slightly improved chance of acquiring the best drivers, or it would leave extra budget available for the chassis or engine. This does, however, come at the expense of a greater chance of retiring from more races throughout the season. Given that the extra risk is relatively small, we will be choosing $\underline{R = 1}$ to give a chance at slightly improved car performance.

- **Marketing**

  While there is valid reason to opt for a low marketing budget, we have seen that to maximise the chance of winning the teams championship, a marketing budget of $M \in [2.7, 3]$ is suggested. We will choose $\underline{M = 2.8}$ for our marketing budget, although choosing any of the others in the previous interval would likely provide similar results. We could say that the $0.1i$ saved on the reliability budget has gone into the marketing budget, which is arguably a slightly high-risk strategy. By doing this, we give ourselves a better chance at being allocated better drivers, thus taking them away from other teams, which will in turn give us a better chance at winning the championship.

- **Chassis and Engine**

  With the remaining $4.2i$ available to us, we need to decide how to split this between the chassis and the engine. Our analysis of the Monte Carlo simulations tells us that investing more into the engine gives the best chance of winning the championship. If we were to opt for a more even split, we would expect more consistent results across the season, as discussed previously. Given that we have already taken some risks with our budget, particularly with the reliability budget, we will choose $\underline{C = 0.5, E = 3.7}$.

Chapter 5

F1 Strategy Competition - Participation and review

5.1 Race strategies

Lap time model

In the F1 Strategy Competition[19], the lap time model being used is very similar to our simple model (2.1). There is an additional random error term which is simply added to the lap time, as well as a driver performance term.

$$t = t_b + t_{\text{perf}} + p_t + p_f + R_{\text{lap}}. \tag{5.1}$$

The separate components, all measured in seconds, are defined as follows:

t_b: The base lap time around the given track
t_{perf}: The combined effect of the performance of the car, driver, and car setup
p_t: Time penalty due to tyre wear
p_f: Time penalty due to fuel level
R_{lap}: Random component

The tyre wear penalty, p_t, can be one of three distinct functions corresponding to the tyre compound that a driver is using in a stint (either soft, medium, or hard). These functions are unique to each track and thus need to be considered on a track by track basis. In general, soft tyres are quickest at the start of the stint but wear out quickly, hard tyres are slower at the start of a stint but are much more durable, while medium tyres find a middle ground between the two.

The fuel penalty, p_f, has the same form as in our simple model (2.2), where each lap of fuel onboard the car adds a penalty of F seconds to the lap time.

The driver performance t_{perf} term is given by

$$t_{\text{perf}} = -0.15(P_{\text{setup}} + P_{\text{driver}}), \tag{5.2}$$

where P_{driver} is the track performance of a driver given in (4.7) and P_{setup} is the setup rating for each driver. Each team is given a uniform random number in the interval [-0.5,0.5] which is the team setup rating. Then each driver also drivers a uniform random number in the same interval [-0.5,0.5] giving the driver setup rating. These ratings are combined for each driver to give P_{setup} which can take values between -1 and 1, with each driver from a team being no more than 1 setup rating apart. A team would want to have the highest combined value for P_{setup} and P_{driver}, as they are multiplied by the negative constant -0.15.

To emulate driver inconsistencies in real life races, the random term R_{lap} is added to the lap time. This is given by

$$R_{\text{lap}} = 4(x_{\text{RNG}} - 0.4)^3 + 0.3x_{\text{RNG}} - 0.4, \tag{5.3}$$

where $x_{\text{RNG}} \sim \text{Unif}(0, 1)$.

When making a pit stop, the time taken to complete the pit stop is a modification of (2.7) given by

$$t_p = t_s + \frac{1}{2}(l_i - l_{i-1}) + R_{\text{pit}}, \tag{5.4}$$

where t_s is some constant that is interpreted as the time taken to drive through the pit lane without changing tyres or adding fuel, R_{pit} is a random variable given by $R_{\text{pit}} = 2x_{\text{RNG}}$ where x_{RNG} is uniform random number between 0 and 1, and it takes 0.5s to add a single lap of fuel to the car.

Across a stint, both t_b and t_{perf} are constant for a driver, while R_{lap} is a random variable that is the same for all drivers. As a result, when it comes devising our strategies for each of the races, both p_f and p_t will be of significant importance. Unfortunately, given that p_t rarely takes the form of (2.2), that is, the tyre wear law is rarely described using a constant T plus some additional constant to differentiate between the tyre compounds as in (3.2), so we need to use alternative methods for finding the optimal strategies around a given race track.

For each track, we will create a library of stint times for each tyre compound and of each length. Then, to find the total race time for a given strategy, we will simply combine the separate stint times plus the time taken to make the pit stops in between each stints. We will then focus on the strategies that have the quickest race time, and make adaptations where necessary. One of the rules in this competition states that any team that has both drivers pitting on the same lap will see the second driver receive a time penalty of t_s seconds at the end of the race. As a result, we will need to choose strategies that do not clash, so suboptimal strategies will need to be used by at least one of drivers.

