LUISS

Double Degree Master's Thesis

Department of Economics and Finance
Major in Finance

Chair of Econometric Theory

The Optimal Stopping Strategy in Dynamic Contests: Empirical Evidence from Formula 1 Races and Financial Applications

Prof. Santucci De Magistris Paolo
SUPERVISOR

D'Agostino Francesco
CANDIDATE
ID: 729811

Prof. Morelli Giacomo
CO-SUPERVISOR

Academic Year 2021/2022

TABLE OF CONTENTS

Introduction

Motivation and research question

Economic and financial models show agents that maximize their objective functions while interacting in their dynamic and competitive contexts. They establish their goal and identify obstacles that become constraints in a modelling environment. Competition in contests forces agents also to study others' characteristics and status, and sports are one of the most important examples of competition, so they are also important for general audience since they are an analysis of choices and strategies in a dynamic context with constraints.

They perfectly recreate a challenging environment, just like an economic sector with two or more firms competing or a financial market with two or more investors, that may not be in competition, but still have an influence on others' choices; thus, sports can serve as a model for them and for several more.

A very competitive team sport is Formula 1, and modelling and optimization are some essential elements of it. Like most of the sports, first of all the contest is studied before happening, as well as tactics. In Formula 1, even if it seems like the driver is alone in the contest to win the race, there is an entire team that only studies racing strategies, since these can overturn the result. Teams find themselves battling against other nine teams. It is a very dynamic environment, since choices made by strategists must be taken considering others' choices and state. The most important strategies concern the start of the race, for example how to approach the first curves and where to attack or defend the opponent, the behavior between the two drivers of the same team and the pit-stop above all. Even though strategists study tactics before the race, most of them are constructed or modified during the event, because the environment is unpredictable and in constant change, and opponents' moves could impact own choices. So, strategists must be able to adapt their tactics to the new situation and the changing during the race using simulation, algorithms and human intuition overall.

The pit-stop strategy, that is the choice of the moment in which to stop, is the most thought-out choice because the driver's surviving instinct is out of its scope, unlike the approach in the first curves. Due to regulations, drivers are obligated to take at least one pit-stop during the entire race[1], changing their tyres' type. The stopping strategy in F1 is highly effective and a wrong choice could compromise the

[1] This is the case of dry races. In wet (raining) races there is no obligation to stop, meaning that stopping strategies don't necessary show up.

entire race. It must carefully consider own capacities, because in the immediate laps after the pit-stop the driver must be able to recover the time lost due to the stop, circuit's characteristics and it must especially consider others' choices and state. An earlier pit-stop could give the driver the opportunity of the "undercut", a typical F1 strategy in which first-mover advantage could arise by stopping before the direct opponent. The principle of the undercut is that, in a contest between two or more drivers, the one who stops before will have at least one lap (trip around the race circuit) with a fresher tyre that corresponds to more performance, compared to the other drivers that still have to stop. In this way he will gain an advantage in terms of time and this advantage will last until the opponents will stop to level off the "freshness" of tyres. Of course, stopping the driver early comes with some problems: the risk of putting him into the traffic of other cars and the disadvantage of having a more worn-out tyre for the rest of the race, meaning slower laps.
On the other side of the coin, a late pit-stop will put the driver himself at risk of others' undercut, while giving him a fresher and higher-performance tyre for the end of the race. Here, there is a clear trade-off between performance and strategy, two fundamental factors to be balanced with difficulty.

That said, a Formula 1 race can be seen as an optimal stopping problem, that is a problem of optimization concerned with the choice of the moment in which executing a determined action in a finite horizon of time, in order to maximize an expected reward or to minimize an expected cost. Since in Formula 1 there is finite time composed by N laps in which to execute the action of the pit-stop, it is a perfect environment to solve an optimal stopping problem.
The objective of this optimization is to maximize the driver's position at the end of the race with the optimal execution of the pit-stop. On the basis of this execution, the driver will find himself in a better or worse position, also in terms of distance from opponents, since he can gain or lose time because of the pit-stop.
It is possible to find similarities between the choice of the lap in which to stop and when an investor should buy or sell a stock, exercise an option or when a firm should start selling a new product, etc.

My goals in this original work will be to calculate for the first time which is the optimal stopping strategy in a F1 race, to understand how much the stopping strategy influences the probability of winning and which other factors have a higher influence on this probability.
I will use a dual approach. From one side, I will investigate empirically what factors shape the choice of a stopping time and increase the probability of winning. This will be done through several regressions, analyzing the choice for different drivers or circuits but also taking into consideration general results. From the other side, I will construct different formal models and look for an

equilibrium stopping time. The reason is that empirically it's not possible to find a general equilibrium stopping time unless circuits are individually studied. In fact, since different circuits come with different specifications, each one of them has a different equilibrium stopping time. So, only a formal model can specify the equilibrium stopping time, whereas an empirical test can describe the optimal stopping time in a specific circuit. Furthermore, I will introduce competition slowly. Initially, the models will be ideal, assuming that the driver is alone in the circuit and only takes into consideration own conditions. Then, I will analyze how competition and interaction with opponents' impact on the strategy and on the behavior.

Finally, I will introduce the optimal stopping in the context of American options. I will show different approaches, but I will analyze the Markovian one, that is similar to my Formula 1 approach. I will explain American options' payoff and gain function and characterize the value function that contributes to the choice of the optimal stopping time.
Then, similarities and differences between the optimal stopping in Formula 1 and American options will be shown.

This thesis examines the sport of F1 in depth, with more than 100.000 data and several regressions from different points of view. My work contributes to the literature on the topic since it is a completely original work in which optimal stopping is applied to the Formula 1 environment, both theoretically and empirically, for the first time. The similarities between Formula 1 and American options' models show the close linkage between the two contexts, and they are useful to give even more importance to the contribution that my thesis brings to the optimal stopping literature. My models can serve as proxy for a lot of backgrounds, with or without competition.

Literature Review

In general, the literature about optimal stopping, dynamic contests and strategies, optimization and forecasting is very extended. One possible division of the research papers in this subject matter can be between the empirical literature on dynamic contests and the theoretical literature on the optimal stopping time. Sports find vast space in both directions since they recreate challenging environment in which the objective is to maximize the probability of winning and some of them come across optimal stopping problems.

My thesis relates to the empirical literature on dynamic contests and to the formal one on the optimal stopping time. I will use new data, since Formula 1 has never been studied under this scope, and I will introduce a competitive environment in the optimal stopping time problem, with the analysis of interactions between opponents and competitors' status and characteristics.

Generally, all literature papers related to sports or to a realistic setting characterize a formal model at first, then apply the model to the data and finally show results. So, just as my work, they bring a mix of theoretical and empirical methods. In empirical literature researches, a recurring approach is to set up the problem as a Markov game and then use dynamic programming to solve the optimization problem.

Hirotsu and Wright (2003) wanted to find the optimal substitution strategy in football matches by a Markov process model and dynamic programming, in order to maximize the expected number of league points, while Kira, Inakawa, Fujita and Ohori (2015) used the same procedure to study optimal decision making in the baseball field. Hoffmeister (2018) instead used the same approach but in the strategy optimization of beach volleyball.

Percy (2015) went through a similar process by using dynamic learning to select the best strategy and to predict the outcome in different sports, such as cricket, football and badminton. He slightly modified the procedure by considering the sports as stochastic processes.

Romer (2006), too, used dynamic programming to decide whether a football team in the NFL should kick the ball or try to go for a first down. He compared this choice with the ones that firms try to optimize every day.

Brown (2007) analyzed the behavior of players in golf tournaments when a superstar is participating and he compared it to the behavior inside firms, when there is internal competition.

Knoeber and Thurman (1994) compared the production of broilers with the tournament theory, finding consistent evidence with their prediction in the strategies construction.

Another application of Markov chain has been made by Pfeiffer, Zhang and Hohmann (2010). They studied elite table tennis competition, going deep into the interaction between the two players, understanding the best strategies in order to maximize the probability of winning.

Goldman and Rao (2011, 2014) brought dynamic efficiency and optimal stopping to the NBA. First, they studied the optimal decision making in a basketball match through dynamic and allocative efficiency, then they applied optimal stopping to the 24 seconds possession that teams have.

Gauriot, Page and Wooders (2018) found evidence from the biggest tennis competition, Wimbledon, of optimal solution used by players in a Nash equilibrium.

Finally, Jeyapragasan, Karra and Krishna (2019) used deep learning and Markov decision processes to predict the outcomes of English football.

Of course, optimal stopping is strongly relevant also in finance, with both theoretical and empirical studies in these settings. Jacka (1991) and Herdegen (2009) studied the pricing of American put options through an optimal stopping problem solved with a typical approach of parabolic free-boundary problem. Guo Xin (2001) extended the research of optimal stopping problems on the pricing of Russian options, linking a hidden Markov process to the classical Black-Scholes model. Dai, Zhang and Zhu (2011) applied it to the stock market, finding the optimal strategy through a trend following system and thresholds, while Belomestny (2013) used the approach of dual optimization discovered by Rogers (2002) and Haugh and Kogan (2004) in order to solve an optimal stopping problem in the context of option pricing.

Strack and Viefers (2013, 2018) wrote about decision makers in job and stock markets introducing the regret sentiment and factor in this optimal stopping problem.

Strategic and optimal behaviors are also examined in depth in generic dynamic contests and backgrounds, as game theory. Differently from the previous, these papers are theoretically examined, with model constructions and a formal solution.

Dixit (1987) analyzed a Nash equilibrium in a contest between two players that try to find the best solution to win the prize and he showed that there is a first-mover advantage in contests.

Multiround contests were introduced by Tsetlin, Gaba and Winkler (2003) using variability and handicaps as the main factors of study, while Beaumont (2010) used the typical dynamic programming to solve the problem of an optimal strategy in a multi and subsequential rounds contract bridge tournament.

Anderson and Cabral (2007) instead examined the variance in a Markov differential game, while comparing this with the choices that happen inside firms, showing some riskier and other safer.

Moscarini and Smith (2007, 2011) created a scoring function to deeply analyze incentives and trade off with respect to being the leader or the runner-up, in the design of a two players' dynamic contest.

A curious study has been made by Seregina, Ivashko and Mazalov (2019) concerning the TV show "The price is right". They used this competitive environment to model the strategic behavior through optimal stopping.

Ryvkin (2021) characterized a Markov equilibrium and a competition in continuous time, as a Brownian motion. He analyzed effort and luck as main factor influencing the performance.

Hinnosaar (2016, 2022), too, discussed effort in a sequential contest. He also showed the advantages of first and earlier movers and the similarities of these contests with firms' environment.

I will start by characterizing some formal models, just like theoretical papers, in order to find the optimal stopping time solution with and without competition. This solution will be a generic one that can be adapted race-by-race, since the specific optimal stopping strategy depends on the circuit taken into consideration. For this purpose, I will empirically analyze historical results and strategies in specific circuits. This will let me find the stopping lap evaluated as the best by the teams and compare it to the formal one that I previously found.

Also, empirical regressions will take into account all the factors that most influence the probability of winning in every race, and not only the stopping lap.

Finally, a classical theoretical setting will introduce the optimal stopping in American options and some similarities with the Formula 1 models will be underlined. Markov approaches and the analysis of some continuation and stopping values are the most important.

Chapter 1. Basics of Formula 1

Every year the Formula 1 World Championship takes place worldwide. With an average of 20 races every season, 10 teams and 20 drivers face themselves to win the two Championships: the constructor's one and the driver's one, with the latter being the most important. Before the season starts, teams develop their cars to construct the fastest one; the car's speed and reliability are some essential factors to win. Of course, the driver is the one who drives the car, so a good team needs a talented one to have a perfect performance and vice versa.

With a positive performance, probabilities of winning will increase. The races take place on Sundays, but the winning framework starts on Saturdays in the qualifying. During this session, drivers challenge themselves to score the fastest lap around the race circuit. Their starting position, technically called "grid position", in the race of Sunday is based on the best lap they scored in the qualifying of Saturday. So, who had the best lap will start first, with the other drivers following based on their Saturday's ranking; the starting position is a very great advantage.

In the race, the drivers have to complete a pre-determined amount of laps around the circuit.

Every circuit has a different length, and the number of race laps depends on it. Once the leading driver completes all the laps, the chequered flag is waved, and the race ends when all the cars cross it. Points will be awarded to the first 10 drivers, according to their ending order. Of course, as previously written, drivers and teams' skills are essential to win.

But there is a crucial factor that influences the probability of winning: the race strategy. It comprehends the start of the race, how to approach the first curves and where to attack or defend the opponent, the interaction between the two drivers of the same team and in particular the choice of tyres and the moment in which to change them, called pit-stop. This is crucial because, as I will show, there exists an optimal moment in which to execute the stop. If one wrongly moves away from it, he could ruin his race because tyres' duration management must be run properly. If he stops too late, he could lose much time since tyres deteriorate with laps increasing. If he stops too early, he could lose much time because he doesn't exploit all tyres' potential.

Teams are allowed to have several types ("compounds") of tyres at disposal. Tyres can be softer, with a better performance but a shorter duration, or harder, with a longer duration but a worse performance. In this analysis, I will assume that drivers choose the same compounds of tyres, so that the study will focus on the stopping strategy and not on the technical-mechanical strategy.

