Working Paper
No 297

NES Working
Paper series

IMPLICIT CENTIPEDES

Dmitry Dagaev
Daniil Starikov
Gleb Vasiliev

May
2026

Implicit Centipedes

Dmitry Dagaev[1,2], Daniil Starikov[3,*], and Gleb Vasiliev[2]

[1]New Economic School
[2]HSE University, Laboratory of Sports Studies
[3]University of Chicago, Harris School of Public Policy
[*]Corresponding author. Email: dstarikov@uchicago.edu.
Address: 1307 E 60th St, Chicago, IL 60637, USA

Abstract

Sequential games are typically studied in environments where the game tree and payoffs are explicitly specified. In many real-world interactions, however, agents face implicitly formulated dynamic games: while the strategic structure and incentives are well defined, the underlying game form is not directly observed and must be inferred from experience. We study an implicitly formulated version of the centipede game and analyze a natural setting in which professional agents repeatedly engage in a sequential interaction with centipede-like incentives without observing the game tree. Using detailed intra-race data from Formula 1, we show that pit-stop timing decisions between closely competing drivers generate a centipede game once overtaking probabilities are taken into account. We estimate the payoff-relevant components of the game, compute the equilibrium of the resulting interaction, and compare theoretical predictions to observed behavior. We find that drivers respond strategically to each other: pit-stop timing shifts systematically when a nearby rival pits, even after controlling for tire degradation and race conditions. The observed responses broadly follow the equilibrium logic of the model. At the same time, we document substantial heterogeneity across teams: some teams consistently choose pit-stop timings close to the theoretical optimum, while others deviate substantially, and these patterns are associated with differences in success rates. By documenting equilibrium behavior in a complex sequential interaction outside the laboratory, the paper provides new evidence on how equilibrium predictions perform in natural dynamic environments.

Keywords: centipede game, equilibrium behavior, implicit strategic environments, field evidence, Formula 1

JEL classification: C72, C93, Z20

1 Introduction

Sequential strategic interactions arise in many areas of human activity, including political campaigns, international negotiations, economic competition, and sports tournaments. Subgame perfect Nash equilibrium (SPNE) is the standard solution concept for such interactions. However, a large body of experimental evidence documents systematic deviations from SPNE behavior. One of the canonical examples is the centipede game, in which backward induction predicts immediate termination, while laboratory subjects typically delay stopping even in relatively simple settings.

Most existing studies of the centipede game focus on environments in which the game tree and payoffs are explicitly specified. In many real-world interactions, however, agents face dynamic strategic situations whose structure resembles a centipede game but is not presented in an explicit tree form. While the strategic incentives and opponents are well defined, players must infer the effective game form from experience. We study behavior in such environments and refer to them as implicit centipede games.

Although the theoretical solutions of explicit and implicit games coincide, behavior in the two environments may differ substantially. On the one hand, sophisticated players may gain an advantage in implicit centipede games by correctly inferring the strategic structure of the interaction. On the other hand, if the implicit structure is sufficiently complex, even highly skilled players may fail to identify the optimal strategy and instead engage in behavior that differs substantially from equilibrium predictions.

A useful benchmark comes from other strategic environments in professional sports. Consider, for example, penalty kicks in football. Although players are unlikely to compute mixed-strategy equilibria explicitly before taking a shot, empirical evidence suggests that both forwards and goalkeepers behave remarkably close to Nash equilibrium in such situations (Chiappori, Levitt, and Groseclose, 2002, Palacios-Huerta, 2003). Related studies examine equilibrium behavior in other professional sports, including American football, baseball, and tennis (Anderson et al., 2025, Kovash and Levitt, 2009, Walker and Wooders, 2001). Interestingly, professional athletes also tend to play near equilibrium in laboratory versions of simple penalty-type games (Palacios-Huerta and Volij, 2008). At the same time, other studies emphasize an important difference between field and laboratory settings: whereas laboratory experiments often reveal substantial deviations from equilibrium play, professional athletes frequently behave much closer to equilibrium in their natural environments (Levitt, List, and Reiley, 2010).

The centipede game presents a more challenging environment than the penalty game because of its dynamic structure and the need to reason through a sequence of contingent decisions. This raises a natural question: are experienced agents able to approximate equilibrium behavior in implicit centipede games arising in real-world environments?

In this paper, we study this question using a natural setting from Formula 1 racing. When two closely competing drivers decide on each lap whether to make a pit stop or remain on track, their interaction generates incentives that resemble a centipede game. Pitting yields fresher tires and therefore a speed advantage. Moreover, the longer both drivers stay on track, the larger the potential advantage of being the first to pit. At the same time, stopping too early in pursuit of this advantage creates a new risk: the opponent may delay the stop and gain a tire advantage later in the race. As a result, each driver would like to pit slightly earlier than the rival, but postpone the stop as long as possible. This tension creates a sequential strategic interaction in which both drivers decide each lap whether to stop or pass the decision to the next lap, generating incentives analogous to those of a centipede game.

Using detailed telemetry data from Formula 1 races between 2018 and 2025, we estimate the payoff-relevant components of the strategic interaction between closely competing drivers. In particular, we estimate the probabilities of overtaking that arise under different pit-stop timing configurations and use these estimates to characterize the payoffs of the game. We then compute the equilibrium of the resulting centipede-like interaction and compare the theoretical predictions with observed pit-stop decisions in real races.

We find clear evidence that drivers respond strategically to each other's actions. Pit-stop timing shifts systematically when a nearby competitor makes a stop, even after controlling for tire degradation and race conditions. Moreover, the observed responses are broadly consistent with the behavior predicted by the model: drivers tend to adjust their stopping decisions in the direction implied by equilibrium best responses. At the same time, we document an asymmetry in these responses. When the pursuer pits after the leader – corresponding to an overcut attempt – the observed responses tend to deviate somewhat more from the equilibrium prediction.

We also document substantial heterogeneity across teams. Some teams consistently choose pit-stop timings close to the theoretical optimum, while others deviate significantly from equilibrium predictions. These differences in strategic behavior translate into differences in outcomes: teams whose pit-stop timing is closer to the theoretical optimum achieve higher success rates in overtaking their rivals. Importantly, once we control for team fixed effects, we find no systematic relationship between the optimality of pit-stop decisions and drivers' experience or performance. This pattern suggests that, in this implicit centipede environment, the ability to approximate equilibrium play is primarily a team-level rather than a driver-level phenomenon.

Our paper contributes to several strands of literature. First, we contribute to the literature on the centipede game and equilibrium behavior in sequential interactions. Most existing empirical evidence on the centipede game comes from laboratory experiments in which the game tree and payoffs are explicitly specified (Bornstein, Kugler, and Ziegelmeyer, 2004, Krockow, Colman, and Pulford, 2016, Levitt, List, and Reiley, 2010, McKelvey and Palfrey, 1992, Nagel and Tang, 1998). By contrast, we study an environment in which the strategic structure resembles a centipede game but is not explicitly presented to the players. Professional drivers repeatedly face this interaction in real races and must infer the effective game form from experience. By estimating the payoff-relevant components of the interaction and comparing theoretical predictions with observed behavior, we provide field evidence on equilibrium behavior in an implicit centipede game.

Second, our paper contributes to the broader literature that uses professional sports as a testing ground for game-theoretic predictions. Previous studies have documented that professional athletes often behave close to equilibrium in strategic environments such as penalty kicks and other competitive interactions (Anderson et al., 2025, Chiappori, Levitt, and Groseclose, 2002, Kovash and Levitt, 2009, Palacios-Huerta, 2003, Palacios-Huerta and Volij, 2008, Walker and Wooders, 2001). We extend this literature by studying a dynamic sequential game with a much richer strategic structure. Unlike static penalty games, pit-stop timing decisions involve a sequence of contingent choices and strategic responses, making them a natural environment for studying equilibrium behavior in complex dynamic interactions.

Third, we contribute to the growing literature on strategic decision-making in motorsports. Much of the existing work focuses on predicting optimal race strategies by minimizing lap times or simulating race dynamics while treating opponents' behavior as exogenous (Bekker and Lotz, 2009, Heilmeier, Graf, and Lienkamp, 2018, Heilmeier et al., 2020a,b, Heine and Thraves, 2023, Tulabandhula and Rudin, 2014). In contrast, we explicitly model the interaction between competing drivers as a strategic game and evaluate whether real-world decisions approximate equilibrium behavior. This approach allows us to study the strategic outcomes of pit-stop duels rather than simply the time-minimizing race strategy, complementing recent game-theoretic approaches to motorsport strategy (Aguad and Thraves, 2024, Deck, Deck, and Zhu, 2014).

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the theoretical model of pit-stop interactions. Because the analytical solution of the game is cumbersome, we evaluate the equilibrium numerically using overtaking probabilities estimated from Formula 1 data. Section 4 describes the telemetry dataset covering the 2018–2025 seasons. Section 5 estimates overtaking probabilities, which are incorporated into the game tree in Section 3. Section 6 derives theoretical predictions, and Section 7 compares them with observed strategic behavior. Section 8 concludes.

2 Literature Review

The centipede game was first introduced by Rosenthal (1981). In its standard formulation (see also Figure 1), two players, X and Y, play a $2T$-period sequential game. In each period $1, \ldots, 2T-1$, one of the players chooses between terminating the game and passing the move to the opponent in the next period. Player X moves first. In the final period $2T$, player Y chooses between terminating the game and opting for the status quo. Let x_n and y_n denote the payoffs of players X and Y, respectively, if the game ends in period n. The sequences x_{2n-1} and y_{2n} are increasing, so that both players have incentives to delay termination. It is typically also assumed that $x_{2n} < x_{2n-1}$ and $y_{2n-1} < y_{2n}$ for each $n = 1, \ldots, T$, meaning that terminating the game oneself in the next period yields a higher payoff than allowing the opponent to terminate it in the current

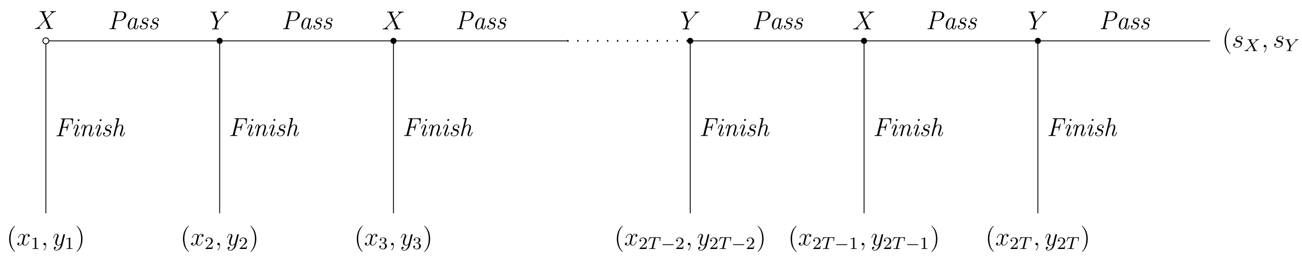

Figure 1: The centipede game.

period.