Philosophy of strategies

In section 3.5, we discussed what factors may govern the strategies a Formula 1 team will use in a race, outside of the tyre wear and fuel load penalties. Generally for this competition, the overcut is the preferable choice as refuelling is allowed during the races. If the optimal strategy involves making the first pit stop on lap 24, it may be worth extending the first to lap 25 or 26, or even beyond, to take advantage of a potential overcut. Unfortunately, we are unable to be reactive to other team's strategies in this competition, as the strategies are fixed before the race. Therefore, we have to "predict" the strategies other teams may use during the race.

Each track is given an "overtake delta" which essentially determines how easy/difficult it is to overtake at a given track, where a higher value suggests it is more difficult to overtake. For instance, Monaco has a 1 second overtake delta, while Italy's is 0.7s. At tracks where the overtake delta is high, securing track position in the short term is often more important than having a quicker potential race time. As such, at these tracks we generally go for more conservative strategies that either make one less pit stop, or attempt to make use of an overcut. There is a strong argument for qualifying on softer tyre compounds than what would be optimal, as this will likely lead to a better grid position, helping to limit the amount of traffic a car would experience in the first stint.

5.2 Season report

In this section, we will discuss the strategies we use at each of the races, along with some extended discussion of some of the races. We will also cover the chances of our team being successful when compared to the other teams in the competition both before and after the drivers have been assigned to each team.

Comparison of budgets

As discussed in the previous chapter, we have chosen the following budget for this competition: $R = 1.0, M = 2.8, C = 0.5, E = 3.7$. To see our expected finishing position before and after the driver drafting process has been completed, we will simulate $n = 1,000,000$ seasons with each team's budget given in Table 5.1, where Team A is our team.

Team	R	M	C	E	Driver 1	Driver 2
Team A	1.0	2.8	0.5	3.7	4.4	3.7
Team B	1.1	2.0	1.9	3.0	5.0	3.5
Team C	1.2	0.3	0.3	6.2	1.4	1.1
Team D	0.9	0.2	0.0	6.9	0.9	0.7
Team E	2.2	0.0	2.3	2.3	0.0	0.0
Team F	1.3	0.1	0.0	6.6	0.6	0.4
Team G	1.4	0.8	0.0	5.8	2.2	2.0
Team H	0.8	1.1	3.0	3.1	2.9	2.7
Team I	1.3	0.5	3.0	3.0	2.6	1.8
Team J	1.3	0.4	0.0	6.3	1.6	1.5
Team K	1.2	2.8	2.0	2.0	4.0	3.9
Team L	1.3	1.5	2.7	2.5	3.4	3.2
Team M	1.2	2.4	3.8	0.6	4.8	4.2

Table 5.1: Teams and their assigned drivers in the F1 Strategy Competition.[23]

Figure 5.1 tells us that according to our simulations, our team has a high variability in its predicted finishing positions, with an 8.5% chance of winning the championship. This is the 6$^{\text{th}}$ best of all teams, while Team M has the best chance with a 26.7% win rate. However, according to our simulations, our team has a 38.2% chance of finishing 7$^{\text{th}}$ or lower, which demonstrates the extreme variability in our predicted results. The most frequent finishing position for our team in our simulations was 3$^{\text{rd}}$, which occurred 12.4% of the time. Given this variability, it is very difficult to predict in what position we will finish.

Table 5.1 also shows the drivers that were assigned to each team by the drafting process described in the previous chapter. Our budget of $M = 2.8$ into marketing yielded a positive return on investment for both drivers, who have a performance rating of 4.4 and 3.8 respectively. This is one of the strongest driver pairings on the grid and it exceeds our expected driver performance ratings, according to Figure 4.10, although this has come at a cost to our chassis and engine budgets. We can now see how our team is likely to perform with the drivers we have been assigned.