When it rains, wet or intermediate tyres are available for the teams. In dry races, drivers are obligated to stop at least once to change their tyres' compound. When the teams will ask them to do it, they

will enter the box, namely the area with the garage, and mechanics will substitute the four tyres. In the box lane, speed limiter on cars is active: cars can move between 60 and 80 kilometers per hour (depending on the circuit) due to safety rules. Normally, without considering the time to go through the box lane, a pit-stop is executed in 2-3.5 seconds; a longer one is considered slow. For this reason, mechanics need to be fast since a slow pit-stop can cost racing positions.
So, teams try their best to minimize the break time and to optimize the moment in which execute the stop.
They must think very well about their stopping strategy. In Formula 1 races, overtake "on the circuit" is quite difficult. In fact, circuits are narrow, an overtake can be dangerous and if one car isn't way faster than the one preceding it, the overtake will be almost impossible. This explains even more why the pit-stop strategy is essential. An overtake can be executed not only on the circuit, but also through a pit stop.
There are two main strategies. The most common one is called the "undercut". In a contest between two drivers, this is typically used by the runner-up, that acts like the first mover. He chooses to bring the pit-stop forward, and to stop before the leader, in order to have fresher tyres with respect to the opponent (still with old and consumed tyres). During the "out-lap", that is the lap after the pit-stop, there will be a significant difference in performance due to the difference in tyres (runner-up's new vs. leader's worn). The runner-up must exploit this to gain an advantage in terms of pace, that will allow him to recover time and distance, and hopefully will let him gain the position after the leader's pit-stop.
The negative side is that the driver who undercuts will have more consumed tyres for the end of the race, meaning a worse performance because these tyres were put on very early. Also, it is not easy to have a good out-lap, because tyres need to be heated, and if this process is not well executed by the driver, he will fail to gain time.
Furthermore, the earlier you stop, the higher is the risk to get into the traffic of other cars. The reason is that the break at the boxes costs a lot of time due to the speed limiter. Of course, the undercut is very effective for the runner-up, so the leader may choose to be the first mover and to anticipate himself the pit-stop in order to avoid the opponent's undercut.
Another possible strategy is the opposite one, "going long", technically called "overcut", a late pit-stop that will give the driver a fresher and higher-performance tyre for the end of the race but will put him at risk of others' undercut. The overcut is very strong if the other driver struggles to heat the new tyres or if he finds himself in the traffic of other cars.

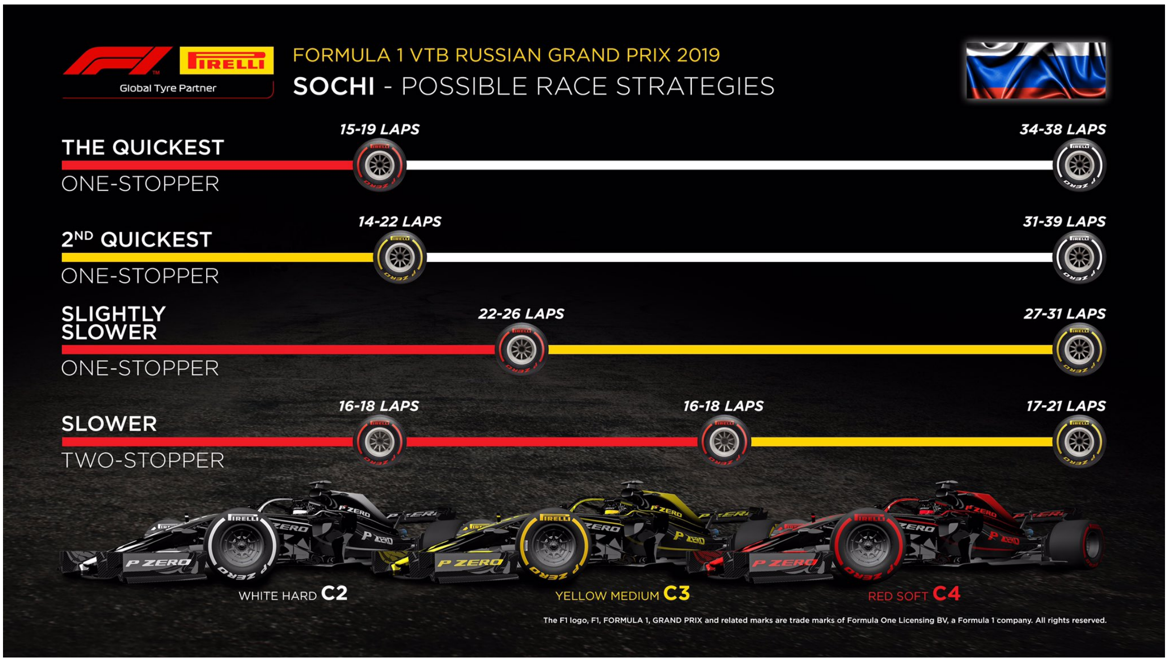

Figure 1 2019 Russian Grand Prix Stop Strategies

During a Formula 1 season an average of 20 races are run. These races take place in different circuits, each one with a different configuration and layout. The most basic difference between circuits is the number of laps, due to the length of the circuit. More specific differences comprehend the temperature of the air and the stress brought to the tyres, based on whether the curves are slow or fast. All these factors affect the ideal lap of pit-stop. Of course, in a hot-weather circuit with a lot of stressful curves, like Brasil, tyres will deteriorate faster, while in a cold-weather circuit, like Russia or Austria, they will last more.

This is why by using historical data it's impossible to find a broad optimal stopping lap. Circuits with less laps and high degradation will have an earlier optimal lap of stop with respect to circuits with more laps and less degradation. Therefore, only a formal model can define an always valid optimal stopping moment in Formula 1 races, with variables adapted as needed, while an empirical model can define circuit-specific optimal stopping lap.

So, the objective of these theoretical models will be to generalize the optimal stopping time, whereas the empirical model will show the actual optimal stopping lap for individual circuits.

Chapter 2. Methodology

2.1 Formal theoretical model setup

2.1.1 Individual decision-making

As for other optimal stopping problems, it is easy to associate the stopping decision-making process with a Markov Decision Process (MDP), in which there are finite states and actions. Formula 1 races are composed by a finite number of laps in which to execute a finite number of pit-stops. Then, there is a driver, the decision maker, that faces an environment whose state changes in response to the actions made by the agent. Also, the changing of the status of the environment determines the immediate reward/payoff obtained by the decision maker. At each time, the agent observes the current state and executes an action, namely executing the pit-stop or continuing the race. So, the final objective of the agent will be to maximize his total reward through his decisions, that is maximizing the points in the end of the race.

Defining the action space of the MDP as the actions of stopping and continuing, and the state space as the current lap, the current driver position (that is particularly important to estimate the expected payoff) and tyres' condition, it is possible now to setup the theoretical optimal stopping model.

To simplify the model and its tractability I will start by using the concept developed by Goldman and Rao (2011 and 2014), thanks to similarities between their work, "Optimal Stopping in the NBA", and mine. Initially, I will assume an ideal individual decision-making in which the strategy depends solely on own's conditions and state, as if the driver is racing alone in the circuit and his goal is to finish the race as fast as possible. The environment defined by Goldman and Rao is less competitive than mine because they only consider choices that can be made by players of the same team. So, I will introduce interaction with opponents in a competitive decision-making in the next sub-chapter.

In a 24 seconds possession, basketball players must decide when to shoot. So, the action to be executed in their optimal stopping problem is the shooting in a finite time of 24 seconds, whereas in mine it is the pit-stop in a finite time of N laps. At every lap, each driver has the possibility to stop. I will necessary assume that in the first 9 laps of the races, drivers don't stop. This would be inefficient both because in the first race's stages a stop would mean to get directly into the traffic of other cars and because the tyres' performance and potential will not be completely exploited. Furthermore, if a driver stops in the first laps it is probably due to an accident since the race is more chaotic in the beginning. Since my model must be generally applicable, I want to exclude accidental and random observations.

Also a pit-stop in the final moments is inefficient because tyres have an almost precisely defined life and they can't last for all the race, since they would break or explode.

That said, I will assume a finite time laps $N \in \{10, 11, \ldots, 50\}$ to execute the stop. Now, at every lap the driver and his team can decide whether to stop or to continue the race to the next lap when the same decision will be made. Of course, their decision will be based on maximizing the expected points obtained in the end of the race, by ranking higher.

So, the decision's process during every lap is based on the fact that the driver and his team observe an unbiased measure S_n (stopping value) of expected points they would get by stopping at that moment, with $n \in N$ being the current lap in which the choice is taken.

This measure, as properly stated by Goldman and Rao, is drawn from a continuous distribution with pdf f_n that has positive density on its support in order to have the existence of a unique inverse cdf.

In every lap, a reservation threshold value R_n is established, and it represents the marginal (in the sense of worst) opportunity of pit-stop lap, a continuation value. Here, the concept of dynamic efficiency comes in help to clarify this topic. As basketball players must efficiently allocate the shoot in Goldman and Rao's paper, in mine teams must efficiently allocate their pit-stop during the race. Therefore, they will execute it if and only if $S_n \geq R_n$. This means that, being R_n a future opportunity of pit-stop, the driver will continue to race if this value exceeds the expected points by stopping in that lap. So, a stop is executed when there is no better future opportunity that occurs with the continuation of the race. This is the standard of dynamic efficiency.

This continuation value will be chosen in order to maximize this objective function, as stated by Goldman and Rao:

$$V_n = V_{n-1} + \int_{R_n}^{\infty} (x - V_{n-1}) df_n(x) \quad (1)$$

$$V_{-1} = 0,$$

With V_n defined by Goldman and Rao as the expected points value of an unused stopping time.

Dynamic efficiency requires to be set $R_n = V_{n-1}$.

It is possible to find another way to define an optimal stopping time or, even better, to combine these two methods. What matters to drivers is the time used to complete the laps because by minimizing it the performance increases and so does the probability of winning. By simplifying and assuming a perfect lap without any external event, such as traffic and mistakes made by the driver, the time will depend only on the condition of the tyres, that is $c \in C = [0,1]$. Then, the time of a lap will be

$$t : C \rightarrow [0,\infty)$$

Clearly, there is a negative effect between tyres' condition and lap's time, with time decreasing:

$$c > c' \Longrightarrow t(c) < t(c')$$

Then, let's think about $t(c)$ in unit terms. A correct way of characterizing this function could be:

$$t(c) = \frac{1}{c}, \tag{2}$$

$$t(c) \equiv time\ for\ a\ lap\ when\ tyres'condition\ is\ c$$

but it is also important to specify a tyres' degradation function, ignoring external forces such as heat and sunlight, and so by assuming that they only get ruined because of the race proceeding.

$$f : C \rightarrow C, \qquad f(c) < c, \qquad f(c) \equiv tyres'\ degradation\ function$$

$$f(c) = \frac{c}{2} \tag{3}$$

So, by combining (2) and (3), the lap's time function becomes:

$$t(f(c)) = \frac{1}{f(c)} = \frac{1}{\frac{c}{2}} = \frac{2}{c} \tag{4}$$

Now, the total time of the race with the stop at lap n is:

$$T(n) = t(1) + t\big(f(1)\big) + \cdots + t\big(f(n-1)\big) + \\ t(1) + t\big(f(1)\big) + \cdots + t(f(N-n)) \tag{5}$$

The first part of the equation is the time until the stop at lap n. The second part is the time after the stop at n.

The objective of teams and drivers will be to minimize this total time T obviously. By minimizing this, the result will be a maximum exploitation of the tyres' performance for each "stint", that is the portion of race between pit-stops. This could be considered a naïve strategy because it doesn't consider the opponents at all. With the introduction of the other drivers, it can be used as a basis and be improved.

2.1.2 Competitive decision-making

Now, I introduce the opponents' behavior.

Ideally, a driver's strategy decision should be based on the previous models, only considering own tyres' condition, exploiting maximum performance from them, and stopping once the worsening condition is decreasing the lap time and the continuation value is lower than expected stopping value. This reasoning is even more valid once the driver is not involved in a direct fight against another driver. In fact, if he is quite safe from others' attacks and undercut strategies, in terms of time gap, he can keep racing and stop when is better for him.

But Formula 1 races are run in a highly competitive environment, and I'm interested in a dynamic and competitive analysis. For this reason, interaction and others' condition and states must be necessarily considered. The outcome of a strategy depends not only on own actions but also on the actions of others.

Simplifying the execution of a race, let's assume that a driver directly competes against only one other driver. I will then construct a model of an instantaneous game in which two players compete to gain an advantage with respect to each other. One is the leader, the other one is the runner-up. They take actions accounting for other's actions and own conditions. Obviously in my game, the action to be taken is the execution of the pit-stop.

Of course, the leader has a position advantage. He is fine with the current situation and would like not to change the state. Ideally, the leader could think like in an individual decision-making setting and stop by using the previous models. Instead, clearly the runner-up wants to change the situation; he can't reason like in an individual setting. If both players think ideally, the race will end almost surely in the current positions, with the runner-up ending the race behind the leader. For this reason, the runner-up is more aggressive, while the leader is more conservative, and he typically imitates runner-up's strategies to protect himself.