In the simplest variant of the game, where $s_X = s_Y = 0$, $x_1 > 0$, $y_2 > 0$, and $x_{2n} = y_{2n-1} = 0$ for each $n = 1, \dots, T$, players strictly prefer to terminate the game at any node they actually reach. Consequently, the subgame perfect Nash equilibrium (SPNE) predicts that player X terminates the game in the first period. More generally, the first player whose payoff exceeds the status quo prefers to terminate the game immediately before the opponent's payoff exceeds their own status-quo level.

A large experimental literature has examined whether individuals behave according to this equilibrium prediction. McKelvey and Palfrey (1992) show that depending on the parameters of the centipede game, only 1–15% of players stop the game immediately. They also document learning effects and faster stopping in high-stakes environments. Additional evidence of deviations from equilibrium behavior is reported in experiments with simultaneous-move versions of the centipede game, where only a small fraction of participants choose equilibrium strategies (Nagel and Tang, 1998). At the same time, some studies suggest that cognitive ability and strategic sophistication may affect equilibrium behavior. For example, Palacios-Huerta and Volij (2009) find that stronger chess players are better able to exploit suboptimal strategies when playing the centipede game against weaker opponents. Similarly, Bornstein, Kugler, and Ziegelmeyer (2004) show that groups tend to stop the game earlier than individuals and therefore behave closer to SPNE predictions. Levitt, List, and Reiley (2010) compare behavior in the centipede game with another sequential game, "Race to 100," which is primarily a game of skill, and question the interpretation that deviations from equilibrium necessarily reflect failures of backward induction.

Several theoretical explanations have been proposed for these deviations. One possibility is that the standard assumptions of rationality do not hold. Aumann (1995) shows that weakening the notion of rationality by assuming that players act rationally only in subgames that are actually reached may lead to outcomes different from SPNE. However, for the particular case of the centipede game the equilibrium outcome remains unchanged even under this weaker assumption (Aumann, 1998). Another explanation emphasizes bounded reasoning. Kawagoe and Takizawa (2012) analyze the centipede game using level-k reasoning and agent quantal response equilibrium models. They find that level-k reasoning performs better in predicting behavior in increasing-pie centipede games, whereas agent quantal response equilibrium better explains behavior in constantpie versions of the game. García-Pola, Iriberri, and Kovářík (2020) design an experiment aimed at distinguishing among several explanations of non-equilibrium behavior and conclude that deviations from SPNE are driven by heterogeneous reasoning processes, with violations of common knowledge of rationality and bounded reasoning abilities playing a particularly important role. A detailed overview of early experimental work on the centipede game is provided in Krockow, Colman, and Pulford (2016).

Sports competitions are frequently used as a testing ground for game-theoretic predictions because they provide high-stakes environments in which participants have strong incentives and substantial experience. In many sports, detailed data allow researchers to observe strategic interactions at a very fine level of granularity. Modern motorsport is particularly well suited for such analysis because high-frequency telemetry data record the timing of pit stops, positions of nearby competitors, and other race conditions. These features make it possible to study sequential strategic interactions between drivers competing for track position during a race.

Previous research on motorsport performance has largely focused on factors determined before the race, such as starting positions or historical driver performance (Mourão, 2017). Although such factors are clearly correlated with race outcomes, intra-race tactics and strategic decisions also play an important role. Several studies therefore attempt to characterize optimal race strategies. One approach is to construct predictive models that determine optimal strategies for stints – segments of the race between pit stops – or even at the lap level (Aguad and Thraves, 2024, Bekker and Lotz, 2009, Heilmeier, Graf, and Lienkamp, 2018, Heilmeier et al., 2020a,b, Heine and Thraves, 2023, Tulabandhula and Rudin, 2014). These approaches typically determine strategies that minimize lap times by adjusting baseline performance according to tire degradation, fuel load, and other factors. However, the behavior of opponents is usually treated as exogenous, which limits the ability of such models to analyze strategic interactions between competing drivers.

More recent work incorporates explicit game-theoretic frameworks to capture the strategic nature of close racing battles. For example, Deck, Deck, and Zhu (2014) analyze pit-stop decisions in NASCAR races during caution periods and develop a sequential model in which drivers decide whether to pit or stay on track. While drivers often follow optimal strategies, they also frequently rely on simple heuristics such as copying the leaders' decisions. Aguad and Thraves (2024), building on Heine and Thraves (2023), model driver interactions as a two-driver feedback Stackelberg game in which the leader moves first and the pursuer responds. In this framework, drivers choose tire compounds before the race and then make sequential decisions to maximize either their time advantage or their probability of winning. This papers demonstrates that the presence of a nearby rival can substantially affect optimal strategy and that drivers' decisions depend on the actions of their competitors. However, the parameters of this model are calibrated rather than estimated from race data, and the empirical analysis does not examine whether real-world drivers actually behave in accordance with the model's equilibrium predictions.

In this paper, we take a different approach by focusing explicitly on the strategic interaction between closely competing drivers deciding when to make a pit stop relative to each other. We show that this interaction can be modeled as a sequential two-player centipede-like game. Using a granular telemetry dataset, we estimate the payoff-relevant components of this interaction and compute the equilibrium of the resulting game. Unlike most previous studies, which focus on minimizing lap times while treating competitors' actions as exogenous, we analyze the strategic outcomes of pit-stop duels between drivers and compare the model's equilibrium predictions with observed behavior in real races.

3 The Model

We consider two drivers who, without loss of generality, start the race in close proximity to each other. The driver in front is the leader (L) and the one behind is the pursuer (P). The race consists of n laps. Each driver must make exactly one pit stop, taken at the end of one of the laps $1, \ldots, n-1$.[1] Thus, each driver has $n-1$ potential pit-stop opportunities.

Suppose no driver has pitted before the end of lap i, where $1 \leq i \leq n-1$. At the end of lap i, the leader first decides whether to stop. The pursuer observes this choice and then makes his own decision. These decisions determine the state of the race on lap $i+1$. Four situations may arise:

1. If only the leader stops on lap i, the pursuer subsequently chooses the lap of his own stop.
2. If only the pursuer stops on lap i, the leader later chooses when to commit his stop.
3. If both drivers stop on lap i, the game ends immediately and no further decisions are made.
4. If neither driver stops on lap i, the same sequential structure repeats on lap $i+1$, except that both drivers now have tires that are one lap older and one fewer lap remains in the race.

Payoffs reflect the probability of finishing ahead of the opponent. Let P_L and P_P denote the leader's and pursuer's probabilities of winning, with $P_L + P_P = 1$. Let l_L and l_P be the laps on which the leader and pursuer stop, respectively. The sequential strategies described above determine (l_L, l_P), but the converse is not true: a strategy includes counterfactual behavior off the equilibrium path. We now specify how payoffs depend on (l_L, l_P).

Case 1: $l_L < l_P$. In this case the leader pits first and the pursuer postpones his stop, thereby attempting an *overcut*. By staying out longer, the pursuer obtains fresher tires for the final stint, which may increase the chance of overtaking the leader near the end of the race.

[1] These simplifying assumptions closely reflect common features of modern Formula 1 strategy. While drivers may make more than one stop, only a single stop is required by regulation, and a large share of races in our sample are effectively one-stop events. For instance, in 78 out of 168 races in our sample, the median number of stops among finishing drivers is one. Restricting the sample to races unaffected by rain, red flags, or safety cars yields an even higher share: 51 out of 83 races are one-stoppers. Drivers can also choose among different tire compounds, but competitors running in close proximity typically follow similar compound strategies. Modeling all twenty cars would greatly expand the state space without changing the core strategic trade-off. Focusing on a two-driver, one-stop interaction therefore captures the essential structure of the problem while keeping the model analytically tractable.

Let $P_1^o(l_L, l_P)$ denote the probability that the pursuer emerges ahead immediately after completing his pit stop on lap l_P. This probability should be relatively small, as the leader enjoys a tire advantage in the interval between the two stops, but it may be positive since the pursuer benefits from clean air and may gain enough time to complete the overtake.

Conditional on whether this initial overtake succeeds, the pursuer may or may not retain the lead until the finish. Let $P_2^o(l_L, l_P \mid O = 1)$ be the probability that the pursuer wins the race given that he was ahead when rejoining from his stop, and let $P_2^o(l_L, l_P \mid O = 0)$ be the probability of eventually winning if the overcut attempt fails. The pursuer's total payoff is therefore

$$P_P^o(l_L, l_P) = P_1^o(l_L, l_P) \cdot P_2^o(l_L, l_P \mid O = 1) + \big(1 - P_1^o(l_L, l_P)\big) \cdot P_2^o(l_L, l_P \mid O = 0), \tag{1}$$

and the leader's payoff is $P_L^o(l_L, l_P) = 1 - P_P^o(l_L, l_P)$.

Case 2: $l_L > l_P$. Here the pursuer pits first and attempts an *undercut*. By stopping earlier, the pursuer acquires fresher tires and may pass the leader before the latter commits his stop. Let $P_1^u(l_L, l_P)$ denote the probability that the pursuer is ahead when the leader exits from his stop on lap l_L. This probability is strictly positive because the tire advantage increases the pursuer's pace between laps l_P and l_L.

If the pursuer succeeds in overtaking before the leader's stop, the leader gains the subsequent tire advantage, which affects the likelihood of maintaining the lead until the end. Let $P_2^u(l_L, l_P \mid O = 1)$ be the probability of the pursuer winning conditional on a successful undercut, and let $P_2^u(l_L, l_P \mid O = 0)$ be the probability of winning if the undercut fails. The pursuer's payoff is

$$P_P^u(l_L, l_P) = P_1^u(l_L, l_P) \cdot P_2^u(l_L, l_P \mid O = 1) + \big(1 - P_1^u(l_L, l_P)\big) \cdot P_2^u(l_L, l_P \mid O = 0), \tag{2}$$

with the leader's payoff given by $P_L^u(l_L, l_P) = 1 - P_P^u(l_L, l_P)$.