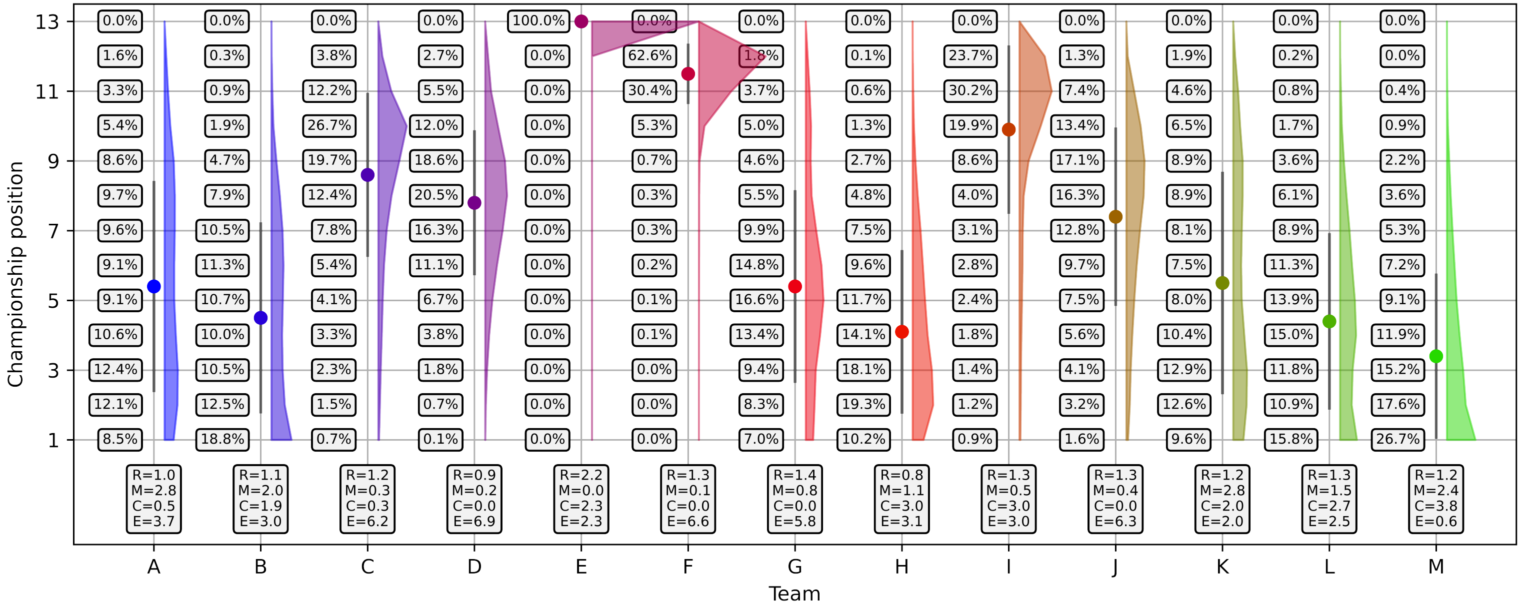

Figure 5.1: Summary of 1,000,000 simulated seasons with budgets given by Table 5.1, before the drivers have been allocated. The budget chosen in Chapter 4 has been assigned to Team A.

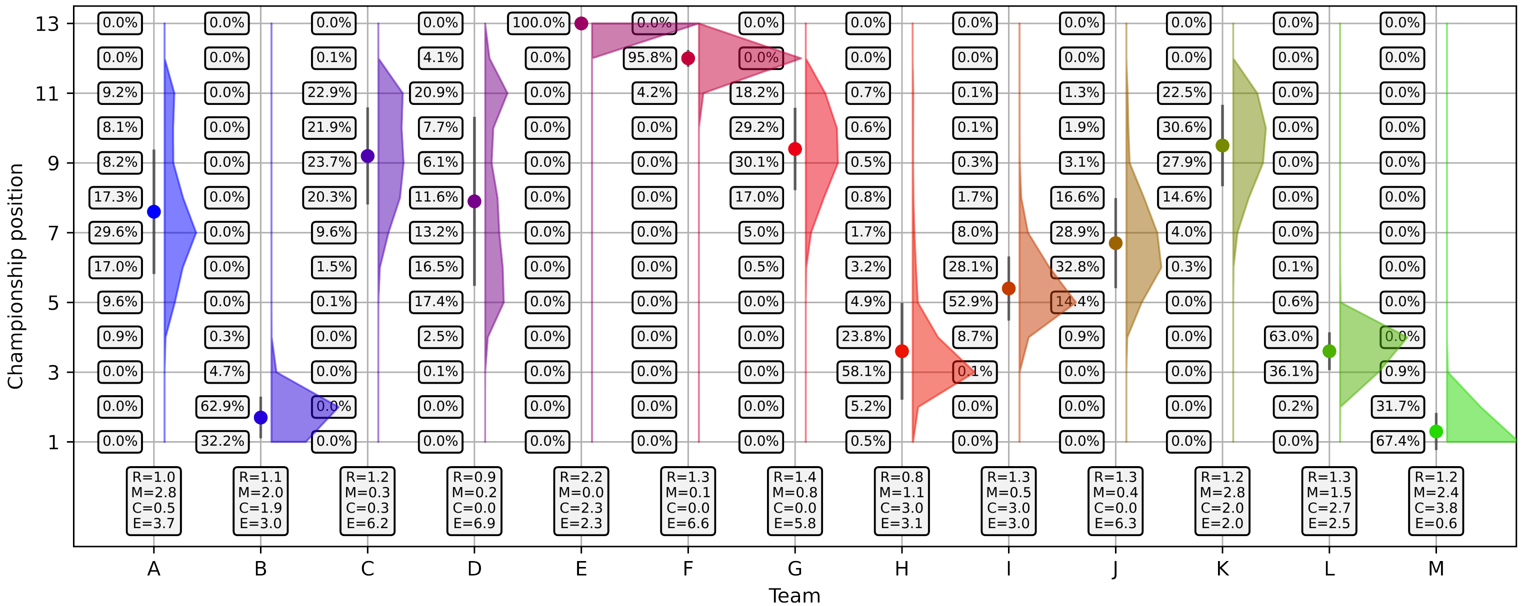

Figure 5.2: Summary of 1,000,000 simulated seasons with budgets given by Table 5.1, after the drivers have been allocated. The budget chosen in Chapter 4 has been assigned to Team A.