Then, the runner-up will try to come up with something different. As a matter of fact, usually, the following driver uses the opposite strategy with respect to the preceding driver. He has the advantage of seeing what the leader does. If the leader stops in front of him, he will keep racing; if the leader continues the race, he will stop.

In general, the best for him would be the undercut, trying to anticipate the stop and gaining time during the out-lap. This strategy is the most effective. But the leader knows this, and by keeping racing and using the ideal strategy (by considering only his conditions) he will fall in the trap of the opponent's undercut. So, he will in turn try to anticipate the stop even more. At this point, the runner-up will not find the stop useful anymore, because the situation would stay unchanged, in favor of the leader. Then, he will choose to continue the race by going long and by using a different strategy, that is stopping later.

It's now possible to start by formally constructing a simplistic 2-by-2 instantaneous game, for which I established numerical payoffs.

Runner-Up

Leader	Stop	Continue
Stop	(5,0)	**(3,3)**
Continue	(0,5)	(3,0)

Figure 2 2-by-2 instantaneous game: strategies and payoffs

Here, we have pure strategies: when a player makes a choice, the other one knows exactly which choice is better for him. The operating environment is deterministic.

Nevertheless, there is one important fact to highlight: I will only take into consideration actions and payoffs in lap n. This means that I will not consider the reactions of players in the lap n+m (with *m>0*). It is significant because in Formula 1 the leader is somehow automatically the first mover since

he crosses the line as first and, by doing so, he is the first to start a new lap. In lap n, he decides to perform or not to perform the action of the stop. The runner-up can react by doing the same.
Instead, in the case of the runner-up being the first executing the pit-stop at lap n, the leader can react (if he wants) by stopping only at least at lap n+1 (or even later), because he already crossed the pit-stop zone. Anyway, since in normal condition one lap of difference in tyres' freshness is enough for the runner-up to recover and gain some time with respect to the leader, the payoffs at lap n are much in favor for the runner-up: (0,5).
So, by assigning numerical positive payoffs, it is clear that we have a Nash Equilibrium in the up-right position, the Stop-Continue strategies. The leader has a dominant strategy in stopping. The runner-up knows it and he positions himself in the continuation strategy. Actually, this isn't his optimal strategy, since he would have a greater payoff by stopping; anyway, the leader knows that, and he would stop in turn.
So, the equilibrium is found when the leader stops, and the runner-up continues. This is easily detectable in real races: the leader is afraid of the following driver's undercut, so he anticipates the stop; then, the runner-up comes with something different and decides to keep racing and stopping later.
The payoffs also show the two different reasonings depending on the position. The leader wants to imitate the opponent's strategies, whereas the runner-up finds more utility by differentiating the strategy.
Clearly, actual and more advanced payoffs should consider utility depending on other factors, such as the previously mentioned continuation values of Goldman and Rao's model, the tyres' degradation function (so, how the driver is able to manage the tyre's life) and the gap between the drivers, all ingredients that make the undercut even more effective. Further future studies and analyses are needed to include utility-payoffs.

Chapter 3. Data and empirical specification

3.1 Data description and sample restriction

As previously said, the theoretical model is useful for general results without taking into consideration specifics of races and circuits. To account for these elements and to check the influences of many other factors, it is necessary to use historical data and then restrict them in several different empirical analyses.

The data were downloaded on *kaggle.com*, a huge database website that comprehends a lot of datasets. The Formula 1 dataset[2] consists of a lot of observations about races, drivers, teams (constructors), qualifying (starting grid position), circuits, lap times and pit-stops. It includes information from 1950, the first Formula 1 championship, to 2022, the current and last one so far. It was necessary to make a first essential restriction: the data were downloaded starting from 2011. The reason is that previous data on pit-stops are not available because rules were different: there were seasons with not mandatory pit-stops or seasons with mandatory pit-stops due to refueling.

The observations are composed as combination of data. For example, Lewis Hamilton (driverId = 1), ranked first in the championship (rank = 1), driver of a Mercedes (constructorId = 131) started the race in the third position (grid = 3) and ended the race in the first place (position = 1), by scoring 25 points (points = 25) in the Italian Grand Prix of 2018 (raceId = 1002), composed by 53 laps (laps = 53). He scored his fastest lap at lap 30 (fastestLap = 30), by stopping once (stop = 1) at the twenty-eighth lap (lap = 28) in a break of 23,728 seconds (duration = 23,728). He ended the race in a normal way (statusId = 1) without any accident or problem to the car. All this result is gathered together in the resultId = 24486, making it easy to select, filter or use data.

Another necessary data restriction was to eliminate random or abnormal events. That's why I selected only data with statusId equal to 1 (the driver ended the race without problems) or with the statusId that represents a driver overtaken by one or more laps. So, I didn't consider drivers that retired because they could mislead the results, for example by giving an unrealistic meaning to some data, and I didn't even consider drivers that ended the race but with some accidents and other problems because their final result will be inevitably influenced by the problems and would again mislead the results. In fact, a driver that has an accident or problems to the car will probably end the race in a

[2] https://www.kaggle.com/datasets/rohanrao/formula-1-world-championship-1950-2020

worse position and, in addition, he will probably be forced to have a pit-stop in an unexpected and unusual moment and not in the optimal one. For the same reason, I ignored data with an unusual long pit-stop due to problems at the box or to the car.

Finally, as explained in the formal theoretical model section, I will exclude all the observations with pit-stop executed before lap 10 and after lap 50.

In the end, after the data restrictions, my dataset comprehends a total of more than 7600 observations corresponding to more than 100.000 data. For more specific analyses, I will restrict data even more, for example related to a single circuit, to better analyze which is the optimal stopping lap in a determined race.

3.2 Empirical model setup

3.2.1 Individual decision-making

All the empirical models that I will use are multiple linear regressions

$$Y = \beta_0 + \beta_i X_i + \epsilon \qquad (6)$$

In the first models, I will study the influence of several variables on the probability of winning. The stopping lap is included in these variables. Then, I will restrict the data to a single circuit, analyzing how in actual races the stopping lap is chosen depending on different factors and I will study the optimal stopping lap according to historical data – demonstrating the prediction for which there exists a reasonable optimal stopping lap only when we talk about circuit-specific analyses.

In order to study the main influences on the probability of winning I had two ways of reasoning, both correct but different in the choice of the dependent variable. Of course, in Formula 1 only the first positioned driver is the winner. However, in my model, but also in reality, racing in a higher-ranked position increases the probability of winning, as higher positions bring more points at the end of the race. Because of this, both ending positions and points can be use as dependent variable that represents the win.

From a broad and general point of view, it is possible to think about maximizing the probability of winning not only a race, but the entire Championship. In the end of the season, each driver has all races' points summed up. Of course, the one with more points wins the Championship. Then, this

objective is achieved through achieving the more possible points during each race. However, the distribution of points at the end of each race don't display linearity, as instead assumed in my empirical specification. In fact, according to the ending positions, the points are given as follows: 25, 18, 15, 12, 10, 8, 6, 4, 2, 1.
After the 10^{th} place, drivers don't get any point. So, ending 11^{th} or 20^{th} would be the same with the points as dependent variable, but actually the probability of winning is higher when ending 11^{th} with respect to the 20^{th} position.
Given these considerations, I decided to use the ending position as dependent variable, as it displays linearity, too.
Clearly, in contrast to points, when we talk about position it is better to have a lower value, so dependent variables will have the opposite effect. I decided to run both regressions in order to show this difference in the results.

This model will take into account only own conditions and characteristics.
As X_i I will make use of several explanatory variables, but since I will run different analyses, sometimes I will omit some variables, or even use just one regressor in a simple linear regression, in order to establish the behavior of just one independent variable with respect to the dependent one.
So, my explanatory variables will be the driver, the constructor team, the starting grid position, the lap of the driver's fastest time, the rank of the driver in the championship, the number of stops, the lap of the pit-stop and finally the duration of the pit-stop.
Some variables aren't really explanatory, but my intention is also to demonstrate this.
After having assessed the probability of winning, I will analyze the optimal stopping lap, running different regressions using the lap of stop as dependent variable. The independent regressors will be the drivers, the constructors, the grid position and the rank. The other variables must be omitted because of simultaneity bias and reverse causality problems. Also these regressions will be run on all races first, and then on a circuit-specific framework.

Of course, my models have some limits. The drivers and the cars (constructors) are one of the main factors that bring to the victory. Unfortunately, it will be not numerically clear how an increase in these factors influence the probability of winning. The reason is that the drivers and teams' data have not a qualitative numerical order based on strengths and abilities. They are enumerated without a precise logic. To put this right, the rank of the current championship will come in help, because surely the better are the driver and the team, the higher will be their rank. Anyway, multicollinearity doesn't arise, because the rank is only related to the current season, and in Formula 1 power balances changes from one championship to the next, with a team or a driver that can be high-ranked in a year and low-ranked in the next one.

3.2.2 Competitive decision-making

To internalize the competitive nature of the sport, it is necessary to include rivals' status in the right-hand size of the regression equation. To do this, it is possible and useful to introduce a variable that takes into account the opponent's ability and the interaction between him and the driver taken into consideration. The most significant variable to explain this is the gap, that is the distance in time between the two drivers.

The gap incorporates a lot of information and easily shows the ability of the drivers on the track circuit: an increasing gap means that the driver in front is gaining time, whereas a decreasing gap means that the runner-up is gaining time. The gap doesn't need a Boolean variable to indicate if the driver in consideration is the leader or the runner-up. This characteristic is intrinsic in the gap, and it is the norm in Formula 1 terminology. If it has negative sign, it represents the distance that the leader has with respect to the runner-up; instead, if it has positive sign, it is the distance that the runner-up needs to recover.

Of course, the relationship between the lap in which to stop and the gap is a negative relation, since the increasing gap should encourage a driver to anticipate the stop in order to recover distance.

Unfortunately, gap data are not publicly stored, and I will not use them in the regressions. To account for opponents' status, condition and choices, further future studies and analyses are needed.

After the regressions, some results will be plotted, including variables coefficients, test statistics and distribution of residuals. The models will be run in R Studio software.

Chapter 4. Results

The results of this thesis satisfy all the expectations and predictions. All variables except one have significance in the general model that includes all races results from 2011 to 2022, without any specification in terms of circuits or drivers, but with the exclusion of random events.

```
Residuals:
     Min       1Q   Median       3Q      Max 
-12.5825  -1.9713  -0.3751   1.7061  14.8259 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)   -1.5426317  0.3181227  -4.849 1.27e-06 ***
driverId       0.0008381  0.0001039   8.063 8.82e-16 ***
constructorId  0.0017872  0.0004959   3.604 0.000316 ***
grid           0.4049706  0.0081461  49.713  < 2e-16 ***
fastestLap     0.0107167  0.0036338   2.949 0.003198 ** 
rank           0.3710390  0.0091447  40.574  < 2e-16 ***
stop           0.3729474  0.0562309   6.632 3.57e-11 ***
lap           -0.0022833  0.0045608  -0.501 0.616655    
duration       0.0618432  0.0092030   6.720 1.98e-11 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 3.105 on 6352 degrees of freedom
  (1 observation deleted due to missingness)
Multiple R-squared:  0.6747,	Adjusted R-squared:  0.6743 
F-statistic:  1647 on 8 and 6352 DF,  p-value: < 2.2e-16
```

Figure 3 General Regression Model Results: Y = ending position

```
Residuals:
    Min      1Q  Median      3Q     Max 
-16.230  -3.196  -0.194   2.788  35.070 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)   19.9021262  0.5032338  39.548  < 2e-16 ***
driverId      -0.0022302  0.0001644 -13.564  < 2e-16 ***
constructorId  0.0097512  0.0007845  12.429  < 2e-16 ***
grid          -0.5362141  0.0128862 -41.611  < 2e-16 ***
fastestLap    -0.0148631  0.0057483  -2.586  0.00974 ** 
rank          -0.4553491  0.0144659 -31.478  < 2e-16 ***
stop          -0.5124323  0.0889508  -5.761 8.76e-09 ***
lap            0.0113002  0.0072147   1.566  0.11733    
duration      -0.0891418  0.0145582  -6.123 9.72e-10 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 4.911 on 6352 degrees of freedom
  (1 observation deleted due to missingness)
Multiple R-squared:  0.5712,	Adjusted R-squared:  0.5706 
F-statistic:  1058 on 8 and 6352 DF,  p-value: < 2.2e-16
```

Figure 4, General Regression Model Results: Y = points

Even if the model with points being the dependent variable isn't really correct in empirical specifications, it is still interesting to use it in a comparison with the model with ending position being the Y.

Obviously, in both models the same variables have similar significance. This demonstrates that these two are correlated, since the points assigned depend on the ending position.

However, except for drivers and constructors that, as previously written, aren't qualitatively ordered by numbers, we can see that the regressors have opposite effects on the dependent variable. In fact, when an explanatory variable has a positive impact on the ending position, it has a negative impact on points (the lower the value of the position, the higher the points obtained).

It was predictable that the driver, the team ("constructorId"), the starting position ("grid") and the rank are the factors that have the highest significance on the general winning probability in Formula 1. In fact, the driver and the car are the two main actors, the starting position is a great advantage in motorsport, being difficult to the overtake, and the rank represents the performance over the year of the driver, meaning that a high-ranked driver should obviously have more chances to perform well.