Case 3: $l_L = l_P$. When both drivers stop on the same lap, neither gains a tire advantage in any phase of the race. Overtaking becomes more difficult for the pursuer, but remains possible due to potential pit-crew errors, on-track mistakes, or small performance differences. Let $P_P^s(l_L, l_P)$ denote the pursuer's probability of winning in this simultaneous-stop case; the leader's payoff is $P_L^s(l_L, l_P) = 1 - P_P^s(l_L, l_P)$.

We now outline the qualitative expectations for how the probabilities introduced above depend on the pit-stop laps (l_L, l_P). In the next sections, we examine whether these patterns are borne out in the data.

1. $P_1^u(l_L, l_P)$

   This is the probability that the pursuer successfully completes an undercut before the leader's stop. It should increase in both the magnitude of the tire advantage at l_L and the length of the interval between the stops, that is, in $l_L - l_P$.

2. $P_2^u(l_L, l_P \mid O = 1)$

   Conditional on a successful undercut, the leader subsequently has the tire advantage. The pursuer is therefore more likely to lose the position when many laps remain. Thus, this probability should decrease in $n - l_L$ and in the tire-age gap $l_L - l_P$.

3. $P_2^u(l_L, l_P \mid O = 0)$

   If the undercut fails, the pursuer remains behind and has the older compound. In this case the likelihood of overtaking is extremely small, so this probability should be close to zero.

4. $P_1^o(l_L, l_P)$

   When attempting an overcut, the pursuer runs on an older compound between the two stops. As a result, overtaking in this phase should be unlikely, so this probability should be low. If there is any systematic dependence on the state variables, it should be decreasing in both the length of the between-stop interval, $l_L - l_P$, and in the magnitude of the pursuer's tire disadvantage, l_L.

5. $P_2^o(l_L, l_P \mid O = 1)$

   If the overcut succeeds, the pursuer enjoys a tire advantage and holds track position. Losing the lead becomes unlikely, so this probability should be close to one.

6. $P_2^o(l_L, l_P \mid O = 0)$

   If the overcut fails, the pursuer re-emerges behind but with fresher tires. The likelihood of overtaking should increase with the tire-age advantage $l_P - l_L$ and with the remaining distance $n - l_P$.

Taken together, these monotonicity assumptions describe how relative pit-stop timing maps into overtaking probabilities and thereby into payoffs. In Section 5, we estimate these objects from Formula 1 data, and in Section 6, we examine the equilibrium of the game using the estimated probabilities.

4 Data

This section describes the dataset used in the empirical analysis. We first outline the sources of the telemetry data and the structure of the raw information. We then explain how the main variables are constructed and introduce the notation that links the data to the theoretical model. Next, we describe the sample restrictions imposed to ensure comparability of strategic situations and consistency with the model assumptions. Finally, we present the resulting datasets and provide descriptive statistics that summarize the key empirical patterns motivating the subsequent analysis.

4.1 Data collection

We assemble official Formula 1 telemetry data for all Grand Prix races from 2018 to 2025, excluding the 2018 Australian Grand Prix, the 2018 Bahrain Grand Prix, and the 2020 Austrian Grand Prix.[2] The raw data are recorded at the driver-lap level and include, for each lap and driver, the compound in use, the identity and distance of the nearest car ahead, a pit-stop indicator, and information on track and session conditions. Our primary interest lies in all pit stops across these races, the positions of the closest cars ahead and behind at the moment of the stop, the strategic situation surrounding each stop, and the age of the tires being replaced.

We reorganize the data so that a single observation corresponds to a pair of neighboring drivers who interact strategically through the timing of their pit stops. Whenever two drivers have made the same number of stops up to a given lap, a stop by the pursuing driver before the leading driver constitutes an *undercut*; a stop by the leader first constitutes an *overcut*; and simultaneous stops form a *same-lap stop*. The strategic interaction between the pair ends either when one of the drivers makes another stop or when the race finishes. Each observation therefore maps to one of three outcomes: undercut, overcut, or same-lap stop.

For each interaction we record the lap of each driver's stop, three indicators (who led before the first stop, who led after both stops, and who led at the end of the interaction), tire ages and compounds at the moment of each stop, and the resulting tire-age advantages. We also observe subsequent pit stops, the inter-driver distance on the lap preceding the stop, the local track and traffic situation, and the pair's relative positions for the remainder of the race.

Finally, we merge in information on Constructors' Championship standings.[3] This allows us to compute team-strength differentials based on the final standings of each season. In addition, we collect a set of driver-level characteristics measured at the beginning of each season, including career totals of races, points, podiums, race wins, and championships. These variables are used to study whether and how strategic pit-stop decisions vary systematically with drivers' experience and past performance.

4.2 Construction of new variables

We use the telemetry data to construct several variables that summarize the strategic environment faced by each pair of drivers. Throughout, let n denote the total number of laps in the race, and let l_L and l_P be the laps on which the leader and pursuer make their pit stops. Let l'_L and l'_P denote the next pit stops of the leader and pursuer after (l_L, l_P) if they occur.

[2] Data are retrieved using the *fastf1* Python package. These three races are omitted because official telemetry is unavailable. We additionally exclude sprint races, which are shorter and do not require pit stops, and therefore do not generate comparable strategic interactions.

[3] Each Formula 1 team is represented by two cars; points scored by both drivers contribute to the Constructors' Championship.

Share of race between the drivers' stops. We compute the share of the race that passes between the first and the second stop within the pair. Define

$$\Delta\text{Laps (between the stops)} = \frac{|l_L - l_P|}{n}.$$

For undercuts ($l_P < l_L$) this equals $(l_L - l_P)/n$; for overcuts ($l_L < l_P$), $(l_P - l_L)/n$; and for same-lap stops it is zero by construction. This variable proxies the time during which one driver operates on fresher tires before the opponent completes his stop.

Remaining share of the race after both stops. Next, we measure the time available for on-track battles after both drivers have completed the focal stops. If neither driver pits again, the remaining race share is

$$\Delta\text{Laps (after the stops)} = \frac{n - \max\{l_L, l_P\}}{n}.$$

If at least one driver makes additional stops, we take the earliest such stop, $\min\{l'_L, l'_P\}$, as the end of the interaction:

$$\Delta\text{Laps (after the stops)} = \frac{\min\{l'_L, l'_P\} - \max\{l_L, l_P\}}{n}.$$

This ensures that the interaction window reflects the time during which the pair can meaningfully compete on track before the next pit stop alters strategies.

Tire-age advantages. For each driver $d \in \{L, P\}$, let $a_d(t)$ denote the age (in laps) of the tire compound used on lap t. We compute two tire-age differentials:

$$\Delta\text{Tire (between)} = \frac{|a_L(l_{\min}+1) - a_P(l_{\min}+1)|}{n}, \qquad \Delta\text{Tire (after)} = \frac{|a_L(l_{\max}+1) - a_P(l_{\max}+1)|}{n},$$

where $l_{\min} = \min\{l_L, l_P\}$ and $l_{\max} = \max\{l_L, l_P\}$. The first captures the tire advantage during the interval between the stops; the second captures the advantage after both drivers have pitted.

Effectiveness measures. To evaluate strategic outcomes, we record which driver is ahead (i) immediately after the second stop in the pair and (ii) at the end of the interaction (the end of the race or the next pit stop of either driver). We define a pursuer-success indicator for overtake:

$$\text{Overtake} = \begin{cases} 1, & \text{if the pursuer is ahead at the specified moment,} \\ 0, & \text{otherwise.} \end{cases}$$

This variable is used to measure the effectiveness of undercuts, overcuts, and same-lap strategies.

Stints and pit-stop counts. For each driver, we compute the total number of pit stops made during the race and the ordinal number of the considered stop within that sequence. We refer to this ordinal number as the *stint*, following standard motorsport terminology.

4.3 Data cleaning

To estimate the probabilities defined in Section 3, we impose a series of sample restrictions to remove observations in which pit-stop decisions or their consequences deviate from the strategic environment captured by our model.

We begin by excluding all races affected by rain or red-flag interruptions.[4] Red flags and heavy rain are infrequent, but they dramatically alter optimal strategy, whereas our theoretical framework does not incorporate such events.

We also remove all observations in which at least one driver makes a stop under yellow flags, Safety Car, or Virtual Safety Car conditions.[5] Because SC/VSC periods reduce the time cost of pitting, they distort the undercut/overcut incentives we aim to measure.

Next, we drop all pairs in which at least one driver fails to finish or finishes more than one lap behind the leader. Such outcomes typically reflect mechanical issues or severe driver errors rather than strategic interactions.

To ensure that we analyze only situations in which drivers are sufficiently close to each other, we retain only those observations in which the inter-driver distance on the lap preceding the first stop is below a circuit-specific threshold. To construct this threshold, we collect qualifying data from all circuits in 2018–2025, compute the top average speed, and calculate the distance covered in three seconds at that speed. We use this "three-second" distance as the cutoff for defining close strategic interactions. Because race pace is slower than qualifying pace, the threshold corresponds to slightly more than three seconds in race conditions; we account for this in robustness checks using alternative thresholds.

We further restrict the sample to pairs of drivers who follow strategies with the same total number of pit stops and the same number of remaining stops at the moment of interaction. This ensures that the strategic options available to the two drivers are comparable.

Finally, we exclude pit stops made in the first or last 5% of the race. Very early stops often reflect technical issues, while very late stops are frequently made to set the fastest lap of the race rather than to influence position, and therefore do not correspond to the strategic behavior modeled in Section 3.

[4]A red flag is shown when track conditions become unsafe or when debris cannot be cleared quickly. The race is suspended and later restarted with all time gaps eliminated. Under current regulations, drivers may change tires during the red-flag break, and such a tire change counts as the mandatory stop. These features make red-flag races incompatible with the strategic structure of our model.

[5]A Safety Car (SC) or Virtual Safety Car (VSC) is deployed when conditions require drivers to slow down but are expected to be resolved quickly. Under SC/VSC, pit stops are much less costly in terms of time lost, sharply reducing the strategic role of tire-age advantages.