Figure 5.2 shows that the predicted finishing positions teams have much less variability after the drivers have been assigned. Unfortunately for us, our team now has a very low chance of winning a championship, and we are very unlikely to finish in 4$^{\text{th}}$ or above. Now, our team is most likely to finish between 5$^{\text{th}}$ and 11$^{\text{th}}$, with 7$^{\text{th}}$ being the most likely result. The championship fight is now expected to be between teams B and M, with no other team winning any notable number of championships in our simulations.

One reason for our team now having next to no chance of succeeding is largely down to the driver assignment. Because it is a random process, the "luck of the draw" plays a pivotal role in the success of teams. From Table 5.1, we can see that teams B and M both had a lower marketing budget than our team, but ended up with better drivers on average. This gives us an immediate disadvantage to these teams, as they have gained a large amount of expected performance for each race compared to our team, thus reducing our chance of winning the team championship.

Race discussions

Race	Optimal Strategy	Driver 1 Strategy	Driver 2 Strategy
Spain	H24 M48 M52 M66	H25 M46 M66	H26 M46 M66
France	H17 M27 M37 M48 M59 M70	H16 M27 M38 M49 M59 M70	H17 M28 M39 M50 M60 M70
Monaco	H30 S46 S62 S78	H34 M56 M78	M34 M56 M78
Belgium	H18 M31 M44	M15 S25 S35 S44	H20 S31 M44
Italy	H23 S35 H53	M18 S30 S42 S53	S15 S28 S41 S53
Hungary	H27 S40 S53 M70	H28 S42 S56 S70	H31 H52 M70
USA	H20 H37 H54 H71 H88	H18 S30 H45 H60 H75 M88	H22 H39 H56 H72 H88
Japan	H26 S38 H53	H27 M41 S53	H23 S33 S43 S53
Canada	H22 S34 S46 S58 S70	M21 S34 S46 S58 S70	H26 S39 M57 S70
Australia	H28 S43 S58	M24 M43 S58	M26 S42 S58

Table 5.2: Summary of strategies across the season with optimal and used strategies given for each driver. 'H25 M46 M66' indicates the driver started the race on the hard tyres, pitted on lap 25 for mediums, then pitted again for mediums on lap 46 for the final stint, completing the race of 66 laps.

Race	Winning Strategy	Driver 1		Driver 2	
		Quali.	Race	Quali	Race
Spain	H26 M46 M66	23	8	19	16
France	H12 H29 H43 H56 H70	11	4	11	11
Monaco	M31 M54 M78	10	3	10	5
Belgium	H20 H35 S44	25	15	14	9
Italy	H23 H40 S53	7	9	3	12
Hungary	H28 H49 H70	23	8	17	12
USA	S13 H31 H49 H68 H88	15	DNF	12	9
Japan	H26 S38 M53	18	8	13	9
Canada	H24 H47 H70	10	5	9	11
Australia	H29 H58	21	11	20	17

Table 5.3: Qualifying and race results for both drivers at each race along with the winning strategy.

Tables 5.2 and 5.3 shows a summary of the race results across the season for our drivers. In general, the optimal strategy consists of starting on the hardest tyres and then usually splitting the remaining laps into some number of equal (or nearly equal) stints on the same tyre compound, with only a few exceptions. As mentioned before, we cannot pit both drivers on the same lap, else the second driver would receive a penalty, so deviations from the optimal strategy are required.

Across the season, the philosophy behind the strategy choices was adapted according to results in previous races by seeing what worked and what did not. At Spain and France, where the strategies were both decided before any race had been completed, more conservative strategies were used that had potential race times close to the fastest possible strategy. In these races, we qualified on the hard tyres with both cars, which had a negative effect on qualifying performance, but the race performances were noticeably better. As these were balanced tracks, our results were largely as expected when considering the budgets and drivers of the other teams.

Monaco is a track where overtaking is notoriously challenging and as a result, the fastest strategy may in fact be terrible for the race! By making 3 pit stops, as suggested by the optimal strategy, track position would be sacrificed when the pit stops were made, and it may not have been possible to overtake enough cars to make up for this, due to the nature of the track. The importance of track position also increases the importance of qualifying, as it is inherently better to start closer to the front. This justifies qualifying on faster tyres, even if it results in a slower strategy. As a track that rewards the better drivers, we expected and achieved strong results at Monaco with both drivers, including a podium, as our team had one of the strongest driver line-ups.