Focusing on the regression with ending position, of course the grid and the rank have a very positive effect on the Y. The higher these two variables, that means a worse starting position and rank, the higher the expected ending position, meaning a worse result.

The coefficients of driver and constructor should be ignored. They aren't explicative because of their origin.

Less obviously, the number of pit-stops ("stop") and the duration of them explain the dependent variable of ending position. Both of them still have a positive numerical effect on it: the more pit-stops a driver executes, the worse he will rank; the slower the pit-stop, the worse will be his final position. Both these factors correspond to a loss of time. This highlights necessary ability and technique by the mechanics in executing a quick pit-stop and also the need of intelligent choices by the team in constructing the pit-stops strategy.

Oddly, the lap in which the driver executes his fastest lap has some statistical relevance. Probably, this is just fortuity since in the reality the fastest lap is not really related to the winning but is more related to the general conditions of the car.

In relation to pit-stops, the lap in which to stop is the least explanatory variable in this general model, with no statistical relevance at all. As highlighted and predicted previously, the stopping lap depends on circuits' characteristics; so, generally talking, it makes no sense to say that a specific stopping lap can influence the probability of winning in Formula 1. The intuitions are confirmed.

Instead, the analysis can be carried out singularly and restricted to a specific circuit to find out if the stopping moment has a higher impact related to a precise track.
I gathered together all Monza races from 2011 and by only analyzing the last eleven Italian Grand Prix results on regressors are much different, indeed.

```
Residuals:
    Min      1Q  Median      3Q     Max 
-6.4439 -1.4741 -0.2354  1.3925 13.0515 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)   -9.0622064  2.0360658  -4.451 1.30e-05 ***
driverId      -0.0002236  0.0004373  -0.511    0.610    
constructorId  0.0013917  0.0020917   0.665    0.506    
grid           0.4063261  0.0339126  11.982  < 2e-16 ***
fastestLap     0.0041194  0.0238784   0.173    0.863    
rank           0.3484721  0.0373421   9.332  < 2e-16 ***
stop           1.9317956  0.4113245   4.697 4.43e-06 ***
lap           -0.0315340  0.0222861  -1.415    0.158    
duration       0.3636066  0.0631263   5.760 2.53e-08 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 2.541 on 243 degrees of freedom
Multiple R-squared:  0.7586,	Adjusted R-squared:  0.7507 
F-statistic: 95.46 on 8 and 243 DF,  p-value: < 2.2e-16
```

Figure 5, Circuit Specific (Monza) Regression Results: Y = ending position

A surprising result, at first sight, is the statistical irrelevance of the stopping lap on the probability of winning. Actually, it isn’t surprising because every driver tries to stop at the optimal stopping lap of Monza (that will be defined later). Due to this, the stopping lap loses importance on the probability of winning.

These results confirm that the starting position is an evergreen element in Formula 1 races, regardless of a generic analysis or a circuit-delimited one. The Championship rank, too, has statistical importance, whereas the driver and the car have a minor influence on winning from a micro point of view; they count more on the long-term view in the entire Championship. This is reasonable since generally it’s unlikely that a driver and a car are skilled on a specific circuit through years. In Formula 1, except for special cases, there is a big turnover in drivers and also in title contenders from one year to the next. Furthermore, the sample (11 races) is too small to highlight the dominance of great champions.

Instead, the rank finds significance because typically balance of power doesn't change over the same year. A well-assembled driver-car pair will have more chance of winning regardless of the circuit.
The duration of pit-stop affirms its importance even in this circuit-specific model, showing that is crucial for mechanics to perform a good pit-stop.
The fastest lap, as predictable, doesn't find significance in this model, highlighting a more realistic result.
Also, this regression shows that the Italian Grand Prix is a one-stop race since an increase of 1 unit in stops has an effect of almost 2 negative units on position (a very heavy worsening in F1), whereas in the general regression it didn't have this impact.
Both of these models predict a large proportion (around 57% and 75% respectively) of variance of the outcome variable of position: they are a good fit. Furthermore, looking at p-values, in both models the F-statistic is statistically significant: the regression models fit the data better than the intercept itself with no independent variables.
In the appendix, residuals are plotted, and they have a normal distribution, meaning that my assumptions are valid as well as the models.

Now, to identify an empirical optimal stopping lap, I needed to use the lap in which to stop as the dependent variable. To start, some regressors used in the previous regressions must be taken out, since problems of reverse causality and simultaneity bias arise.
In fact, the fastest lap depends on when the pit-stop has been executed, giving the driver a greater opportunity to score a fastest lap.
The ending position, that was the previous dependent variable, must be eliminated, too. It is the effect of the choice of the stopping lap, and it can't be an explanatory variable. The duration of the pit-stop is an effect, too, and it can't be used as a regressor.
The number of stops shows simultaneity bias. As a matter of fact, the number of stops that a team wants to undertake influences the laps in which to stop (if there is a second pit-stop, it will surely be in the ending laps of the race), and vice versa the stopping lap has an impact on the number of pit-stops (if one anticipates the stop, he will probably need to stop again to have fresh tyres in the end of the race).
Then, the correct variables to be selected are the driver, the constructor, the starting position and the rank.

```
Residuals:
    Min      1Q  Median      3Q     Max 
-31.419 -11.671  -0.309  10.345  50.352 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)   26.8535718  0.3745741  71.691  < 2e-16 ***
driverId      -0.0006086  0.0004338  -1.403    0.161    
constructorId  0.0092681  0.0020661   4.486 7.37e-06 ***
grid           0.2600473  0.0343659   7.567 4.27e-14 ***
rank          -0.4244834  0.0379923 -11.173  < 2e-16 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 14.28 on 7616 degrees of freedom
Multiple R-squared:  0.01811,   Adjusted R-squared:  0.0176 
F-statistic: 35.12 on 4 and 7616 DF,  p-value: < 2.2e-16
```

Figure 6 Optimal Stopping General Regression Results: Y = lap

In this first regression, I used all the available data. All the variables, except for the driver, have statistical significance. The F-statistic confirms that the model fits the data better than the intercept alone.

Analyzing why the driver is an insignificant variable, while the constructor is a significant one, it can be thought that the team influences the strategy of both drivers, whereas the driver individually is not having an impact on it. It’s important to highlight again that the coefficients of these two variables aren’t quantitively defined.

The grid confirms the prediction of the leader anticipating the stop and the runner-up stopping later. The data demonstrate that the better the driver is positioned, the earlier he will execute the stop, trying to avoid others’ undercut.

That said, a general optimal stopping lap isn’t realistic when looking at all circuits, because each one has different characteristics, including degradation and number of laps. That’s why it is more important to individually look at circuits when empirically defining an optimal stopping lap.

```
Residuals:
     Min       1Q   Median       3Q      Max 
-24.4281  -5.1379  -0.3825   5.8291  25.3209 

Coefficients:
                Estimate Std. Error t value Pr(>|t|)    
(Intercept)   24.5871432  1.4295742  17.199   <2e-16 ***
driverId      -0.0009979  0.0016590  -0.602    0.548    
constructorId -0.0028373  0.0080144  -0.354    0.724    
grid           0.1059845  0.1274766   0.831    0.407    
rank          -0.1090627  0.1409499  -0.774    0.440    
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 9.808 on 247 degrees of freedom
Multiple R-squared:  0.005146,	Adjusted R-squared:  -0.01096 
F-statistic: 0.3194 on 4 and 247 DF,  p-value: 0.8648
```

Figure 7 Optimal Stopping Circuit-Specific (Monza) Regression Results: Y = lap

The first difference with the general regression is the p-value of the F-statistic. It demonstrates that the regressors have no predictive capability. This means that when looking at individual circuits, there exists a well-defined optimal stopping lap that only depends on the circuit's characteristics themselves, regardless of the starting position, the abilities (rank) and the driver-car couple. Of course, this lap should be modified race by race according to competition and tyres' condition, data that can't be incorporated since they are not available.

Nonetheless, it was possible to find the ideal stopping lap in the Italian Gran Prix (as it is possible to find it in the other circuits), demonstrating that when studying a single circuit, there is an optimal stopping strategy.

So, the optimal stopping strategy in Monza would be to stop either at lap 24 or 25, when considering a normal tyres' management and an individual decision-making. With the introduction of interaction and competition, this result could be slightly modified.

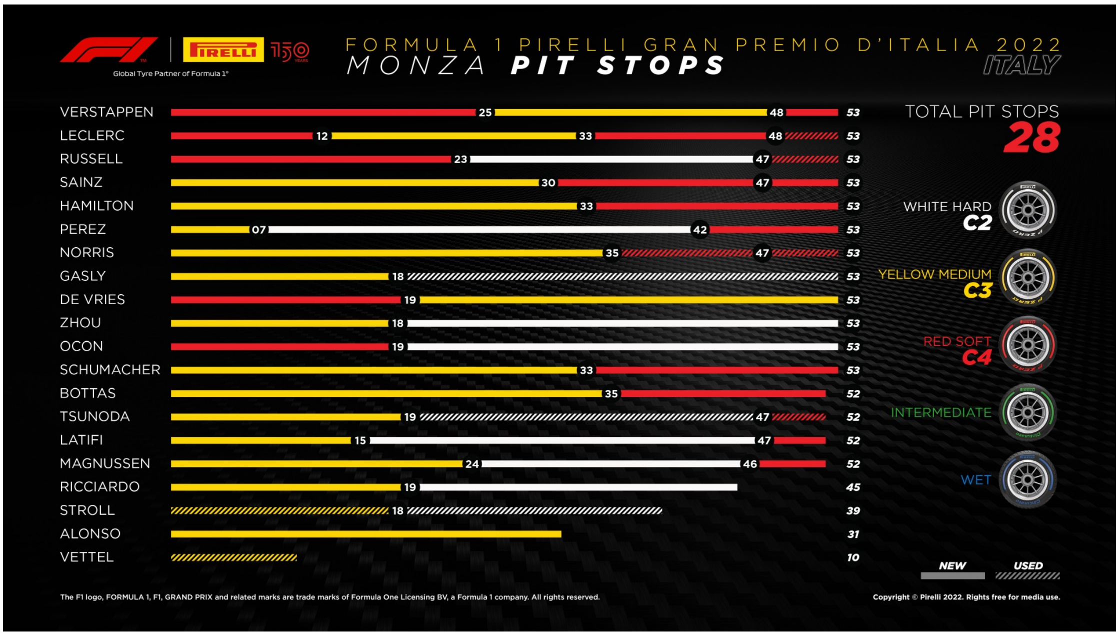

Figure 8 Italian Grand Prix 2022: pit-stop summary

This year's Italian Grand Prix (not included in the data) confirms the results and the theoretical models, too. It's important to highlight that 2022 regulation is slightly different from the past. This was a two-stop race, since there was a Safety Car in the end (lap 47), an advantage for who stops. Without considering that late stop and the early ones due to problems, a lot of drivers that were battling stopped between laps 18 and 19, anticipating the stop to try the undercut strategy. The leader, Verstappen, that had a clear pace advantage and was safe from others' attacks, had the opportunity to construct his strategy as an individual decision-making: he indeed stopped at lap 25.

PIT STOP SUMMARY					
CAR	DRIVER	START	PIT 1	PIT 2	PIT 3
3	RIC	C3n	C2n (22)		
4	NOR	C3n	C2n (24)		
77	BOT	C2n	C3n (26)		
16	LEC	C3n	C2n (26)		
11	PER	C3n	C2n (26)		
55	SAI	C3n	C2n (26)		
18	STR	C3u	C2u (25)		
14	ALO	C3n	C2n (25)		
63	RUS	C3n	C2n (26)		
31	OCO	C3n	C2n (26)		
6	LAT	C3n	C2n (23)		
5	VET	C3u	C2u (22)	C3n (26)	
99	GIO	C3n	C2n (1)	C3n (25)	
88	KUB	C2n	C3n (26)		
47	MSC	C3n	C2n (23)		
9	MAZ	C3n	C4u (24)	C2n (28)	C4n (39)
44	HAM	C2n	C3n (25)		
33	VER	C3n	C2n (23)		
10	GAS	C2n			
22	TSU	C3n			

C2 = Hard C2 | C3 = Medium C3 | C4 = Soft C4
n = new | u = used

Figure 9 Italian Grand Prix 2021: pit-stop summary

The Italian Grand Prix in 2021, with typical regulation about tyres and cars, confirms the results even more: the drivers (without accident or technical problems) stopped around lap 24-25.

Here it's also possible to point out the precision of the formal competitive decision-making model in the battle between Verstappen (VER) and Hamilton (HAM). Taking into consideration the optimal stopping lap (24.5 according to empirical results), the leader (VER) decided to anticipate the pit-stop at lap 23, whereas the runner-up (HAM), seeing the decision of the leader, reacted later with a stop at lap 25.

In the appendix, also these residuals are plotted, and they have a normal distribution.

The empirical results confirm the predictions presented in the theoretical models. Formal models are valid in general but must be adapted with drivers' conditions and circuit-by-circuit because, as demonstrated empirically, results are different when looking at individual races with respect to entirety.

Each circuit has its own optimal stopping lap. The optimal one can be found empirically, but, depending on competition and degradation, as established by theoretical models, it can and perhaps must be modified. The dynamicity of the Formula 1 environment determines an uncertainty factor for which, even if it is possible to state the ideal stopping lap, teams must be ready to adapt their strategies to single race's conditions.