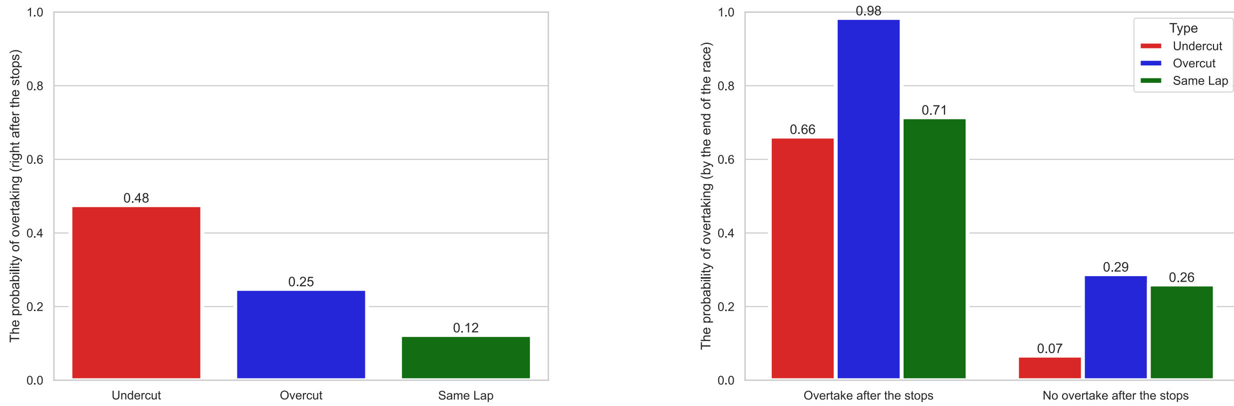

(a) Overtaking probabilities after the stops (b) Overtaking probabilities by the end of the race

Notes: This figure reports unconditional overtaking probabilities for undercuts, overcuts, and same-lap pit stops. Panel (a) reports the probability that the pursuer overtakes the leader immediately after both drivers have completed their pit stops. Panel (b) reports the probability that the pursuer finishes ahead by the end of the race, conditional on whether an overtake occurred immediately after the stops. The sample is restricted to pit stops that satisfy the data-cleaning procedure described in Section 4.3 and are the only stops within the drivers' strategies.

Figure 2: Overtaking Probabilities by Pit-Stop Strategy

4.4 Descriptive statistics

After assembling the raw telemetry, constructing additional variables, and applying the restrictions described in Subsection 4.3, we create three datasets for use in our empirical analysis. Each has distinct advantages.

Our primary dataset consists of pit stops that are the *only* stops in each driver's strategy. This sample most closely matches the theoretical environment in Section 3, where each driver pits exactly once. Its limitation is size: relying solely on single-stop strategies yields relatively few observations and may reduce statistical precision.

To ensure robustness, we therefore employ two complementary datasets. First, we use the full set of pit stops remaining after data cleaning. This maximizes statistical power, but introduces a potential concern: intermediate stops in multi-stop strategies may not reflect genuine attempts to overtake the opponent before the next scheduled stop, potentially biasing estimated strategic responses. Second, we construct a sample containing only the *last* stops in each driver's strategy. These stops necessarily precede the final on-track phase of the race for the pair, ensuring that drivers have strong incentives to overtake before the interaction ends. This sample is substantially larger than the one-stop dataset while avoiding the interpretational issues associated with intermediate strategic phases.

Table 1 reports summary statistics for our main dataset (single-stop strategies). The first two rows of the table, together with Figure 2, summarize the empirical frequencies of overtakes immediately after the stop and by the end of the race. Undercuts exhibit the highest immediate

Table 1: Descriptive Statistics by Pit-Stop Interaction Type

	Undercut		***Overcut***		***Same lap***	
	Mean	SD	Mean	SD	Mean	SD
$\mathbb{I}$***(Overtake):***						
– right after the stops	0.476	0.501	0.248	0.433	0.123	0.331
– in the end of the stint/race	0.350	0.479	0.461	0.499	0.316	0.469
Δ***Laps:***						
– between the stops	0.155	0.161	0.172	0.147	0.000	0.000
– after the stops	0.505	0.165	0.480	0.151	0.702	0.136
Δ***Tire:***						
– between the stops	0.350	0.132	0.355	0.125	0.006	0.016
– after the stops	0.151	0.157	0.170	0.144	0.006	0.016
Other characteristics:						
– Average distance (meters)	99.775	44.670	84.319	46.280	107.898	49.433
– ΔTeam's strength	−0.364	2.676	0.291	2.487	−0.088	2.165
Observations	143		254		57	

Notes: This table reports descriptive statistics for key race-state variables and outcomes, separately for undercuts, overcuts, and same-lap pit stops. The sample consists of driver pairs involved in close pit-stop interactions within one-stop strategies, conditional on both drivers finishing the race and satisfying the data-cleaning procedure described in Section 4.3. The unit of observation is a pair of drivers involved in a close pit-stop interaction within a one-stop strategy. An *undercut* occurs when the pursuer pits before the leader, an *overcut* when the pursuer pits after the leader, and *same-lap* when both drivers pit on the same lap. $\mathbb{I}$(Overtake) indicates whether the pursuer overtakes the leader either immediately after the pit stops or by the end of the stint (or race). ΔLaps measures the normalized difference in laps either between pit stops or after both drivers have stopped, while ΔTire captures the corresponding normalized tire age advantage. Additional characteristics include the average on-track distance between drivers, the number of pit stops, stint indicators, and the difference in team strength. Means and standard deviations are reported.

success rate (48%; panel (a)). Panel (b) shows that when a position swap occurs immediately after the stop, the probability of retaining the new position until the finish is extremely high for both overcuts (98%) and undercuts (66%). Conditional on no immediate swap, overcuts still yield higher subsequent overtaking rates (29%) than undercuts (7%) or same-lap stops (26%). While overcuts appear more effective overall, this pattern must be interpreted jointly with the race-state variables ΔLaps and ΔTire, since observed success rates reflect both strategic choices and the state of the race. A further observation is that both undercuts and overcuts outperform same-lap stops, consistent with the empirical fact that competing drivers rarely pit on the same lap. Finally, potential overcut situations occur more frequently than potential undercuts; the underlying reasons are explored in later sections.

Online Appendix Figure A1 presents distributions of pit-stop timing for leaders and pursuers in one- and two-stop races, scaled by race distance. In one-stop races (panels (a)–(b)), leaders tend to pit earlier, typically between 30–40% of race distance, while pursuers display a broader distribution with more late stops. This pattern reflects the prevalence of overcut attempts. Few pit stops occur outside the 10–80% race-distance window. In two-stop races (panels (c)–(d)), the number of second stops is smaller than the number of first stops; this arises from our requirement that drivers be sufficiently close for inclusion, a condition less frequently satisfied later in the race.

Using the pit-stop data, we construct the main state variables ΔLaps and ΔTire, whose distributions are shown in Table 1 and Online Appendix Figure A2. For both undercuts and overcuts, the distribution of the share of the race between stops (panel (a)) is right-skewed, with a peak near zero; the peak is more pronounced for undercuts, indicating that leaders often respond immediately to a pursuer's early stop. Because overcuts typically occur later in the race, the tire-age difference after the first stop is slightly larger for overcuts (panel (b)). Panel (c) shows that the remaining share of the race after both stops is roughly bell-shaped with a peak around 45%, implying ample remaining distance for on-track overtaking. Finally, the distribution of tire-age differences after the stops (panel (d)) mirrors the distribution of ΔLaps between stops.

We also report additional characteristics in Table 1: average distance between drivers on the lap before the stop and the difference in Constructors' Championship rankings between the teams. These variables display no meaningful differences between undercuts and overcuts. Descriptive statistics for the full-sample and last-stop subsamples are reported in Online Appendix Table A1. Panel A summarizes all stops satisfying the criteria of Subsection 4.3; Panel B reports statistics for last-stop observations only.

5 Estimation

This section links the theoretical framework to the data by estimating the key probability objects that determine payoffs in the model. Guided by the structure of the game in Section 3, we examine how the likelihood of an overtake depends on the length of the strategic window between pit stops and on the magnitude of temporary tire advantages. Our empirical strategy exploits within-race variation in pit-stop timing among closely competing drivers to identify the functional relationships underlying the probability components introduced in the theoretical model.

Our baseline specification is a random-effects linear probability model of the form

$$y_{ijsct} = \beta_0 + \beta_1 \Delta\text{Laps}_{ijsct} + \beta_2 \Delta\text{Tire}_{ijsct} + \gamma X_{ijsct} + \mu_c + \varepsilon_{ijsct}, \tag{3}$$

where y_{ijsct} is an indicator equal to one if the pursuer i overtakes the leader j by the end of a given race period associated with their s-th pit stops on circuit c in season t. The relevant period is either (i) the interval between the two drivers' pit stops or (ii) the interval following both pit stops until the end of the race (or the end of their interaction, as defined in Section 4).

The variable $\Delta\text{Laps}_{ijsct}$ measures the share of race distance covered during the considered period, while $\Delta\text{Tire}_{ijsct}$ captures the tire-age advantage between the two drivers over the same interval. The vector X_{ijsct} includes additional controls. In our baseline specification, it contains the difference in team strength; in extended specifications, it also includes season fixed effects and controls for differences in tire compounds. The term μ_c denotes a circuit-specific random effect, and ε_{ijsct} is an idiosyncratic error term clustered at circuit level. As robustness checks, we additionally estimate specifications with circuit fixed effects and specifications without circuit-level controls.

To allow for non-linear effects of tire advantage and the duration of this advantage, we also estimate a specification that includes quadratic and interaction terms:

$$y_{ijsct} = \beta_0 + \beta_1 \Delta\mathrm{L}_{ijsct} + \beta_2 \Delta\mathrm{T}_{ijsct} + \beta_3 \Delta\mathrm{L}^2_{ijsct} + \beta_4 \Delta\mathrm{T}^2_{ijsct} + \beta_5 \Delta\mathrm{L}_{ijsct} \times \Delta\mathrm{T}_{ijsct} + \gamma X_{ijsct} + \mu_c + \varepsilon_{ijsct}, \tag{4}$$

where all variables are defined as in equation (3). For notational convenience, we use $\Delta\mathrm{L}_{ijsct}$ and $\Delta\mathrm{T}_{ijsct}$ to denote $\Delta\mathrm{Laps}_{ijsct}$ and $\Delta\mathrm{Tire}_{ijsct}$, respectively.