Each of Belgium, Italy and USA rewarded teams with large engine budgets, which we opted for in favour of the chassis budget. However, the teams that chose a low marketing budget were rewarded much more than our team at these races, as they generally had a much larger engine budget as a result. At Italy we opted for an aggressive strategy with an additional pit stop compared to the optimal strategy as overtaking was easier at this track. Unfortunately, this did not result in strong finishing positions. This may be a combined result of a poor strategy choice and other teams gaining more from the engine bias at these tracks. Our result at USA was hampered by an unfortunate retirement for one of our drivers while he had the potential to score a good number of points.

Both Hungary and Japan rewarded teams with high chassis budgets, which we decided to opt against. This was the case for a number of other teams, however, so our performances were largely unaffected by this. The last 2 races were both balanced tracks, like the first 2, so our results were again as expected, even with the terrible qualifying results on the medium tyres at Australia.

5.3 Post-season discussion

Overall, our team collected 131 points which was enough to secure P5 in the Constructors' Championship which is around where we expected to finish. Our drivers collected 80 and 51 points respectively, enough to finish in P4 and P16. According to our predictions, we only had a 10.5% chance of finishing in P5 or better, so we performed about as well as we could given the competing teams and driver draw.

We could argue that our team was particularly unlucky with the driver draw, as the two teams with better driver line-ups both had a lower marketing budget which put our team on the back foot before the season even had even started. Of course, we could have used different strategies at each of the races, evidenced by the fact that the winning strategies were often different to the strategies we used, although it would be very difficult to find a strategy that is able to overcome the deficit to teams with outright quicker cars/drivers. Team M won the Constructors' Championship, and given that they opted for a similar marketing budget, this justifies the budget we chose. They were, however, the only team that prioritised the chassis heavily. This, combined with their strong driver line-up, gave them an immediate advantage at Monaco, Hungary, and Japan, and their overall strong car gave them an advantage at the 4 balanced tracks.

At Belgium and Australia, our results in the race were hampered by poor qualifying results. When devising strategies for a race, it is near enough impossible to predict exactly how much traffic our cars will experience in each stint. As a result, when a driver qualifies poorly, they are unable to make the most of their tyres as they end up stuck in traffic for at least the first stint. By qualifying better, drivers are less likely to be stuck in track simply because there are fewer cars ahead of them. We could argue that we were particularly unlucky at these races, and if they were run again we could see significantly different results. Of course, this is the same for all teams, as occasions on which we qualified well could have been helped by other teams qualifying poorly.

5.4 Adjustments to strategy with perfect knowledge

With our knowledge of the budgets of the other teams, we can use Monte Carlo simulations to find the optimal team budget(s) for competing against the 12 other teams in the competition. We will consider 7,500 possible budgets, allocated by the same method as in section 4.3, and then simulate 2,000 seasons with each of these budgets against the 12 other teams in the competition. We will then order these budgets by the number of championships they won (equivalent to win rate as all teams took part in 2,000 seasons) and look for any patterns between the teams that perform best.

Rank	R	M	C	E	Champs. won	**Win rate**
1	1.1	1.4	5.5	0.0	707	**0.353**
2	1.0	1.4	5.6	0.0	683	**0.342**
3	1.2	1.9	4.9	0.0	667	**0.334**
4	1.1	1.8	5.1	0.0	666	**0.333**
5	1.0	1.8	5.1	0.1	665	**0.333**
6	1.1	1.7	4.9	0.3	664	**0.332**
7	1.2	1.7	5.0	0.1	663	**0.332**
8	1.0	1.3	5.0	0.7	661	**0.331**
9	1.1	1.5	4.8	0.6	656	**0.328**
10	1.1	1.5	5.3	0.1	653	**0.327**

Table 5.4: Top 10 teams according to championship win rate after Monte Carlo simulations with 12 of the teams fixed across all simulations. Each team in the table competed in exactly 2,000 seasons.

Table 5.4 shows the randomly selected teams that performed best against the 12 other teams in the competition. Clearly, opting for $R = 1.1$ is one of the best options, as we noticed in the previous chapter. This is because such a reliability budget appears to provide the best balance between pace and reliability. It is important to note that this reliability budget is lower than the majority of the other teams in the competition, which shows that this team would have a little more budget to allocate to other aspects of the car without risking too many more retirements. This would give this team a slight advantage in performance relative to the other teams.