Chapter 5. Financial Applications

The presented models find applications in several backgrounds. The optimal stopping problem in Formula 1 shows typical characteristics that can be found elsewhere, and they are the maximization of a payoff, strategies, dynamicity and interaction with the surroundings.
Financial markets are a very dynamic environment in which countless new data and information keep coming continuously and influence the markets by changing trading directions, securities prices and volumes. Strategy is essential to achieve positive investment results. Certainly, the other agents have a big impact on one's investment decisions, but differently from Formula 1, agents are not in competition.
In contrast to that, both investors and drivers try to maximize their objective function in their dynamic contexts, visualizing the situation in a specific moment in time as a Markov Decision Process.
Drivers want to optimize their ending position every race, while investors will try to maximize the realized profit.
In financial markets, a lot of instruments' payoffs depend on when they are executed. For instance, it is easy to think about American options. Then, we can illustrate the choice of when executing American options as an optimal stopping problem, just as Formula 1 pit-stops, that is finding the stopping time that maximizes the expected gains.

5.1 American Options

An option is a financial instrument whose price is derived from an underlying asset. When an investor buys an option, he enters a contract in which he has the right, but not the obligation, to buy or sell a pre-determined amount of the underlying security at a pre-determined price (known as strike price) on or before a specified date. If he buys a call option, he is in a long position and he has the right to buy the underlying; whereas, if he buys a put option, he is in a short position and he has the right to sell it. In order to have the right to exercise the option, the holder pays an option premium to the seller.
It's important to differentiate between American options and European options. American options can be exercised in every moment from the opening of the contract to the expiration date, whereas European options can be exercised only at the expiration date.

I will focus on American options, since they are obviously the only type for which is possible to find an optimal stopping time.
Since the call option is a long position, the investor will profit when the price of the underlying asset increases. The final payoff is the difference between the spot price (namely the current market price) and the strike price at the expiration date (if it is negative, the investors doesn't exercise the right to buy the security, and the payoff will be zero). Vice versa, when an investor buys a put option he will profit when the price of the underlying asset decreases. Here, the final payoff is the difference between the strike price and the spot price at the expiration date (if it is negative, the investors doesn't exercise the right to sell the security, and the payoff will be zero).
The intrinsic value of an option is the value it would have if exercised at the moment.
So, considering an American option issued at $t=0$ with expiration date $T \in (0,\infty)$, the intrinsic value at $t \in [0,T]$ is equal to

Call option $$(S_t - K)^+ = \max(0, S_t - K) \tag{7}$$

Put option $$(K - S_t)^+ = \max(0, K - S_t) \tag{8}$$

With S_t = *spot price, K = strike price.*

Now, defining $\tau = T - t$ as the option's residual time, it is possible to express the value of the option as a function of the underlying price, the strike price and of the residual time.
For a call option, the value at t is $C\ (S_t, K, \tau)$, whereas for a put option it is $P\ (S_t, K, \tau)$.
Then, at the expiring date $\tau = 0$, we have

Call option $$C\ (S_t, K, 0) = (S_t - K)^+ \tag{9}$$

Put option $$P\ (S_t, K, 0) = (K - S_t)^+ \tag{10}$$

Instead, before expiry, the option value must be at least equal to its intrinsic value.

When an option has a positive intrinsic value is said to be in-the-money (ITM). When it hasn't, the option is out-of-the-money (OTM). When the strike price is the same as the underlying price, the option is said to be at-the-money (ATM).
The time value instead is any premium in excess of the intrinsic value before the expiration date. It is the additional amount the investor is willing to pay over the intrinsic value for additional time until the expiration. In fact, the more the time remaining, the greater the probability for the option to be in-the-money. As the expiry date gets closer, the time value diminishes until zero on the expiration.
It's important to highlight that comparing two equal options (same strike price, same underlying price) with two different expiring date, the one with the longer residual life will have a higher value.
So, in an American option the timing is essential to maximize the extraction of value from the option.
To solve an optimal stopping problem one can use two possible approaches. The first one is the Martingale approach, for which there are two methods: the backward induction and the essential supremum. Typically, the first one is used for discrete time, whereas the latter is used for continuous time. The second approach is the Markovian one.
When choosing the model, one should look at the probabilistic structure of the stochastic process. I decided to use exhibit the Markovian approach that is very similar to the model used for the Formula 1 application. This technique gives both the optimal stopping time τ in which to stop the process and the optimal value, called value function V, that is the smallest supermartingale dominating the gain process, a parallel approach to the Martingale one in which the same function is executed by the Snell envelope.
My objective will be to solve the optimal stopping problem using this approach and relating it to that of Formula 1, with continuation and stopping values, to demonstrate similarities and applications.
Just like in Formula 1, the optimal stopping problem in American options can work as a Markov Decision Process. Then, at every moment in time, the investor will observe his opportunity of expected profit obtained by closing the contract at that specific moment.
Let's start by defining a Markov chain $X = (X_t)_{t\geq 0}$ on a probability space $(\Omega, \mathcal{F}_t, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ that takes values in the measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$.
It is necessary to define a value function and a gain function that have a very similar purpose to my stopping value and continuation value.
Introducing a risk-free asset with value $B(t) = e^{rt}$, with interest rate r, and defining $\tau \in [0,T]$ as the stopping time in American options, it is possible to define the arbitrage-free price of American options at time $t=0$. But first, it is necessary to state the fundamental theorem of asset pricing.

Let $\mathbb{P}$ be a risk-neutral probability measure and $\{\mathcal{F}_t\}_{t\geq 0}$ be the filtration of the underlying probability space. With *T>0,* let *D* be a $\mathbb{P}$-integrable and $\mathcal{F}_t$-measurable random variable.
Therefore, the arbitrage-free price of a derivative D at time t ∈ *[0, T]* is

$$D(t) = \mathbb{E}[e^{-r(T-t)}D|\mathcal{F}_t] \tag{11}$$

The probability measure $\mathbb{P}$ makes the discounted derivative price $e^{-rt}D(t)$ a martingale.

Then, the arbitrage-free price of an American option is

Call option
$$C = \ sup_{\tau \in [0,T]} \ \mathbb{E}[e^{-r\tau}(S_t - K)^+] \tag{12}$$

Put option
$$P = \ sup_{\tau \in [0,T]} \ \mathbb{E}[e^{-r\tau}(K - S_t)^+] \tag{13}$$

To describe the gain and the value functions, we first assume that $S_t = x$ is the value of a stock with volatility σ.
Now, identifying the time coordinate as *E = [0, T] × (0,∞),* we define the gain function *G : E* ⟼ *[0, K]* by

Call option
$$G\ (t,x) \coloneqq e^{-rt}(x - K)^+ \tag{14}$$

Put option
$$G\ (t,x) \coloneqq e^{-rt}(K - x)^+ \tag{15}$$

By discounting, we can simplify the gain function as

Call option $$g(t,x) = e^{rt}G(t,x) = (x-K)^+ \tag{16}$$

Put option $$g(t,x) = e^{rt}G(t,x) = (K-x)^+ \tag{17}$$

Instead, for *(t,x)* $\in$ *E,* the value function is

Call option $$\begin{aligned} V(t,x) = \; & sup_{\tau\in[0,T-t]}\, \mathbb{E}_{(t,x)}[G(X_\tau)] \\ = \; & sup_{\tau\in[0,T-t]}\, \mathbb{E}_{(t,x)}\left[e^{-r(t+\tau)}(S_\tau - K)\right] \end{aligned} \tag{18}$$

Put option $$\begin{aligned} V(t,x) = \; & sup_{\tau\in[0,T-t]}\, \mathbb{E}_{(t,x)}[G(X_\tau)] \\ = \; & sup_{\tau\in[0,T-t]}\, \mathbb{E}_{(t,x)}\left[e^{-r(t+\tau)}(K - S_\tau)\right] \end{aligned} \tag{19}$$

also written as

Call option $$v(t,x) = e^{rt}V(t,x) = \; sup_{\tau\in[0,T-t]}\, \mathbb{E}_x[e^{-r\tau}(S_\tau - K)^+] \tag{20}$$

Put option $$v(t,x) = e^{rt}V(t,x) = \; sup_{\tau\in[0,T-t]}\, \mathbb{E}_x[e^{-r\tau}(K - S_\tau)^+] \tag{21}$$

for which we define an optimal stopping problem, with *T* being the upper boundary of time coordinate *E* and τ being the stopping time in *[0, T-t]*.

Finally, using the Markovian framework, we state

Continuation set $$C \coloneqq \{(t,x) \in [0,T) \times (0,\infty) : V(t,x) > G(t,x)\} \tag{22}$$

Stopping set $$S \coloneqq \{(t,x) \in [0,T] \times (0,\infty) : V(t,x) = G(t,x)\} \tag{23}$$

As anticipated, the value sequence $(V^{T-t}(X_t))_{\{0 \le t \le T\}}$ is the smallest supermartingale dominating the gain sequence $(G(X_t))_{\{0 \le t \le T\}}$ under $\mathbb{P}$ for x fixed.

In the end, the optimal stopping time is

$$\tau_D = inf\{0 \le t \le T : X_t \in S\} \tag{24}$$

5.2 Formula 1 and American options

This whole reasoning is very similar to the Goldman and Rao's one. In both Formula 1 and American options, the Markov process is essential. Drivers and investors observe states and payoffs, and according to that they decide when to stop and execute the action.
In Goldman and Rao's model, a reservation threshold value is set through dynamic efficiency. That threshold is the continuation value, that is the marginal or worst opportunity of pit-stop lap. When the stopping value, that is observed at lap *n*, exceeds or equals the continuation value, the action of pit-stop is executed.
Here, the gain function is observed at time *t* as the payoff of the American option. Then, one stops when the stopping time τ makes the value function, that is the smallest supermartingale dominating the gain process, equal to the gain function, that means to extrapolate maximum gain from the option.
The difference is that in American options the time is the cause: it is chosen in order to maximize the payoff and it is a key factor in the decision of stopping. Instead, in Formula 1 it is a consequence or

just a surrounding element: drivers and team choose their best opportunity to stop in the race, depending only on the payoffs. Then, they visualize which was their optimal stopping lap.

In spite of the differences, the models could be interchangeable with some adjustments due to the specifics of the backgrounds. Of course, in Formula 1 the competition is one of the main elements. Instead, in financial markets competition doesn't show up clearly.

Conclusion

In every competitive environment, whether it is a financial context or a generic dynamic contest, agents work towards the win confronting themselves with competitors. They set an objective and, by knowing own characteristics, choices are taken through optimization. The optimal stopping is a problem of optimization: the decision of the moment in which executing an action.

In my work, I used Formula 1 as the main environment, and I implemented optimization and theories that could be used in a financial or economic context, such as American options.
Formula 1 is a very competitive field with 20 agents that simultaneously aim for the win. My objective was to find a theoretical way to optimize their choices with particular attention to the pit-stop, to empirically test what factors more influence the winning probability and how drivers and teams behave in reality when it comes to pit-stop.
The focus was on the optimal stopping time. The results, that can be deepened and continued in further studies, confirm the expectations: the optimal stopping time is found by formal models if general conclusions are needed, and can be adapted race by race, circuit by circuit, according to specific characteristics; instead, it is found by the empirical model and historical data if one wants circuit-specific results.

So, three formal models were presented, one of them developed by Goldman and Rao (2011, 2014) in their work on the optimal stopping problem in the NBA.
In the first model, every lap the strategists keep an eye on a reservation threshold (a marginal opportunity of pit-stop lap), that is the continuation value of the race and compare it to the pit-stop (stopping) value in terms of expected points. They will execute the stop once the continuation value is less than or equal to the pit-stop (stopping) value. It is based on dynamic efficiency and marginality.
The second model instead is based on minimizing the entire time of the driver's race by choosing the optimal stopping time according to the tyres' condition function only, assuming that this is the only factor influencing the entire race time and that the driver is racing alone around the circuit.
This leads to the third model, that introduces the competition and the interaction with the opponents. Here, a 2-by-2 instantaneous game is presented. Two drivers have two possible actions: to stop or to continue the race. One of them is the leader, the other is the runner-up. With numerical payoffs, the

Nash equilibrium is found when the leader chooses to stop, and the runner-up decides to continue to race. In fact, the leader will anticipate the stop to protect himself from others' undercut strategies, whereas the runner-up will try to come up with something different by differentiating the strategy and not stopping: he will go long and stop later.
The first model is adaptable to every race; the second one is usable when the driver isn't threatened by opponents' attack and has a considerable advantage on chasers. The third model studies the interaction between drivers.
All of these models work and find confirmation in reality. The best solution would be to use them together and to mix them with empirical ones.

My empirical tests were multiple regressions in order to assess which factors influence the most the probability of winning, to study how much the stopping lap counted in historical results and how the stopping lap is actually chosen.
Furthermore, I found that factors of influence have different weights whether the analysis is related to the entire history or to a single circuit. Drivers and cars have a higher influence on the long-term and on the totality of the Championship, whereas they have less on specific circuits. In fact, the drivers and teams' dominance are more related to eras and years than to circuits, that find themselves several protagonists.