We estimate the probabilities defined in Section 3 using different race periods and subsamples. The key probabilities are summarized below.

1. $P_1^u(l_L, l_P)$: The probability of overtaking by undercut immediately after the stops. We restrict the sample to potential undercuts. The outcome variable is a dummy for successful overtaking one lap after the leader's stop. $\Delta\mathrm{Laps}_{ijsct}$ measures the race share between the two drivers' stops, and $\Delta\mathrm{Tire}_{ijsct}$ is the tire advantage of the pursuer during this period.

2. $P_2^u(l_L, l_P|O = 1)$: The probability of overtaking by undercut by the end of the interaction, conditional on the pursuer being ahead immediately after both stops. The sample consists of potential undercuts where the pursuer overtakes the leader one lap after the leader's stop. The outcome variable is a dummy for successful overtaking by the end of the interaction. $\Delta\mathrm{Laps}_{ijsct}$ measures the race share between the leader's stop and the last lap of the interaction, and $\Delta\mathrm{Tire}_{ijsct}$ is the tire advantage of the leader.

3. $P_2^u(l_L, l_P|O = 0)$: The probability of overtaking by undercut by the end of the interaction, conditional on the pursuer being behind immediately after both stops. This sample includes potential undercuts where the leader remains ahead one lap after his stop. The outcome variable, $\Delta\mathrm{Laps}_{ijsct}$, and $\Delta\mathrm{Tire}_{ijsct}$ are the same as above.

4. $P_1^o(l_L, l_P)$: The probability of overtaking by overcut immediately after the stops. We restrict the sample to potential overcuts. The outcome variable is a dummy for successful overtaking one lap after the pursuer's stop. $\Delta\mathrm{Laps}_{ijsct}$ measures the race share between the two drivers' stops, and $\Delta\mathrm{Tire}_{ijsct}$ is the tire advantage of the leader.

5. $P_2^o(l_L, l_P|O = 1)$: The probability of overtaking by overcut by the end of the interaction, conditional on the pursuer being ahead after both stops. The sample consists of potential overcuts where the pursuer overtakes the leader one lap after his stop. The outcome variable

is a dummy for successful overtaking by the end of the interaction. $\Delta\text{Laps}_{ijsct}$ measures the race share between the pursuer's stop and the last lap of the interaction, and $\Delta\text{Tire}_{ijsct}$ is the tire advantage of the pursuer.

6. $P_2^o(l_L, l_P | O = 0)$: The probability of overtaking by overcut by the end of the interaction, conditional on the pursuer being behind after both stops. This sample includes potential overcuts where the leader remains ahead one lap after the pursuer's stop. The outcome variable, $\Delta\text{Laps}_{ijsct}$, and $\Delta\text{Tire}_{ijsct}$ are the same as above.

7. $P_2^s(l_L, l_P)$: The probability of overtaking by the end of the interaction when both drivers pit on the same lap. The sample consists of same-lap stops. The outcome variable is a dummy for successful overtaking by the end of the interaction. $\Delta\text{Laps}_{ijsct}$ measures the race share between the drivers' stops and the last lap of their interaction, and $\Delta\text{Tire}_{ijsct}$ is trivially zero since the drivers pit simultaneously and have the same tire age.

The estimation results are reported in Tables 2–8. All tables follow the same structure. Column (1) reports estimates from the linear specification in equation (3) without circuit-level effects. Column (2) adds random effects for circuits. Column (3) replaces random effects with circuit fixed effects and includes additional controls for seasons and tire compounds. Columns (4)–(6) repeat these three specifications for the quadratic model in equation (4): without circuit effects, with circuit random effects, and with circuit fixed effects and additional controls, respectively.

Our preferred specification is column (5). It allows for flexible, non-linear effects of tire advantage and race length while accounting for within-circuit correlation through random effects. This approach balances flexibility and parsimony: unlike the linear models, it captures key non-linearities implied by the theory, and unlike fixed-effects specifications, it avoids introducing circuit-specific intercepts that would complicate the calibration of the theoretical model.

Table 2 reports estimates for the probability of a successful undercut immediately after both drivers have stopped. The preferred specification reveals a concave relationship between undercut effectiveness and the share of the race passed between the pursuer's and leader's stops. The implied turning point occurs at approximately 0.31 of race distance $(2.9/(2 \cdot 4.6))$. As shown in panel (a) of Online Appendix Figure A2, this threshold lies above the bulk of the empirical support, so the relevant region of the relationship is monotonically increasing and concave. Tire advantage enters positively and significantly only through its linear term, indicating that fresher tires increase the likelihood of an immediate overtake, but without strong evidence of non-linearities. Taken together, these estimates highlight the trade-off faced by the pursuer when attempting an undercut: stopping earlier increases the duration of the strategic window (ΔLaps), while stopping later increases the tire advantage (ΔTire). As a result, the pursuer faces a non-trivial optimization problem, yet in the data undercuts lead to an immediate overtake in 48% of cases.

Turning to final overtake probabilities, Table 3 shows that conditional on a successful undercut, the pursuer remains ahead until the end of the interaction with a probability close to 0.66. In

Table 2: Intermediate overtake probabilities for undercuts ($P_u^1(l_L, l_P)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.671**	0.595**	0.508	4.057**	2.894*	1.934
	(0.283)	(0.250)	(0.322)	(1.736)	(1.483)	(1.301)
ΔTire	0.818***	0.812***	0.528	1.642	0.851	1.579
	(0.280)	(0.263)	(0.371)	(1.626)	(1.199)	(1.369)
ΔLaps^2				−5.940**	−4.643**	−2.937
				(2.431)	(2.087)	(1.821)
ΔTire^2				−1.133	−0.304	−1.616
				(1.923)	(1.426)	(1.517)
ΔTire × ΔLaps				−2.184	−0.383	0.851
				(2.742)	(2.427)	(2.387)
Constant	0.094	0.115	0.218	−0.165	0.035	−0.017
	(0.128)	(0.108)	(0.188)	(0.343)	(0.263)	(0.311)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	143	143	136	143	143	136
Mean of Dep. Variable	0.476	0.476	0.471	0.476	0.476	0.471
Within R-squared	0.074	0.095	0.117	0.145	0.145	0.167

Notes: This table reports estimates of the probability that the pursuer overtakes the leader immediately after pit stops when employing an undercut strategy. The dependent variable is a dummy equal to one if the pursuer is ahead one lap after the leader's stop. ΔLaps measures the share of the race distance between the pursuer's and the leader's pit stops, while ΔTire captures the pursuer's tire-age advantage during this interval. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential undercuts and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

our preferred specification, neither the share of the race remaining nor the tire advantage enters significantly. This lack of precision is not unexpected given the small size of this subsample. Nevertheless, even with limited statistical power, the estimated coefficient on the tire-age difference is sizeable and negative, indicating that a larger tire disadvantage for the pursuer after the stops tends to facilitate a counter-overtake by the leader. The coefficient on the share of remaining laps is also negative, though smaller in magnitude, suggesting that the probability of staying ahead declines as the leader's post-stop tire advantage persists over a longer remaining distance. Thus, while the estimates are imprecise, their signs are economically meaningful, and we retain the full set of coefficients when calibrating the model, as this improves the accuracy of the implied theoretical predictions.

Table 3: Final overtake probability conditional on a successful undercut ($P_u^2(l_L, l_P|O = 1)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.149	−0.203	0.474	−0.935	−0.935	−0.796
	(0.442)	(0.434)	(0.537)	(1.987)	(1.987)	(4.126)
ΔTire	−0.954**	−1.001***	−0.492	−2.203	−2.203	−0.864
	(0.344)	(0.369)	(0.570)	(2.805)	(2.805)	(5.509)
ΔLaps^2				0.436	0.436	1.269
				(1.837)	(1.837)	(3.146)
ΔTire^2				0.149	0.149	0.152
				(3.345)	(3.345)	(5.104)
ΔTire × ΔLaps				3.033	3.033	0.830
				(3.950)	(3.950)	(8.136)
Constant	0.889***	0.928***	0.735**	1.154*	1.154*	1.022
	(0.231)	(0.228)	(0.264)	(0.594)	(0.594)	(1.301)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	68	68	59	68	68	59
Mean of Dep. Variable	0.662	0.662	0.678	0.662	0.662	0.678
Within R-squared	0.205	0.323	0.422	0.214	0.312	0.426

Notes: This table reports estimates of the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate undercut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the former leader. ΔLaps measures the share of the race distance from the leader's stop to the end of the interaction, while ΔTire captures the leader's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential undercuts, conditional on the pursuer overtaking the leader immediately after the stops, and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

The results in Table 4 indicate that if an undercut fails to generate an immediate position change, subsequent overtaking by the pursuer is rare. On average, pursuers overtake only in about 7% of such cases. While the point estimates suggest that a longer remaining race distance may increase the likelihood of a late overtake, these effects are statistically insignificant across all specifications.

Taken together, Tables 2–4 highlight that undercuts can be highly effective when the pursuer succeeds in overtaking shortly after the stops and protecting the position until the end of the race. This outcome is more likely when the pursuer enjoys a substantial tire advantage and has sufficient time between the stops to exploit it. From the leader's perspective, an early stop by the pursuer

Table 4: Final overtake probability conditional on an unsuccessful undercut ($P_u^2(l_L, l_P|O = 0)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.206	0.121	0.570	2.206	1.896	1.127
	(0.163)	(0.146)	(0.496)	(1.441)	(1.917)	(1.202)
ΔTire	−0.076	−0.013	0.101	1.605	1.119	−0.769
	(0.095)	(0.089)	(0.329)	(1.119)	(1.426)	(1.439)
ΔLaps^2				−1.581	−1.464	−0.582
				(1.125)	(1.538)	(1.087)
ΔTire^2				−0.945	−0.633	0.913
				(0.645)	(0.701)	(1.289)
ΔTire × ΔLaps				−2.918	−1.936	1.198
				(1.976)	(2.662)	(2.782)
Constant	−0.025	0.034	−0.335	−0.626	−0.479	−0.444
	(0.092)	(0.088)	(0.344)	(0.423)	(0.555)	(0.371)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	75	75	69	75	75	69
Mean of Dep. Variable	0.067	0.067	0.043	0.067	0.067	0.043
Within R-squared	0.049	0.006	0.158	0.063	0.011	0.179

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction conditional on an unsuccessful intermediate undercut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. ΔLaps measures the share of the race distance from the leader's stop to the end of the interaction, while ΔTire captures the leader's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential undercuts in which the pursuer does not overtake the leader immediately after the stops and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

creates a strategic trade-off: responding immediately limits the pursuer's opportunity to overtake, while delaying the stop preserves a tire advantage that may facilitate a counter-overtake later in the race.