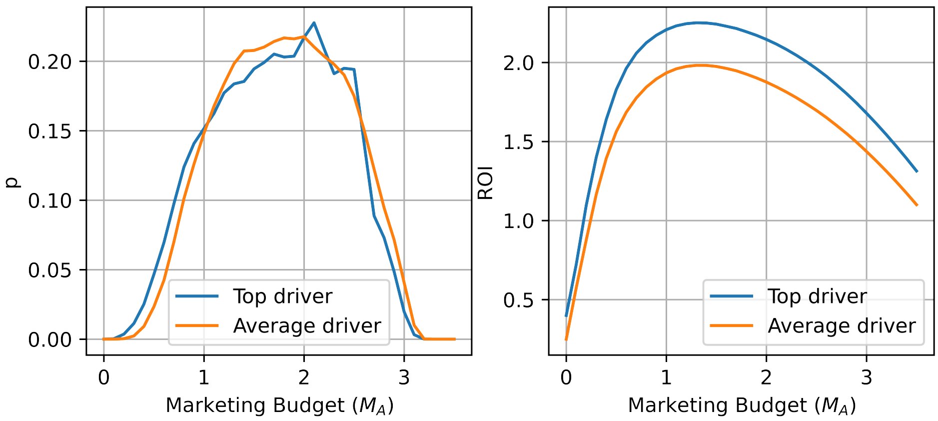

Figure 5.3: Probability of having the best return on investment on driver performance out of all teams (left) and average return on investment (see Example 4.3) for both average and top driver performance rating (right). The marketing budgets for teams B to M have been fixed, and we have repeated the driver draft process 1,000,000 times for each $M_A \in [0.0, 3.5]$, where M_A is the marketing budget of team A. The erratic nature of the curves can be explained by the unequal differences between the ordered driver performances.

We can see that the best option would have been to opt for a marketing budget of $M \approx 1.4$ and also focusing heavily on the chassis rather than the engine. This is a much lower marketing budget than we originally predicted to be optimal which is largely a result of many teams opting for a relatively low marketing budget. Figure 5.3 (left) shows the probability of getting the best return on investment on the drivers given the marketing budget. We can see that this peaks for the average driver performance around $M = 1.8$, but when $M = 1.4$ the probability p is only slightly lower. For the top driver performance, the curve is slightly altered, and even a little erratic, but we can also see that at $M = 1.4$, the probability p is only slightly lower than the apparent peak. From this, we can say that a team can choose $M = 1.4$ and have a good chance at getting the best performance gain from the driver draw. Getting the best return on investment gives a team an advantage over the other teams which increases their chance of winning the championship.

The right plot in Figure 5.3 shows the expected return on investment given the marketing budget of team A. Interestingly, this peaks at a slightly lower budget than the probability plot. This tells that hoping for the best return on investment may not necessarily ensure greater success in the competition, as what matters is the return on investment compared to the other teams, not just the return on investment in absolute terms.

By allocating the majority of the remaining budget towards the chassis, we would have had by far the largest chassis budget which would have resulted in very strong performances at the chassis-biased tracks. The main competition comes from team M who have a lower chassis budget and a similar engine budget. As a result, at chassis-biased tracks, this optimal team would have an immediate advantage over the strongest team in the competition, while at engine tracks they would not lose out by much, if at all. At tracks that favour the driver, team M are likely to gain due to their better chance of being allocated the better drivers, but their drivers would need to have a high enough performance rating to overcome the chassis and engine deficit. The same can be said about balanced tracks, although this time team M would need even stronger drivers to overcome the performance deficit. Overall, the high chassis budget would give the team a much greater chance of winning the competition, as they would outperform the previous best team in most situations.