A relevant outcome was that the starting grid position and the championship rank are the most important factors related to winning in F1. Looking at races in the same circuit, a lot of variables lose importance, the grid position and the rank stays on the top winning factors. Duration of pit-stop and number of pit-stops are both important in a macro and in a micro view, with the latter being particularly impactful when carrying on a circuit-specific analysis.
The stopping lap hasn't a statistical relevance neither in the general model nor in the circuit-related one. It was predictable, since generally it's impossible to find a precise lap number in which to stop: circuits have different total number of laps and characteristics. Furthermore, analyzing an individual circuit, every driver will want to stop at the optimal lap. Due to this, it loses importance; by the way, the optimal stopping lap still exists, and it is studied empirically.

Firstly, using it as the dependent variable of the regression, it was analyzed from a general point of view, in all the races. There, it wasn't possible to establish an optimal one, since it depends on several variables that must be adapted race-by-race, circuit-by-circuit. These variables are the starting grid position, the rank of the driver and the car.

Instead, when looking at individual circuit, an optimal stopping lap arises. As a matter of fact, in the restricted-Monza model, the regressors have no predictive capability: the optimal stopping lap is explained by the intercept only. It has a value of 24.5, meaning that in a normal race in Italy a driver should stop between lap 24 and 25. Of course, this optimal lap can be adjusted taking into consideration formal models, so accounting for tyres' degradation and competition, giving once again importance to a right mix of the two kind of models (theoretical and empirical) that are coherent and find confirmation between each other.

So, even if it loses importance as an impactful factor on the final winning since everyone pursue it, it is shown that the optimal stopping lap exists, found theoretically for general purposes and empirically for specific circuits, and the drivers can't move much away from it, or their race will be ruined. They should instead slightly modify it according to competition and own conditions.

The optimal stopping time in Formula 1 is as important as in a financial or an economical context. The action must be executed by taking into account the entire environment and conditions, by checking every moment what is worth more between continuing and stopping and by looking at historical data.

Important financial instruments for which optimal stopping is crucial are the American options. The reasoning behind their optimal stopping problem is very similar to the one behind my first model in Formula 1 and the problem can be modelled very likewise to the problem in Formula 1, even if the competition is not included in American options.
I found similarities in the use of the Markov framework, since drivers and investors observe states and payoffs, and in the usage of a threshold, that is the continuation value in F1, and the gain function in American options. Both drivers and investors look at these values and are ready to execute the action once a condition is satisfied.
In the case of F1, the stopping value must get greater than or equal to the continuation value; in the case of the American options, the value function must be equal to the gain function.
One big difference instead, is in the usage of time: in Formula 1, it is passive. It is only a consequence of the previously-presented choice.
Instead in American options, time has an active role: it is the cause, not the consequence. It is chosen in order to satisfy the condition and maximize the payoff.
Another difference is the presence of competition: in Formula 1 it is central, whereas in American options it isn't an element to keep into consideration.

Despite the differences, the models can be interchanged because they are based on the same concepts. Formula 1 can be used as a proxy for other contexts where time has an important value and impact on the final payoff and where optimal stopping problem arises. With adjustments related to the specifics of the surroundings and the characteristics, the models shown are much adaptable.

List of References

Anderson A. & Cabral L. M. B. (2007), "Go for Broke or Play It Safe? Dynamic Competition with Choice of Variance", The RAND Journal of Economics, Autumn, 2007, Vol. 38, No. 3 (Autumn, 2007), pp. 593-609, Wiley, Available At SSRN: https://www.jstor.org/stable/25046326

Beaumont N. (2010), "Using Dynamic Programming to Determine An Optimal Strategy In A Contract Bridge Tournament", The Journal Of The Operational Research Society, Vol. 61, No. 5 (May 2010), pp. 732-739, Palgrave Macmillan Journals, Available At SSRN: https://www.jstor.org/stable/40608235

Belomestny D. (2013), "Solving Optimal Stopping Problems Via Empirical Dual Optimization", The Annals of Applied Probability, 2013, Vol. 23, No. 5, 1988–2019

Brown J. (2007), "Quitters Never Win: The (Adverse) Incentive Effects of Competing with Superstars", Job Market Paper

Dai M., Zhang Q. & Zhu Q. J. (2011), "Optimal Trend Following Trading Rules"

Dixit A. (1987), "Strategic Behavior In Contests Author", The American Economic Review, Dec., 1987, Vol. 77, No. 5 (Dec., 1987), Pp. 891- 898, Published By: American Economic Association, Available At SSRN: https://Www.Jstor.Org/Stable/1810215

Downarowicz A. (2010), "The First Fundamental Theorem of Asset Pricing", Revista de Economia Financiera

Gauriot R., Lionel Page L. & Wooders J. (2018), "Nash at Wimbledon: Evidence from Half A Million Serves", Available At SSRN: Https://Ssrn.Com/Abstract=2850919

Goldman M. & Rao J. M. (2011), "Allocative and Dynamic Efficiency In NBA Decision Making", MIT Sloan Sports Analytics Conference 2011

Goldman M. & Rao J. M. (2014), "Optimal Stopping in The NBA: An Empirical Model Of The Miami Heat", Available At SSRN: Https://Ssrn.Com/Abstract=2363709

Guo Xin (2001), "An Explicit Solution to an Optimal Stopping Problem with Regime Switching", Journal of Applied Probability, Vol. 38, No. 2, pp. 464-481, Applied Probability Trust

Herdegen M. (2009), “Optimal Stopping and Applications: American Options”, University of Warwick

Hinnosaar T. (2022), “Optimal Sequential Contests”, Available At SSRN: Toomas.Hinnosaar.Net/Contests.Pdf

Hirotsu N. & Wright M. (2003), “Determining the Best Strategy For Changing The Configuration Of A Football Team”, Department Of Management Science, Lancaster University, UK

Hoffmeister S. (2018), “Sport-Strategy Optimization With Markov Decision Processes”, Universität Bayreuth

Jacka S. D. (1991), “Optimal Stopping and the American Put”, Mathematical Finance 1 (2): 1-14, University of Warwick

Jeyapragasan K., Karra M. & Krishna G. (2019), “Predicting the Game: Using Deep Learning And Markov Decision Processes To Predict English Premier League Outcomes”

Kira A., Inakawa K., Fujita T. & Ohori K. (2015), “A Dynamic Programming Algorithm For Optimizing Baseball Strategies”, Faculty Of Mathematics, Kyushu University

Knoeber C. R. & Thurman W. N. (1994), “Testing The Theory Of Tournaments: An Empirical Analysis Of Broiler Production”, Journal Of Labor Economics, Apr., 1994, Vol. 12, No. 2 (Apr., 1994), pp. 155-179, The University Of Chicago Press

Moscarini G. & Smith L. (2007, 2011), “Optimal Dynamic Contests”, Meeting Papers 249, Society for Economic Dynamics

Percy D. F. (2015), “Strategy Selection and Outcome Prediction In Sport Using Dynamic Learning For Stochastic Processes”, Journal Of The Operational Research Society (2015) 66, 1840–1849

Pfeiffer M., Zhang H. & Hohmann A. (2010), “A Markov Chain Model Of Elite Table Tennis Competition”, International Journal Of Sports Science & Coaching Vol. 5, No. 2

Romer D. (2006), “Do Firms Maximize? Evidence From Professional Football”, Journal Of Political Economy, 2006, Vol. 114, No. 2, By The University Of Chicago

Ryvkin D. (2021), “To Fight Or To Give Up? Dynamic Contests With A Deadline”, Available At SSRN: https://doi.org/10.1287/mnsc.2021.4206

Seregina T. V., Ivashko A. A. & Mazalov V. V. (2019), “Optimal Stopping Strategies In The Game “The Price Is Right””, Proceedings Of The Steklov Institute Of Mathematics, 2019, Vol. 307, Suppl. 1, pp. S127–S141.

Strack P. & Viefers P. (2013, 2018), “Too Proud To Stop: Regret In Dynamic Decisions”, Available At SSRN: https://ssrn.com/abstract=2465840

Tsetlin I., Gaba A. & Winkler R. L. (2003), “Strategic Choice Of Variability In Multiround Contests And Contests With Handicaps”, INSEAD, Working Paper Series, Faculty & Research

D’Agostino F. (2022), “The Optimal Stopping Strategy in Dynamic Contests: The Case of F1 Races”, ICEF, Higher School of Economics, Term Paper

Appendix

Code

```
library(dplyr)
library(sn)
library(car)
library(ggplot2)
library(PerformanceAnalytics)



full_model <- lm(position ~ driverId + constructorId + grid + fastestLap + rank + stop + lap +
duration, data = dataset)
summary(full_model)

full_model <- lm(lap ~ driverId + constructorId + grid + rank , data = dataset)
summary(full_model)

plot(full_model[["residuals"]])
qqPlot(full_model[["residuals"]])
plot(selm(full_model[["residuals"]]~1, family="SN"))
chart.TimeSeries(full_model[["residuals"]])
chart.Histogram(full_model[["residuals"]], methods = "add.normal")

plot(residuals(full_model))

layout(matrix(c(1,2,3,4),2,2))
plot(full_model)
```