We next examine the determinants of overcut success. Table 5 reports estimates for the probability of an intermediate overtake when the leader stops first. The average probability of a successful overtake in these situations is approximately 0.25. In our preferred specification, the estimated probability decreases with the tire-age advantage of the leader during the interval between the two stops. The quadratic specification implies that this negative relationship holds over the empirically relevant range: the implied turning point corresponds to a tire-age difference of

Table 5: Intermediate overtake probabilities for overcuts ($P_o^1(l_L, l_P)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.208	0.208	0.123	−1.317	−1.317	0.853
	(0.312)	(0.312)	(0.298)	(1.576)	(1.576)	(1.579)
ΔTire	−0.341	−0.341	0.431	−3.638*	−3.638*	−1.300
	(0.377)	(0.377)	(0.412)	(1.899)	(1.899)	(1.987)
ΔLaps^2				0.756	0.756	−0.875
				(1.338)	(1.338)	(1.321)
ΔTire^2				3.661*	3.661*	2.522
				(2.093)	(2.093)	(2.209)
ΔTire × ΔLaps				3.999	3.999	−1.183
				(3.637)	(3.637)	(3.666)
Constant	0.327*	0.327*	−0.054	0.989**	0.989**	0.160
	(0.169)	(0.169)	(0.192)	(0.416)	(0.416)	(0.443)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	254	254	253	254	254	253
Mean of Dep. Variable	0.248	0.248	0.249	0.248	0.248	0.249
Within R-squared	0.034	0.014	0.083	0.058	0.020	0.102

Notes: This table reports estimates of the probability that the pursuer overtakes the leader immediately after pit stops when employing an overcut strategy. The dependent variable is a dummy equal to one if the pursuer is ahead one lap after his own stop. ΔLaps measures the share of the race distance between the leader's and the pursuer's pit stops, while ΔTire captures the leader's tire-age advantage during this interval. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential overcuts and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

about 50% ($3.6/(2 \cdot 3.7)$) of race distance, which is rarely exceeded in the data (see subplot (b) of Online Appendix Figure A2). We therefore interpret the effect of ΔTire as monotonically negative over the support of the data. We find no statistically significant dependence of intermediate overtake success on the share of the race passed between the two stops. If present, the estimated relationship mirrors that for tire advantage but is smaller in magnitude. Taken together, these results indicate that overcuts succeed in intermediate overtake with relatively low probability and become increasingly difficult as the leader's tire advantage grows with each additional lap before the pursuer stops.

However, conditional on a successful intermediate overtake, position reversals are extremely rare, as the pursuer then enjoys both a tire and a track advantage. As shown in Table 6, when the

Table 6: Final overtake probability conditional on a successful overcut ($P_o^2(l_L, l_P | O = 1)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.163	−0.163	−0.219	0.725	0.725	0.771
	(0.151)	(0.151)	(0.287)	(0.809)	(0.809)	(1.210)
ΔTire	−0.102	−0.102	−0.254	0.379	0.379	0.198
	(0.108)	(0.108)	(0.297)	(0.416)	(0.416)	(0.477)
ΔLaps^2				−0.752	−0.752	−0.768
				(0.833)	(0.833)	(1.155)
ΔTire^2				0.135	0.135	0.772
				(0.221)	(0.221)	(0.694)
ΔTire × ΔLaps				−1.579	−1.579	−2.269
				(1.627)	(1.627)	(2.329)
Constant	1.086***	1.086***	1.079***	0.874***	0.874***	0.863***
	(0.084)	(0.084)	(0.127)	(0.147)	(0.147)	(0.252)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	63	63	57	63	63	57
Mean of Dep. Variable	0.984	0.984	0.982	0.984	0.984	0.982
Within R-squared	0.034	0.009	0.036	0.085	0.117	0.144

Notes: This table reports estimates of the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate overcut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the former leader. ΔLaps measures the share of the race distance from the pursuer's stop to the end of the interaction, while ΔTire captures the pursuer's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential overcuts, conditional on the pursuer overtaking the leader immediately after the stops, and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

pursuer overtakes the leader after both drivers have completed their pit stops, he remains ahead until the end of the race with a probability of about 0.98. Given this near-deterministic outcome, we find no systematic dependence on race-state variables or strategic choices, consistent with the large and persistent advantage associated with securing track position at this stage of the race.

By contrast, when the pursuer does not succeed in overtaking immediately after the overcut, which is the most common outcome according to Table 5, we obtain robust and statistically significant effects of race-state variables. The corresponding estimates are reported in Table 7. We find that the probability of an eventual overtake increases with the share of the race remaining after the leader's stop, and this relationship is nonlinear. In our preferred specification, the implied

Table 7: Final overtake probability conditional on an unsuccessful overcut ($P_o^2(l_L, l_P|O = 0)$)

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.105	−0.089	−0.420	3.791**	3.978**	4.433**
	(0.282)	(0.300)	(0.393)	(1.845)	(1.837)	(1.923)
ΔTire	0.250	0.252	0.323	3.418	4.113*	5.729**
	(0.393)	(0.400)	(0.464)	(2.424)	(2.306)	(2.197)
ΔLaps^2				−3.263**	−3.316**	−3.737**
				(1.405)	(1.374)	(1.639)
ΔTire^2				−2.627	−3.279	−4.600**
				(2.236)	(2.085)	(2.014)
ΔTire × ΔLaps				−4.586	−5.652	−7.927**
				(4.088)	(3.850)	(3.226)
Constant	0.292	0.276	0.497*	−0.854	−0.954	−1.023
	(0.172)	(0.184)	(0.245)	(0.574)	(0.584)	(0.607)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	191	191	189	191	191	189
Mean of Dep. Variable	0.288	0.288	0.291	0.288	0.288	0.291
Within R-squared	0.077	0.054	0.076	0.096	0.079	0.105

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction conditional on an unsuccessful intermediate overcut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. ΔLaps measures the share of the race distance from the pursuer's stop to the end of the interaction, while ΔTire captures the pursuer's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies involving potential overcuts in which the pursuer does not overtake the leader immediately after the stops and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

turning point corresponds to a remaining race share of about 60% $(4.0/(2\cdot 3.3))$, a range that covers most observations in the data (see subplot (c) of Online Appendix Figure A2). We therefore interpret this effect as increasing over the empirically relevant support.

A similar pattern emerges for tire advantage. The probability of overtaking increases with the pursuer's tire advantage after the leader's stop, but at a decreasing rate. The estimated effect becomes flat once tire advantage exceeds roughly 63% of race distance $(4.1/(2 \cdot 3.3))$. As observations beyond this threshold are rare (subplot (d) of Online Appendix Figure A2), we again focus on the increasing portion of the relationship. Overall, these results indicate that when the immediate overcut fails, successful overtaking later in the race is more likely when the pursuer has

Table 8: Final overtaking probability under same-lap pit stops ($Pr_s^2(l_L, l_P)$))

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.939**	−0.949**	−1.681	−0.724	−0.516	−10.622
	(0.443)	(0.451)	(1.767)	(4.032)	(4.051)	(21.483)
ΔLaps^2				−0.165	−0.333	7.846
				(2.946)	(2.962)	(18.950)
Constant	0.978***	0.975***	1.358	0.912	0.844	3.558
	(0.338)	(0.343)	(1.349)	(1.342)	(1.346)	(5.374)
Circuit	No	RE	FE	No	RE	FE
Other Control	No	No	Yes	No	No	Yes
Observations	57	57	50	57	57	50
Mean of Dep. Variable	0.316	0.316	0.320	0.316	0.316	0.320
Within R-squared	0.092	0.097	0.101	0.092	0.098	0.111

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction when both drivers pit on the same lap. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. Because pit stops occur on the same lap, ΔTire is equal to zero by construction, and ΔLaps measures the share of the race distance from the joint stop to the end of the interaction. The regressions therefore include ΔLaps and its square to allow for non-linear effects. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. The sample is restricted to single-stop strategies with same-lap pit stops and follows the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

both sufficient remaining distance and a meaningful tire advantage.

Taken together, Tables 5–7 indicate that overcuts can be successful when the pursuer secures a sufficiently large tire advantage for the final stint and when enough of the race remains to exploit this advantage on track. At the same time, extending the first stint is costly, as it allows the leader to benefit from fresher tires during the interval between the two stops, reducing the likelihood of an immediate overtake.

Table 8 reports the results for drivers who pit on the same lap. In this case, the period between stops is zero by construction, and tire age differences are absent, so we focus exclusively on the final probability of overtaking as a function of the share of the race remaining and its square. Using the subsample of one-stop races, we do not detect statistically significant effects, likely due to the small sample size. In broader samples, however, the estimated probability of overtaking increases with the remaining race distance at a decreasing rate. This pattern is consistent with the interpretation that a longer remaining race provides more opportunities for overtaking, even in the absence of tire-based strategic asymmetries.

We summarize the estimated coefficients from our preferred specifications for all probability objects in Online Appendix Table A2. We additionally re-estimate the same models using alternative subsamples described in Section 4.4. The corresponding results, reported in Online Appendix Tables A3–A9, are qualitatively similar to those discussed above. As further robustness checks, we replace the linear probability model with logit and probit specifications, which yield comparable estimates (Online Appendix Table A10). Finally, we vary the threshold used to define close pit-stop interactions and obtain similar results under both stricter and more permissive definitions (Online Appendix Table A11).

6 Model predictions

In this section, we derive the predictions of the model introduced in Section 3, using the empirical estimates obtained in Section 5.

The model describes a dynamic interaction in which two drivers make sequential pit-stop decisions on each lap. To discipline the calibration and ensure consistency with observed racing behavior, we impose two restrictions. First, we parameterize the model using coefficients estimated from the random-effects specification on the subsample of races with single-stop strategies, which most closely matches the theoretical environment. Second, although the model can generate predictions for pit stops at any lap, very early and very late stops are rare in the data and therefore associated with imprecisely estimated overtake probabilities. To avoid extrapolation beyond the support of the data, we restrict the set of feasible pit-stop laps using the empirical distribution of pit-stop timing.