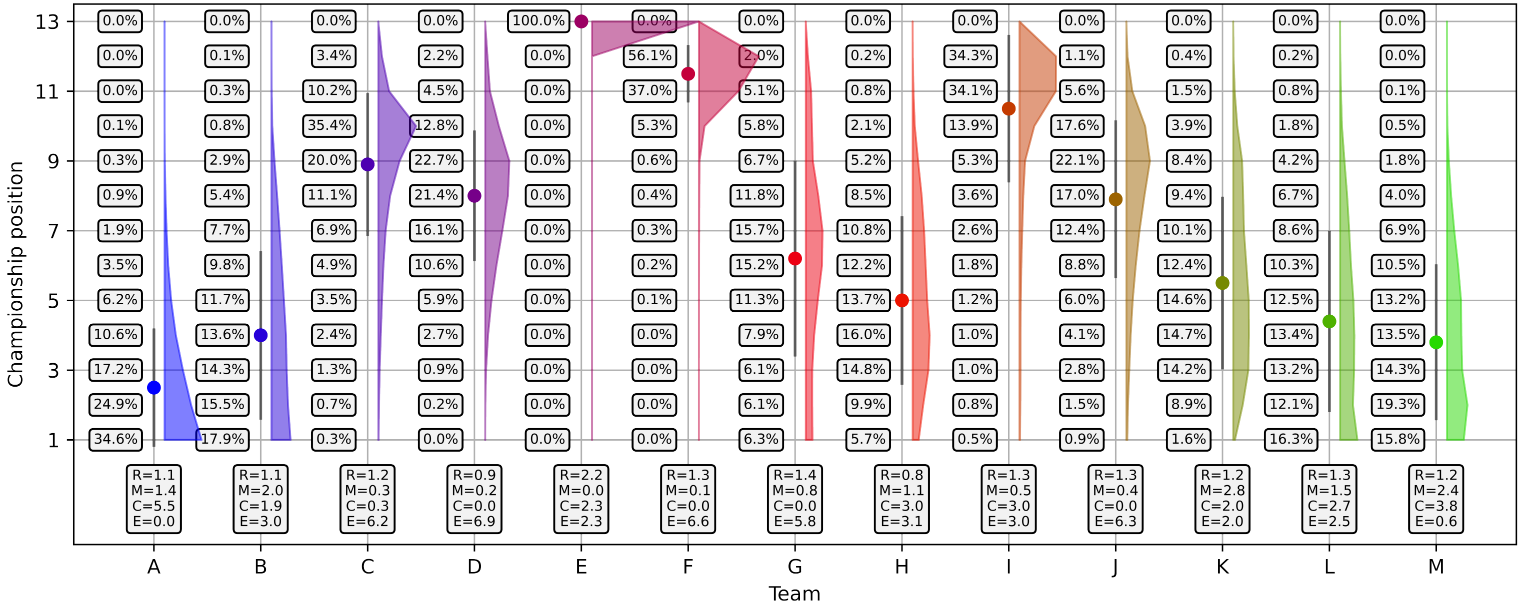

Figure 5.4: Summary of 1,000,000 simulated seasons with optimal budget from Table 5.4 (Team A) against 12 other teams in the competition.

From Figure 5.4, we can clearly see how successful this optimal budget is when competing against the other teams. Across 1,000,000 seasons, the best budget from Table 5.4 won around 34.6% of the championships they took part in, even with the drivers being redrawn each time. This is much greater than the win rate for any team taking part in the original competition, where the best team would only have won around 25% of the time. Team M, who won the championship in the competition, now only wins 15.7% of the time when the optimal budget is included, and they are not even the second most successful team! Team B now has the second best chance of winning with an 18% success rate. Team M loses out when this optimal budget is included because Team A would very likely have the best performance at chassis-biased tracks, and possibly balanced tracks too, thus worsening Team M's results across the season. This further demonstrates the relationship between every teams' budget and the performance of each team, and shows that a successful team in one set of budgets may not be successful in another.

Clearly, opting for one of the budgets in the table above would have resulted in a very good chance of winning the championship. In some ways, we could have predicted that opting for a high chassis budget would have been optimal, as we had good reason to expect most teams to favour the engine. However, it would have been very difficult to predict the marketing budgets of every team, so we could not realistically have expected such a low marketing budget to be so successful.

Chapter 6

Discussion

In this project, we have developed a simple lap time model and used it to formulate strategies for a given race. We then used our model in a real-life scenario by taking on the role of a Strategy Engineer at the 70$^{\text{th}}$ Anniversary Grand Prix, held at Silverstone in 2020. Having used the lap times of a number of drivers from the race and practice sessions, we were able to devise an "optimal" strategy that would have beaten the race winner. We also compared our optimal strategy to both the winning strategy and the fitted winning strategy to see how they deviate. By calculating the total race time for the winning strategy using our model, we found it to be accurate within 0.3% of the winning time. Having seen our model applied in a real world scenario, we covered various in which we could improve the model to capture features in a driver's lap time that were previously ignored.

The remainder of this project covered both the preparation and participation of the F1 Strategy Competition, run by Cédric Beaume. We used Monte Carlo simulations to choose a budget allocation that we felt gave us the best chance of winning the teams championship, involving a number of assumptions about how other teams may act when choosing their own budgets. We finished this project with an analysis of the results of the competition where we discussed our race strategy decisions, and compared our budget allocation to that of the other teams. With our knowledge of every other team's budget, we used further Monte Carlo simulations to find the budget allocation that would have given the greatest chance at success in the competition.

One area into which we could study further is more advanced extensions and modifications to our simple lap time model. We have already discussed a number of possible modifications to our model, for which we could study their direct effect on the strategies that a team would use in a race. For instance, we could consider the effect of the car's setup on its performance on a single lap and across a stint. When creating the model, we did not even mentioned car setup, but it arguably plays the most significant role on a car's lap time. Across a race weekend, drivers and engineers work to find the best car setup that balances qualifying and race performance. A car with a bad setup could be too slow in a straight line or in the corners, or it may experience greater tyre wear rates. We could adapt our model to consider the effect of the setup on the tyre wear of the car, for example.

When selecting our race for the study in Chapter 3, we discarded any races that featured wet weather conditions as we were unable to model it. This was an unfortunate but necessary decision, as strategic decisions in changeable conditions are often some of the most crucial decisions in a whole season. By being on the right tyre at the right time in the wet, a driver can gain several seconds per lap! Given the rarity of wet races in Formula 1, it is very difficult to calculate exactly how much an effect rain has on lap time, and ultimately all that matters is whether you are quicker than everyone else. This, along with other extensions we have suggested are certainly not an exhaustive list, but they leave plenty of opportunity for further study into their combined effect on a driver's lap time. The end goal for a Strategy Engineer is to have an exhaustive model that considers every possible effect on lap time, but such a model would be unwieldy at best, and it likely does not even exist. After all, the joy in all sport comes from what we fail to predict.