Plots

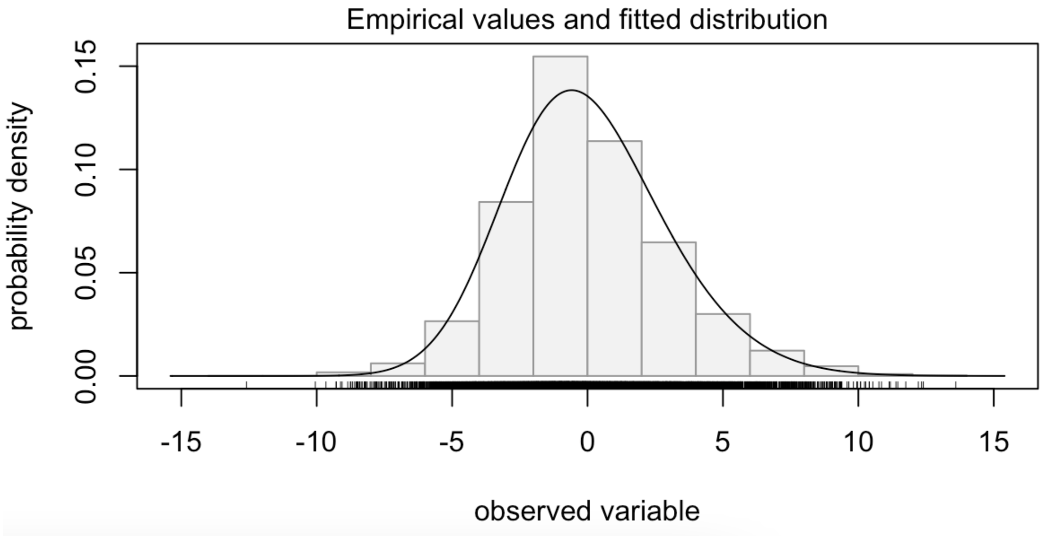

Figure 10 Y=Position Residuals Analysis

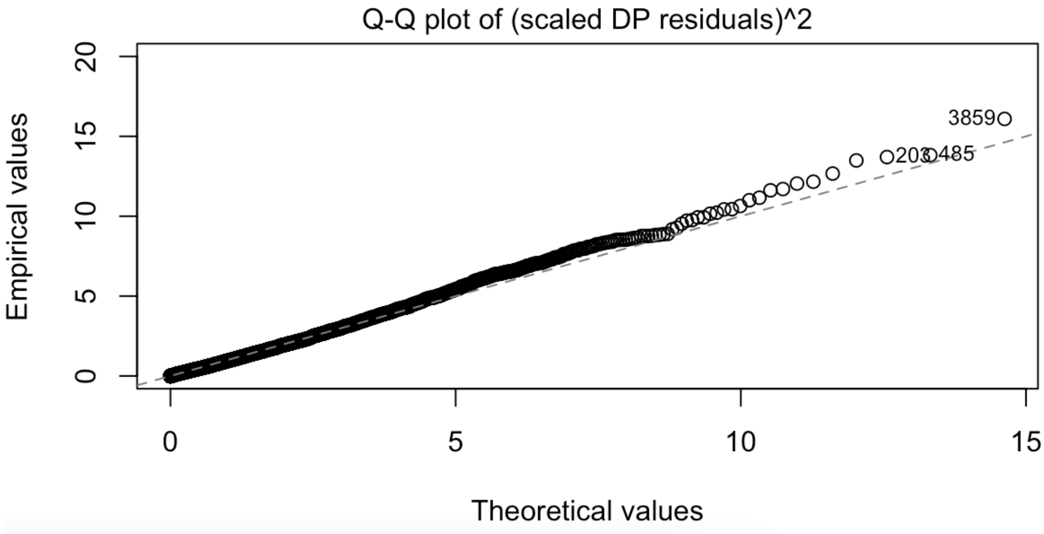

Figure 11 Y=Position Residuals Analysis

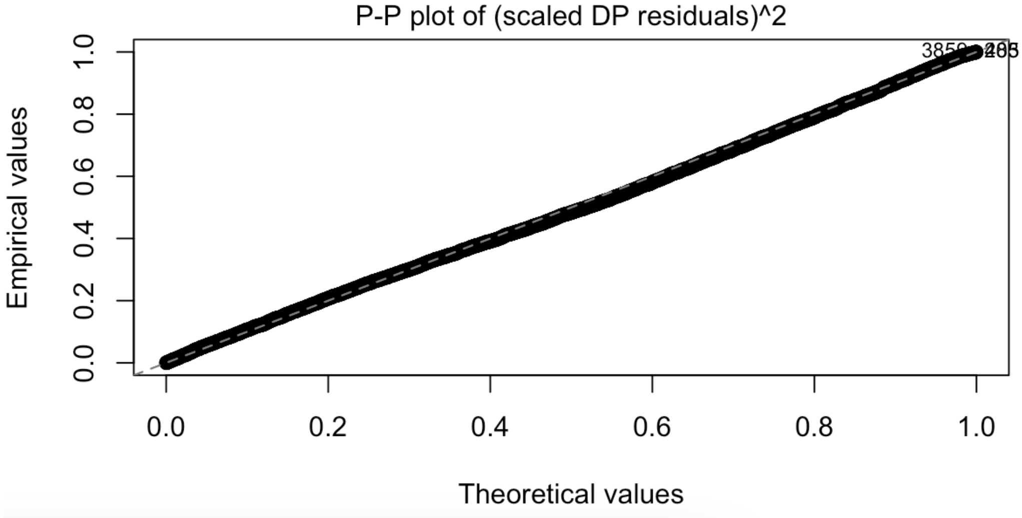

Figure 12 Y=Position Residuals Analysis

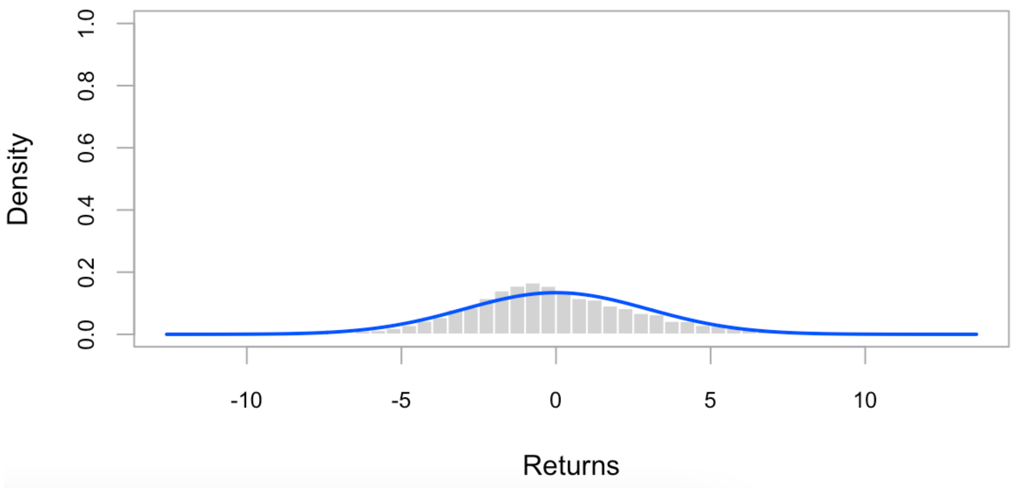

Figure 13 Y=Position Residuals Analysis

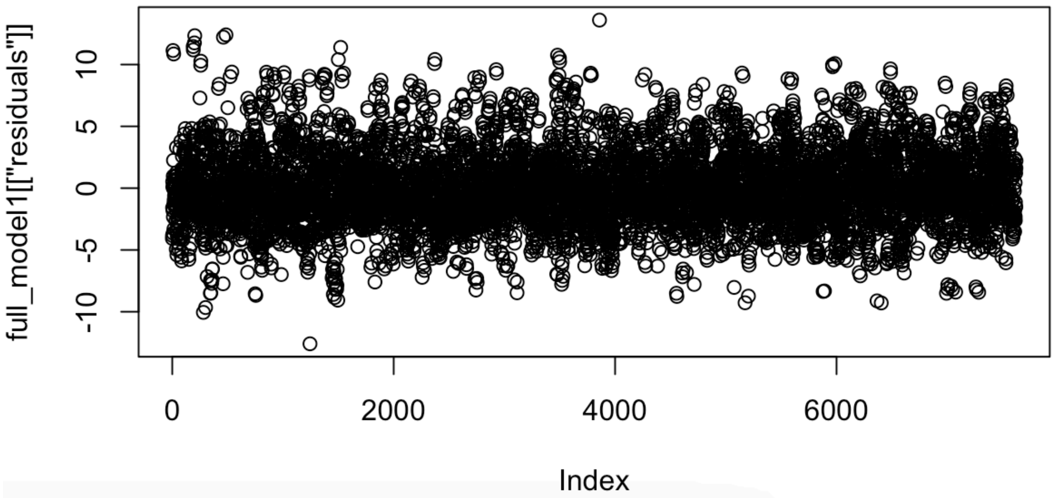

Figure 14 Y=Position Residuals Analysis

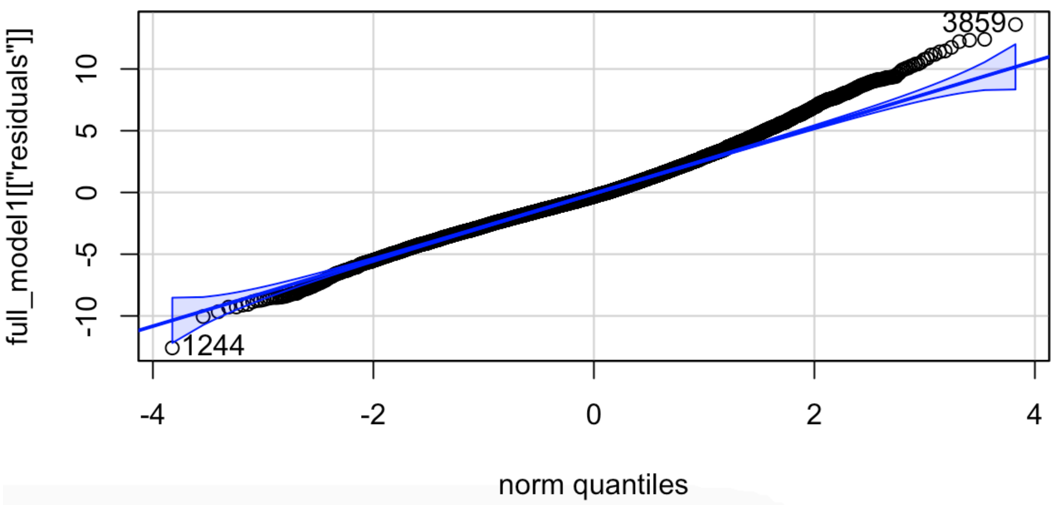

Figure 15 Y=Position Residuals Analysis

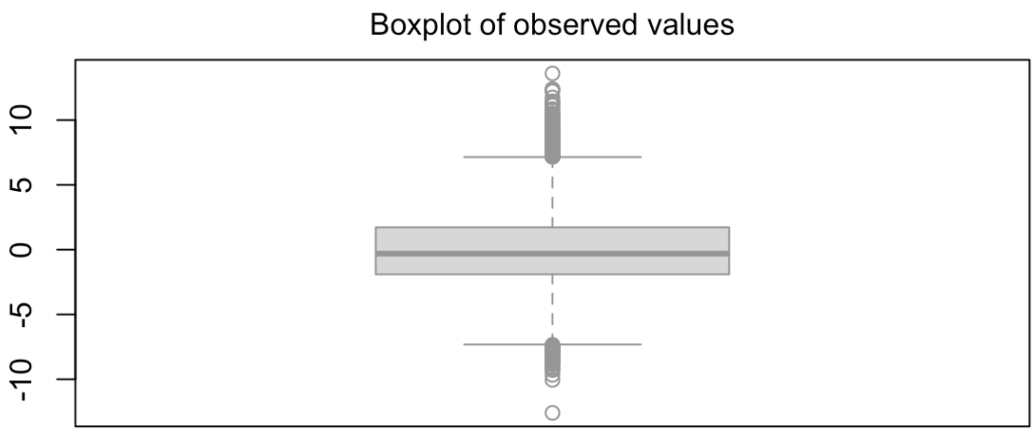

Figure 16 Y=Position Residuals Analysis

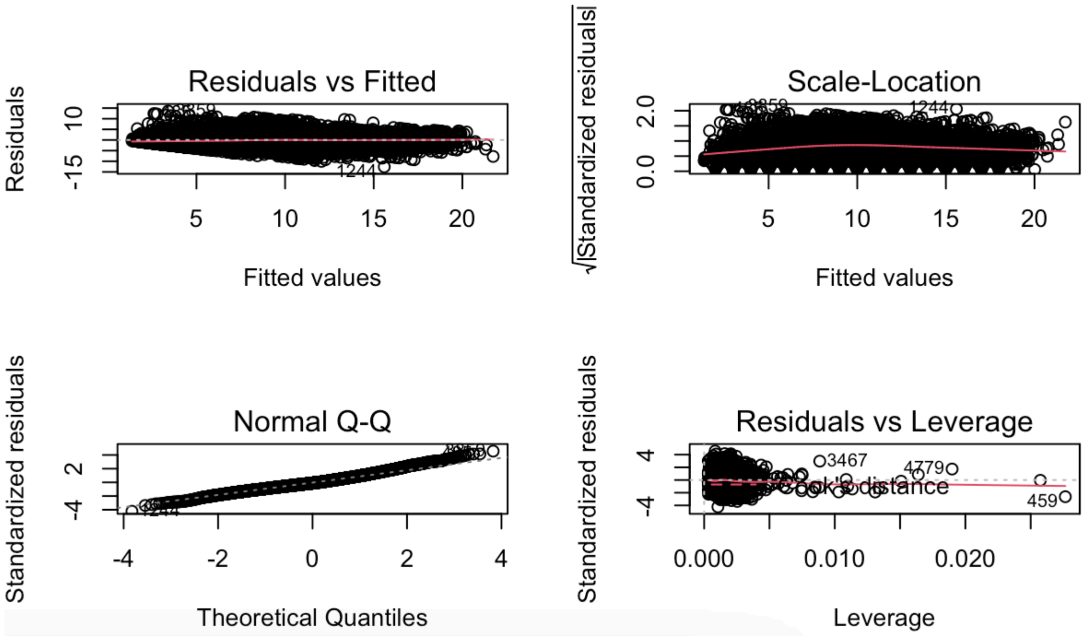

Figure 17 Y=Position Residuals Analysis

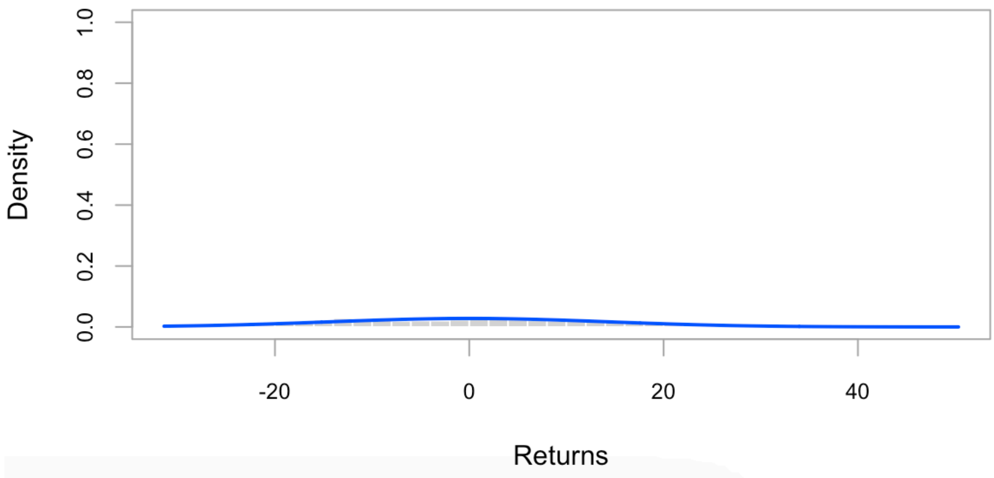

Figure 18 Y=Lap Residuals Analysis

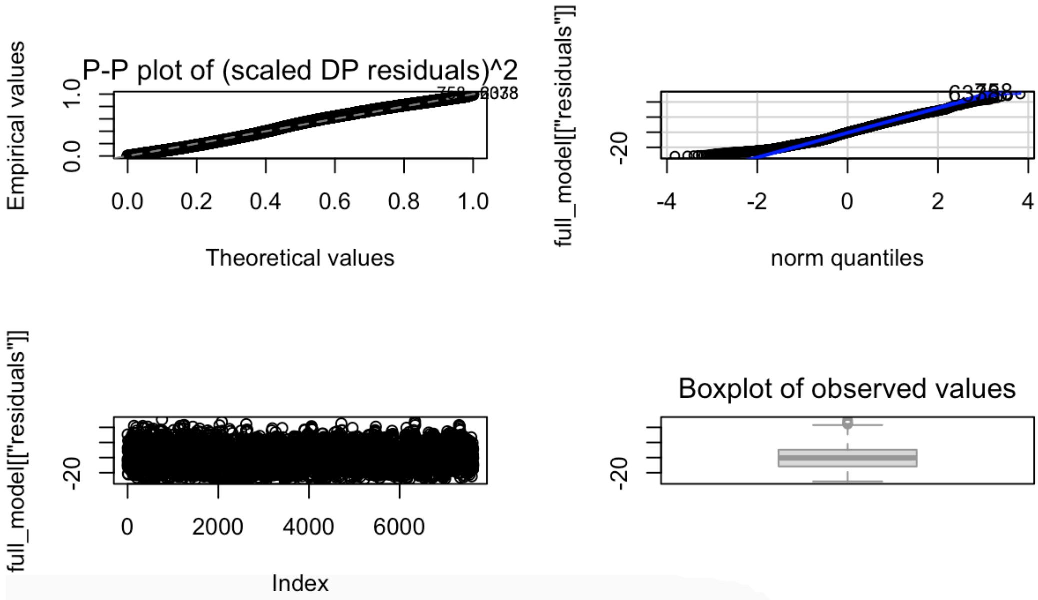

Figure 19 Y=Lap Residuals Analysis

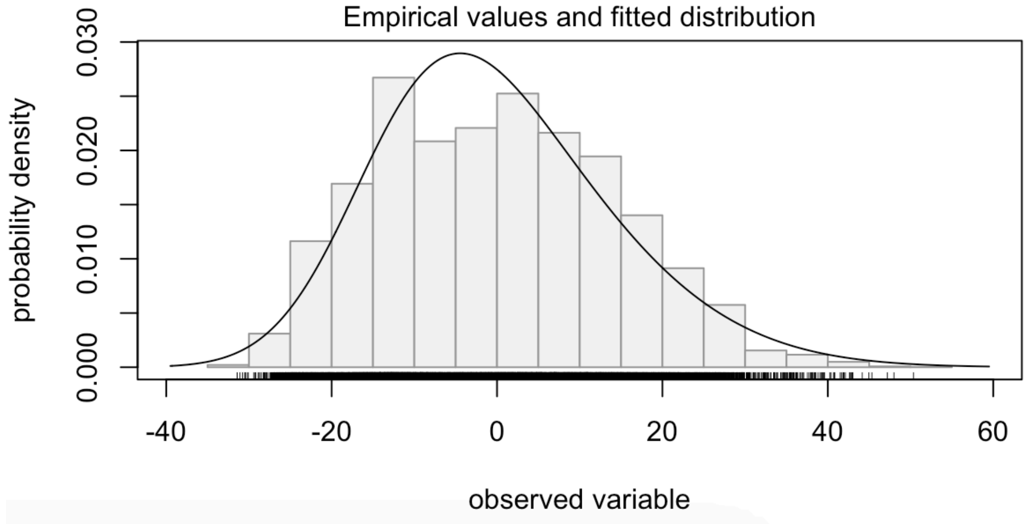

Figure 20 Y=Lap Residuals Analysis

Summary

Economic and financial models show agents that maximize their objective functions while interacting in their dynamic and competitive contexts. They establish their goal and identify obstacles that become constraints in a modelling environment. Competition in contests forces agents also to study others' characteristics and status, and sports are one of the most important examples of competition, so they are also important for general audience since they are an analysis of choices and strategies in a dynamic context with constraints.

They perfectly recreate a challenging environment, just like an economic sector with two or more firms competing or a financial market with two or more investors, that may not be in competition, but still have an influence on others' choices; thus, sports can serve as a model for them and for several more.

A very competitive team sport is Formula 1, and modelling and optimization are some essential elements of it. Like most of the sports, first of all the contest is studied before happening, as well as tactics. In Formula 1, even if it seems like the driver is alone in the contest to win the race, there is an entire team that only studies racing strategies, since these can overturn the result. Teams find themselves battling against other nine teams. It is a very dynamic environment, since choices made by strategists must be taken considering others' choices and state. The most important strategies concern the start of the race, for example how to approach the first curves and where to attack or defend the opponent, the behavior between the two drivers of the same team and the pit-stop above all. Even though strategists study tactics before the race, most of them are constructed or modified during the event, because the environment is unpredictable and in constant change, and opponents' moves could impact own choices. So, strategists must be able to adapt their tactics to the new situation and the changing during the race using simulation, algorithms and human intuition overall.

The pit-stop strategy, that is the choice of the moment in which to stop, is the most thought-out choice because the driver's surviving instinct is out of its scope, unlike the approach in the first curves. Due to regulations, drivers are obligated to take at least one pit-stop during the entire race[3], changing their tyres' type. The stopping strategy in F1 is highly effective and a wrong choice could compromise the entire race. It must carefully consider own capacities, because in the immediate laps after the pit-stop

[3] This is the case of dry races. In wet (raining) races there is no obligation to stop, meaning that stopping strategies don't necessary show up.

the driver must be able to recover the time lost due to the stop, circuit's characteristics and it must especially consider others' choices and state. An earlier pit-stop could give the driver the opportunity of the "undercut", a typical F1 strategy in which first-mover advantage could arise by stopping before the direct opponent. The principle of the undercut is that, in a contest between two or more drivers, the one who stops before will have at least one lap (trip around the race circuit) with a fresher tyre that corresponds to more performance, compared to the other drivers that still have to stop. In this way he will gain an advantage in terms of time and this advantage will last until the opponents will stop to level off the "freshness" of tyres. Of course, stopping the driver early comes with some problems: the risk of putting him into the traffic of other cars and the disadvantage of having a more worn-out tyre for the rest of the race, meaning slower laps.
On the other side of the coin, a late pit-stop will put the driver himself at risk of others' undercut, while giving him a fresher and higher-performance tyre for the end of the race. Here, there is a clear trade-off between performance and strategy, two fundamental factors to be balanced with difficulty.

That said, a Formula 1 race can be seen as an optimal stopping problem, that is a problem of optimization concerned with the choice of the moment in which executing a determined action in a finite horizon of time, in order to maximize an expected reward or to minimize an expected cost. Since in Formula 1 there is finite time composed by N laps in which to execute the action of the pit-stop, it is a perfect and interesting environment to solve an optimal stopping problem.
The objective of this optimization is to maximize the driver's position at the end of the race with the optimal execution of the pit-stop. On the basis of this execution, the driver will find himself in a better or worse position, also in terms of distance from opponents, since he can gain or lose time because of the pit-stop.
It is possible to find similarities between the choice of the lap in which to stop and when an investor should buy or sell a stock, exercise an American option or when a firm should start selling a new product, etc.

My goals in this original work were to calculate for the first time which is the optimal stopping strategy in a F1 race, to understand how much the stopping strategy influences the probability of winning, which other factors have a higher influence on this probability and to compare the stopping in Formula 1 with the stopping in the execution of American options' contracts, whom I studied the pricing, too.
I used a dual approach. From one side, I investigated empirically what factors shape the choice of a stopping time and increase the probability of winning. This has been done through several

regressions, analyzing the choice for different drivers or circuits but also taking into consideration general results. From the other side, I constructed different formal models and looked for an equilibrium stopping time. The reason is that empirically it's not possible to find a general equilibrium stopping time unless circuits are individually studied. In fact, since different circuits come with different specifications, each one of them has a different equilibrium stopping time. So, only a formal model can specify the equilibrium stopping time, whereas an empirical test can describe the optimal stopping time in a specific circuit. Furthermore, I introduced competition slowly. Initially, the models were ideal, assuming that the driver is alone in the circuit and only takes into consideration own conditions. Then, I analyzed how competition and interaction with opponents' impact on the strategy and on the behavior.

Finally, I introduced the optimal stopping in the context of American options. I showed different approaches, but I analyzed the Markovian one, that is similar to my Formula 1 approach. I explained American options' payoff and gain function and characterized the value function that contributes to the choice of the optimal stopping time.
Then, similarities and differences between the optimal stopping in Formula 1 and American options had been shown.