Figure 3 plots the distribution of pit stops in single-stop strategies, excluding races affected by yellow or red flags and rain. Based on this distribution, we identify the smallest interval of race distance that contains 80% of all observed stops and restrict pit-stop choices to this range. This interval spans from 21% to 65% of race distance and captures the region in which pit stops are most likely to occur in practice.

6.1 Predicted probabilities

We begin by illustrating how the predicted probabilities of overtaking vary with the pit-stop timing choices of the pursuer and the leader. For ease of interpretation, we normalize race length to 100 laps and use the coefficients from the random-effects specification estimated on the subsample of single-stop strategies. Figure 4 plots the predicted payoffs to the pursuer for different pit-stop choices, holding the leader's stop fixed.

Specifically, we consider three representative leader stop laps: lap 32 (early; light red), lap 43 (baseline; light blue), and lap 54 (late; light green). These choices divide the empirically relevant interval from 21% to 65% of race distance into four equal segments. For each leader stop, we

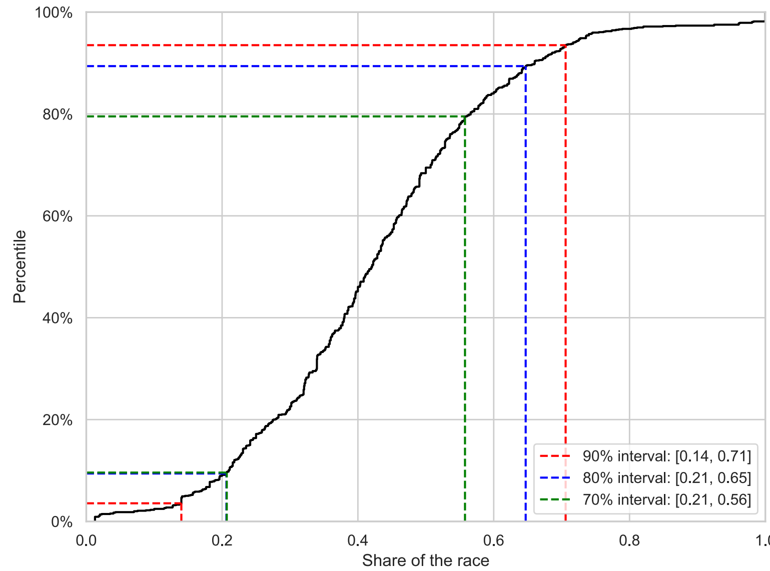

Notes: The figure plots the empirical cumulative distribution function (CDF) of pit-stop timing, measured as the share of race distance completed at the time of the stop. Vertical dashed lines indicate the shortest intervals that contain 70%, 80%, and 90% of pit stops, respectively. These intervals are used to define alternative pit windows in the theoretical analysis. The sample includes all single-stop strategies that satisfy the data-cleaning procedure described in Section 4.3.

Figure 3: Cumulative distribution of pit stops and implied pit windows

plot the predicted payoff for all feasible pursuer decisions, restricting both drivers' choices to the empirically supported interval, indicated by vertical dashed lines.

Figure 4 yields several clear insights. First, for any given leader stop, the most effective undercut involves stopping a few laps earlier than the leader. A window of roughly ten laps maximizes the probability of finishing ahead, as it allows the pursuer sufficient time to exploit the tire advantage. Stopping only one lap earlier typically provides too little time for the tire advantage to materialize, while stopping much earlier exposes the pursuer to a substantial tire disadvantage in the final stint, increasing the likelihood of a counter-overtake by the leader.

Second, the later the leader's stop, the more effective undercuts become. This reflects the fact that later stops imply more degraded tires for the leader and thus a larger temporary tire advantage for the pursuer between the stops.

Finally, overcuts display a markedly different pattern. The predicted probability of success increases monotonically with the pursuer's stop lap, making it optimal to delay the stop as long as possible within the feasible interval. This is consistent with the mechanism underlying overcuts: their effectiveness relies on maximizing tire advantage in the final stint. At the same time, overcuts become less effective when the leader also stops later, as this increases the leader's tire advantage between the stops and reduces the pursuer's relative advantage at the end of the race.

Figure 5 presents the analogous exercise from the leader's perspective. Fixing the pursuer's

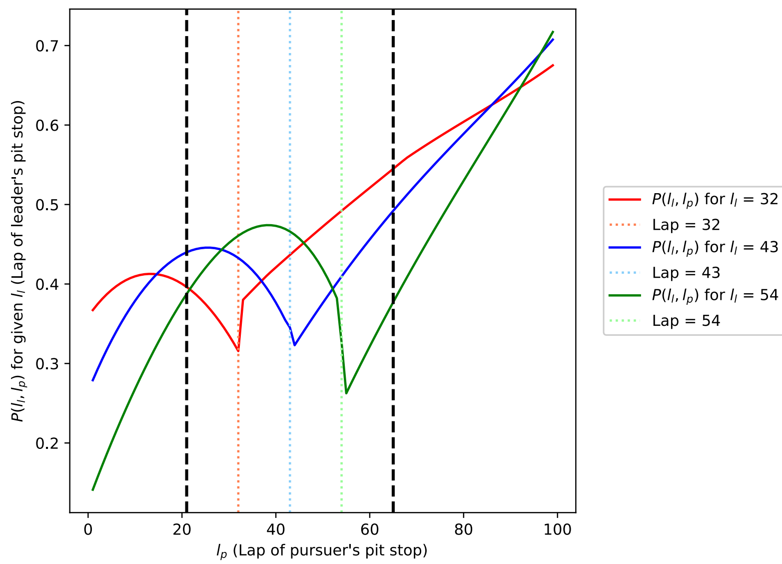

Notes: The figure plots the predicted probability that the pursuer overtakes the leader, $P(l_L, l_P)$, as a function of the pursuer's pit-stop lap l_P, for three fixed leader pit-stop choices. The leader is assumed to stop early (lap 32), at an intermediate point (lap 43), or late (lap 54), corresponding to representative points within the empirically relevant pit window. Vertical dotted lines indicate the leader's pit-stop lap in each case. Vertical dashed black lines mark the boundaries of the feasible pit window, restricted to the interval containing 80% of observed pit stops. Predictions are based on estimates obtained from single-stop strategies only, using the random-effects specification, and are computed for a 100-lap race.

Figure 4: Predicted effectiveness of the pursuer's pit-stop timing

pit-stop choice, we plot the predicted probability of an overtake for different leader stop laps. As before, we consider three representative pursuer strategies: an early stop at lap 32 (light red), a baseline stop at lap 43 (light blue), and a late stop at lap 54 (light green). Since the leader's objective is to avoid being overtaken, we interpret lower predicted probabilities as more favorable outcomes.

A clear pattern emerges: for any fixed pursuer strategy, the probability of an overtake is minimized when the leader stops on the same lap as the pursuer. This highlights why same-lap stops are rarely observed in the data: because the pursuer moves second, he can always avoid simultaneous stopping by adjusting his strategy once the leader's choice is revealed.

Beyond this, the probability of an overtake typically increases with the time between stops, reflecting the growing temporary tire advantage of the pursuer. For sufficiently large stop gaps, particularly when the pursuer stops very early, this relationship becomes non-monotonic. In such cases, further delaying the leader's stop can reduce the overall overtake probability by generating a substantial tire advantage in the final stint, thereby increasing the likelihood of a counter-overtake. Nevertheless, for the pursuer strategies shown in the figure, delaying the stop is less effective for the leader than matching the pursuer's stop, as same-lap stopping minimizes the probability of being overtaken.

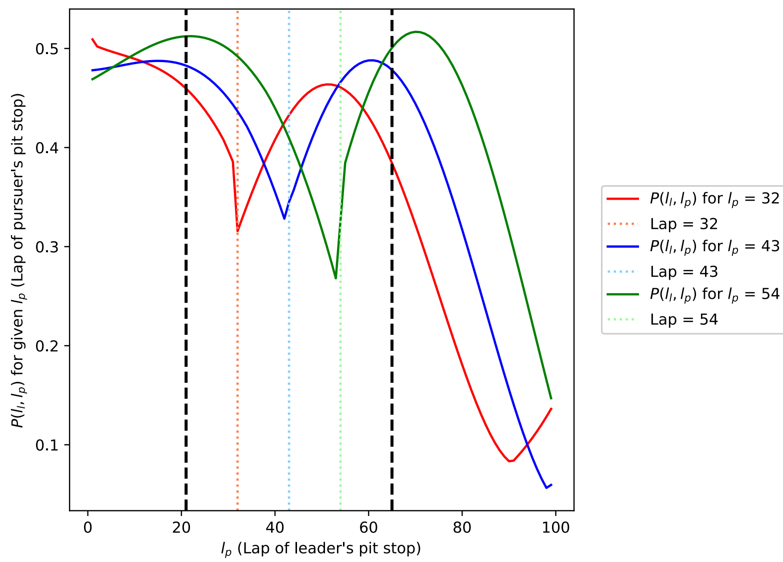

Notes: The figure plots the predicted probability that the pursuer overtakes the leader, $P(l_L, l_P)$, as a function of the leader's pit-stop lap l_L, for three fixed pursuer pit-stop choices. The pursuer is assumed to stop early (lap 32), at an intermediate point (lap 43), or late (lap 54), corresponding to representative points within the empirically relevant pit window. Vertical dotted lines indicate the pursuer's pit-stop lap in each case. Vertical dashed black lines mark the boundaries of the feasible pit window, restricted to the interval containing 80% of observed pit stops. Predictions are based on estimates obtained from single-stop strategies only, using the random-effects specification, and are computed for a 100-lap race.

Figure 5: Predicted effectiveness of the leader's pit-stop timing

To summarize the joint strategic interaction, Online Appendix Figure A3 presents heat maps of the predicted overtake probability for all combinations of leader and pursuer stop laps. Panel (a) shows the full set of combinations, while Panel (b) restricts attention to the empirically relevant interval used in the model. These heat maps reinforce the insights from Figures 4 and 5 and provide a compact visualization of how overtaking probabilities vary across the two-dimensional strategy space.

6.2 Predicted optimal choices

Having characterized the main properties of the estimated overtaking probabilities, we now turn to the solution of the strategic game.

We begin by analyzing drivers' best responses to their opponents' actions. While Figures 4 and 5 already provided intuition based on unconditional comparisons, we now explicitly incorporate the sequential structure of the game and assume fully rational behavior. In particular, given the timing of moves, if the pursuer (leader) commits to stopping on lap X, the leader's (pursuer's) feasible responses are restricted to stopping on lap $X + 1$ (X) or on any subsequent lap.