Regarding the strategy competition, when discussing the budgets, we mentioned that it was closely linked to game theory, as we have incomplete information about the decisions other teams would make. We saw in the competition that the winning team prioritised the chassis budget, when we predicted that prioritising the engine budget would result in a greater chance of winning the championship. This is another example of incomplete information, as we did not know how the other teams would split their budget between the chassis and the engine. When we performed a retrospective analysis of the chosen budgets, we were able to easily find a budget that would comfortably outperform all other teams in the competition. It was also particularly interesting to see that our optimal budget resulted in the previous best budget now falling behind some of the other teams. Clearly, the success rate of a team is highly dependent on its competitors; we have established that the success rate is governed by relative performance compared to other teams, and not the absolute performance. Going into further detail about the relationship between our budget and the other teams' budgets may have resulted in our team performing better in the competition, although there was no guarantee of success.

Bibliography

[1] Formula 1, *Ruth Buscombe On The Role Of The F1 Strategy Engineer*, https://www.youtube.com/watch?v=K7EEBE-2Wf4&ab_channel=FORMULA1, accessed on 16/04/21.

[2] Cédric Beaume, *F1 Strategy Competition - Oct 2020*, http://cbeaume.com/download/MATH3001_Formula1.pdf, accessed on 16/04/21.

[3] FIA, *2021 Formula One Sporting Regulations - May 2020*, https://www.fia.com/sites/default/files/2021_formula_1_sporting_regulations_-_iss_2_-_2020-05-27_0.pdf, accessed on 16/04/21.

[4] *Silverstone*, https://www.formula1.com/en/racing/2020/Great_Britain/Circuit.html, accessed on 21/04/21.

[5] *Red Bull Ring*, https://www.formula1.com/en/racing/2020/Austria/Circuit.html, accessed on 21/04/21.

[6] FIA, *2021 Formula One Technical Regulations - May 2020*, https://www.fia.com/sites/default/files/2021_formula_1_technical_regulations_-_iss_7_-_2020-12-16.pdf, accessed on 06/05/21.

[7] DHL, *DHL Fastest Pit Stop Award*, https://inmotion.dhl/en/formula-1/fastest-pit-stop-award/, accessed on 16/04/21.

[8] Keith Collantine, *2020 70th Anniversary Grand Prix interactive data: lap charts, times and tyres*, https://www.racefans.net/2020/08/09/2020-70th-anniversary-grand-prix-interactive-data-lap-charts-times-and-tyres/, accessed on 19/04/21.

[9] *Lap Time Data*, https://www.kaggle.com/cjgdev/formula-1-race-data-19502017, accessed on 15/10/20.

[10] Mark Hughes, *How shrewd race craft and brilliant tyre management allowed Hamilton to snatch victory in Portugal*, https://www.formula1.com/en/latest/article.how-shrewd-race-craft-and-brilliant-tyre-management-allowed-hamilton-to.4kTjQpByG0SYoqiKf7cbJQ.html, accessed on 03/05/21.

[11] RaceFans.net, *Silverstone - circuit information*, https://www.racefans.net/f1-information/going-to-a-race/silverstone-circuit-information/, accessed on 21/04/21.

[12] Formula 1, https://twitter.com/F1/status/1291759508876865537/, accessed on 21/04/21.

[13] FIA, *70th Anniversary GP Event and Timing Information*, https://www.fia.com/events/fia-formula-one-world-championship/season-2020/formula-1-70th-anniversary-grand-prix, accessed on 20/04/21.

[14] Pirelli, https://twitter.com/pirellisport/status/1292023053778587649/, accessed on 21/04/21.

[15] IMDB, *Alain Prost - Biography - Personal Quotes*, https://www.imdb.com/name/nm1169354/bio?ref_=nmbio_ql_dyk_1#quotes, accessed on 05/05/21.

[16] Pirelli, https://twitter.com/pirellisport/status/1292480389525594117/, accessed on 19/04/21.

[17] Autosport, *F1 terms explained: what box, marbles, DRS, undercut and more mean*, https://www.autosport.com/f1/news/f1-terms-explained-what-box-marbles-drs-undercut-and-more-mean-5477591/5477591/, accessed on 06/05/21.

[18] Edd Straw, *What caused F1's no-grip Friday? The big questions answered.* https://the-race.com/formula-1/what-caused-f1s-strangest-friday-the-big-questions-answered/, accessed on 03/05/21.

[19] Cédric Beaume, *F1 Strategy Competition - Oct 2020*, http://cbeaume.com/download/F1StratCompetition.pdf, accessed on 11/12/20.

[20] Wikipedia, *Error function*, https://en.wikipedia.org/wiki/Error_function, accessed on 11/12/20.

[21] Cédric Beaume, *F1 Strategy Competition - Season 1 Report*, http://cbeaume.com/download/Teaching_F1_S1.pdf, accessed on 17/02/21.

[22] Cédric Beaume, *F1 Strategy Competition - Season 2 Report*, http://cbeaume.com/download/Teaching_F1_S2.pdf, accessed on 17/02/21.

[23] Cédric Beaume, *F1 Strategy Competition - Season 3 Report*, http://cbeaume.com/download/Teaching_F1_S3.pdf, accessed on 11/12/20.