This thesis examines the sport of F1 in depth, with more than 100.000 data and several regressions from different points of view. My work contributes to the literature on the topic since it is a completely original work in which optimal stopping is applied to the Formula 1 environment, both theoretically and empirically, for the first time. The similarities between Formula 1 and American options' models show the close linkage between the two contexts, and they are useful to give even more importance to the contribution that my thesis brings to the optimal stopping literature. My models can serve as proxy for a lot of backgrounds, with or without competition.

The results, that can be deepened and continued in further studies, confirm the expectations: the optimal stopping time is found by formal models if general conclusions are needed, and can be adapted race by race, circuit by circuit, according to specific characteristics; instead, it is found by the empirical model and historical data if one wants circuit-specific results.

So, three formal models were presented, one of them developed by Goldman and Rao (2011, 2014) in their work on the optimal stopping problem in the NBA.

In the first model, every lap the strategists keep an eye on a reservation threshold (a marginal opportunity of pit-stop lap), that is the continuation value of the race and compare it to the pit-stop (stopping) value in terms of expected points. They will execute the stop once the continuation value is less than or equal to the pit-stop (stopping) value. It is based on dynamic efficiency and marginality. The second model instead is based on minimizing the entire time of the driver's race by choosing the optimal stopping time according to the tyres' condition function only, assuming that this is the only factor influencing the entire race time and that the driver is racing alone around the circuit.
This leads to the third model, that introduces the competition and the interaction with the opponents. Here, a 2-by-2 instantaneous game is presented. Two drivers have two possible actions: to stop or to continue the race. One of them is the leader, the other is the runner-up. With numerical payoffs, the Nash equilibrium is found when the leader chooses to stop, and the runner-up decides to continue to race. In fact, the leader will anticipate the stop to protect himself from others' undercut strategies, whereas the runner-up will try to come up with something different by differentiating the strategy and not stopping: he will go long and stop later.
The first model is adaptable to every race; the second one is usable when the driver isn't threatened by opponents' attack and has a considerable advantage on chasers. The third model studies the interaction between drivers.
All of these models work and find confirmation in reality. The best solution would be to use them together and to mix them with empirical ones.

My empirical tests were multiple regressions in order to assess which factors influence the most the probability of winning, to study how much the stopping lap counted in historical results and how the stopping lap is actually chosen.
Furthermore, I found that factors of influence have different weights whether the analysis is related to the entire history or to a single circuit. Drivers and cars have a higher influence on the long-term and on the totality of the Championship, whereas they have less on specific circuits. In fact, the drivers and teams' dominance are more related to eras and years than to circuits, that find themselves several protagonists.

A relevant outcome was that the starting grid position and the championship rank are the most important factors related to winning in F1. Looking at races in the same circuit, a lot of variables lose importance, while the grid position and the rank stays on the top winning factors. Duration of pit-stop and number of pit-stops are both important in a macro and in a micro view, with the latter being particularly impactful when carrying on a circuit-specific analysis.

The stopping lap hasn't a statistical relevance neither in the general model nor in the circuit-related one. It was predictable, since generally it's impossible to find a precise lap number in which to stop: circuits have different total number of laps and characteristics. Furthermore, analyzing an individual circuit, every driver will want to stop at the optimal lap. Due to this, it loses importance; by the way, the optimal stopping lap still exists, and it is studied empirically.

Firstly, using it as the dependent variable of the regression, it was analyzed from a general point of view, in all the races. There, it wasn't possible to establish an optimal one, since it depends on several variables that must be adapted race-by-race, circuit-by-circuit. These variables are the starting grid position, the rank of the driver and the car.
Instead, when looking at individual circuit, an optimal stopping lap arises. As a matter of fact, in the restricted-Monza model, the regressors have no predictive capability: the optimal stopping lap is explained by the intercept only. It has a value of 24.5, meaning that in a normal race in Italy a driver should stop between lap 24 and 25. Of course, this optimal lap can be adjusted taking into consideration formal models, so accounting for tyres' degradation and competition, giving once again importance to a right mix of the two kind of models (theoretical and empirical) that are coherent and find confirmation between each other.

So, even if it loses importance as an impactful factor on the final winning since everyone pursue it, it is shown that the optimal stopping lap exists, found theoretically for general purposes and empirically for specific circuits, and the drivers can't move much away from it, or their race will be ruined. They should instead slightly modify it according to competition and own conditions.

The optimal stopping time in Formula 1 is as important as in a financial or an economical context. The action must be executed by taking into account the entire environment and conditions, by checking every moment what is worth more between continuing and stopping and by looking at historical data.

Important financial instruments for which optimal stopping is crucial are the American options. The reasoning behind their optimal stopping problem is very similar to the one behind my first model in Formula 1 and the problem can be modelled very likewise to the problem in Formula 1, even if the competition is not included in American options.
I found similarities in the use of the Markov framework, since drivers and investors observe states and payoffs, and in the usage of a threshold, that is the continuation value in F1, and the gain function

in American options. Both drivers and investors look at these values and are ready to execute the action once a condition is satisfied.

In the case of F1, the stopping value must get greater than or equal to the continuation value; in the case of the American options, the value function must be equal to the gain function.

One big difference instead, is in the usage of time: in Formula 1, it is passive. It is only a consequence of the previously-presented choice.

Instead in American options, time has an active role: it is the cause, not the consequence. It is chosen in order to satisfy the condition and maximize the payoff.

Another difference is the presence of competition: in Formula 1 it is central, whereas in American options it isn’t an element to keep into consideration.

Despite the differences, the models can be interchanged because they are based on the same concepts. Formula 1 can be used as a proxy for other contexts where time has an important value and impact on the final payoff and where optimal stopping problem arises. With adjustments related to the specifics of the surroundings and the characteristics, the models shown are much adaptable.