For consistency with the previous analysis, we continue to consider a 100-lap race and restrict admissible pit-stop choices to laps 21 through 65.

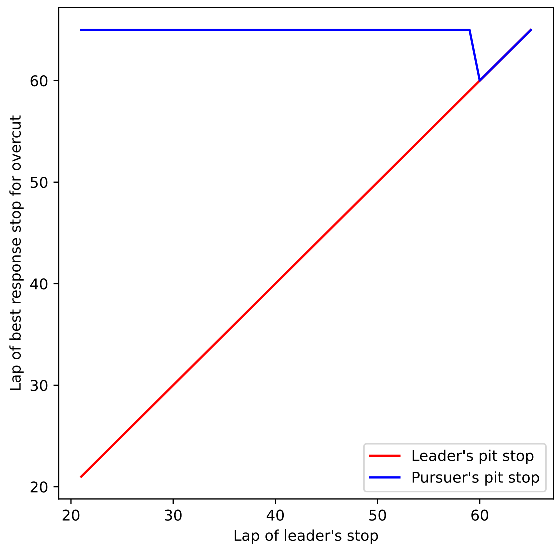

(a) Best response of pursuers

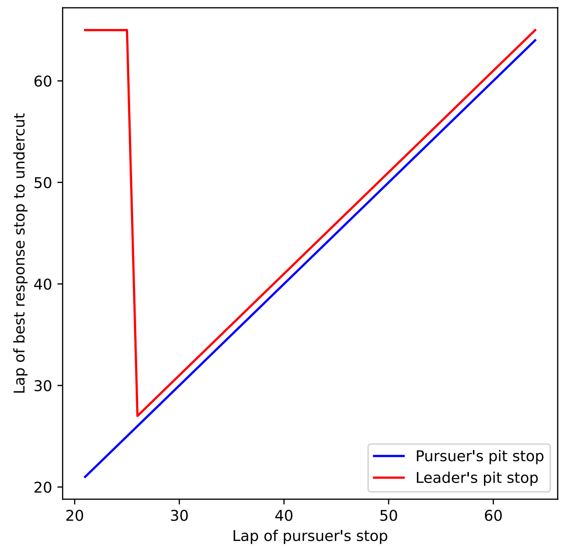

(b) Best response of leaders

Notes: This figure plots best-response functions for pit-stop timing implied by the theoretical model. Panel (a) shows the pursuer's optimal pit-stop lap as a function of the leader's pit-stop decision, while Panel (b) shows the leader's optimal pit-stop lap as a function of the pursuer's decision. The pursuer's decisions are shown in blue, and the leader's decisions in red. Best responses are computed using estimated overtake probabilities from single-stop strategies and are restricted to the empirically relevant pit window in a 100-lap race.

Figure 6: Best-response functions in pit-stop timing

Figure 6 depicts the best-response functions of both drivers. We first consider the case in which the leader stops before the pursuer. For each leader's stopping lap (shown in red), subplot (a) reports the pursuer's optimal response (blue). When the leader stops before lap 59, the pursuer's best response is to delay his stop as much as possible. In this region, the leader's temporary tire advantage between stops is limited, making it optimal for the pursuer to maximize his own tire advantage in the final stint. By contrast, when the leader stops later than lap 59, his tire advantage becomes sufficiently large that delaying the pursuer's stop is no longer effective. In this case, the pursuer's optimal response is to stop immediately.

Subplot (b) considers the opposite scenario, in which the pursuer stops first and the leader chooses a response. If the pursuer stops before lap 27, the leader optimally delays his stop until the last feasible lap, thereby accumulating a substantial tire advantage for the final stint. However, once the pursuer's stop occurs later in the race, the leader's best response switches to stopping as soon as possible, typically on the following lap. In this region, delaying the stop no longer generates a sufficiently large tire advantage to ensure a counter-overtake, so the leader instead minimizes the risk of being undercut by responding immediately.

Using these best responses and assuming fully rational behavior, we compute equilibrium payoffs for all possible stopping laps. For each lap, we consider two scenarios: (i) the leader stops

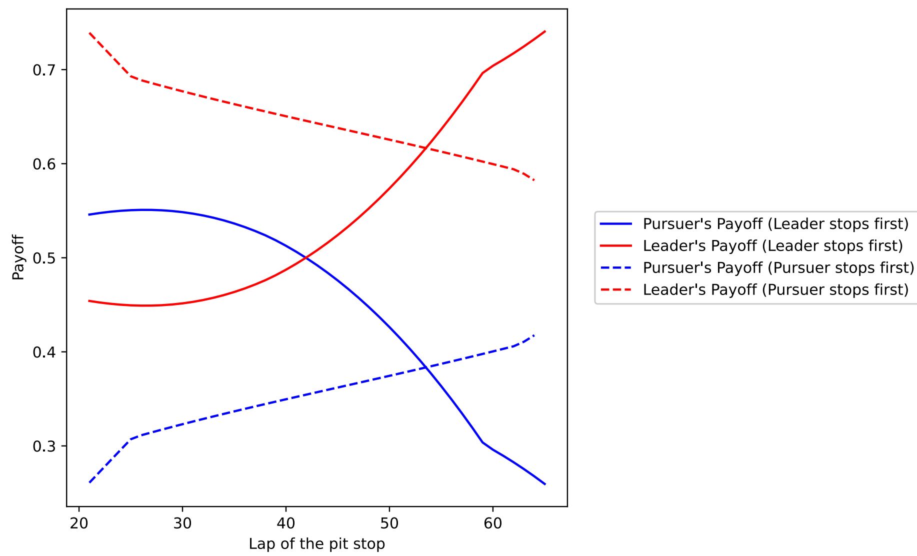

Notes: This figure plots the payoffs implied by the theoretical model for the pursuer and the leader as functions of pit-stop timing. Solid lines correspond to the case in which the leader stops first, while dashed lines correspond to the case in which the pursuer stops first. The pursuer's payoffs are shown in blue, and the leader's payoffs in red. Payoffs are computed using estimated overtake probabilities from single-stop strategies and are evaluated over the empirically relevant pit window in a 100-lap race.

Figure 7: Payoffs for pursuer and leader

first and the pursuer responds optimally, and (ii) the pursuer stops first and the leader responds optimally. The resulting payoffs are shown in Figure 7. Solid lines correspond to scenarios in which the leader stops first, while dashed lines represent scenarios in which the pursuer stops first. The pursuer's payoff, the probability of overtaking, is shown in blue, and the leader's payoff, equal to one minus this probability, is shown in red.

Figure 7 reveals a clear pattern. When the leader stops first, the pursuer's payoff is generally decreasing in the stopping lap (except for the very early laps), while the leader's payoff is correspondingly increasing. Conversely, when the pursuer stops first, his payoff is increasing in the stopping lap. As a result, there exists a unique point at which the pursuer's expected payoff is the same regardless of which driver stops first. The payoff structure generated by this interaction closely resembles a centipede game, in which delaying action benefits both players up to a point but creates incentives for preemption.

We solve the game by backward induction. Early in the race, stopping is unattractive for both drivers, as payoffs are low and increase with each additional lap. However, drivers cannot postpone their stop indefinitely. At some point, the payoff from stopping on the current lap exceeds the payoff from waiting and allowing the opponent to stop first. This point corresponds to the intersection of the payoff curves in Figure 7 and determines the SPNE.

Because stopping decisions occur in discrete time, the intersection typically lies between two integer laps. Suppose the intersection occurs between laps X and $X + 1$. If the leader's payoff

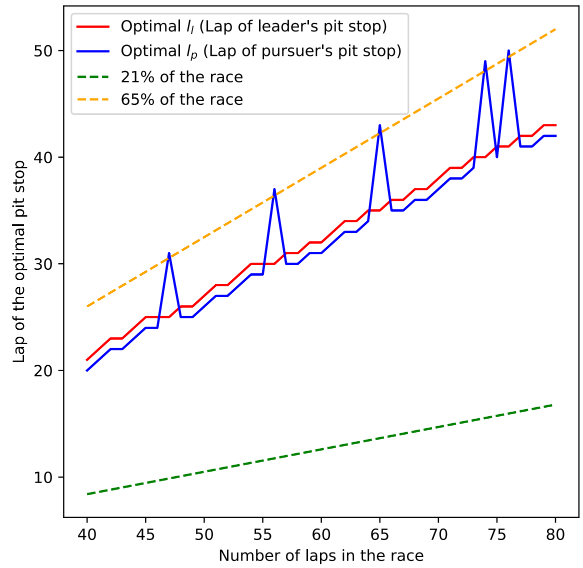

Notes: This figure plots the SPNE pit-stop timing implied by the theoretical model as a function of total race length. The red line shows the leader's equilibrium pit-stop lap (l_L^*), while the blue line shows the pursuer's equilibrium pit-stop lap (l_P^*). Dashed lines indicate fixed fractions of the race distance, corresponding to 21% and 65% of the race, respectively. Equilibrium strategies are computed using estimated overtake probabilities from single-stop strategies and are restricted to the empirically relevant pit window.

Figure 8: SPNE pit-stop timing as a function of race length

from stopping on lap $X+1$ exceeds his payoff when the pursuer stops on lap X, then the pursuer stops on lap X, and the leader responds optimally as described in subplot (b) of Figure 6. This equilibrium outcome corresponds to an undercut by the pursuer. Conversely, if the leader's payoff from stopping on lap $X+1$ is lower than his payoff when the pursuer stops on lap X, then the leader stops on lap $X+1$, and the pursuer responds optimally by attempting an overcut, as illustrated in subplot (a) of Figure 6.

Which of these two scenarios is realized in equilibrium depends on which of the two corresponding payoffs is marginally larger. In practice, however, the difference between them is negligible. We therefore interpret the model as admitting two SPNE in pure strategies that yield nearly identical payoffs.

To illustrate these equilibria, we solve the game for different race lengths and report the results in Figure 8. As shown in the figure, either equilibrium can arise depending on small differences in payoffs. In the equilibrium featuring an undercut, the pursuer stops first and the leader responds immediately by pitting on the following lap. In the alternative equilibrium, the leader stops first and the pursuer optimally attempts an overcut by delaying his stop until the last feasible lap.

Importantly, across all considered race lengths, the timing of the first stop in the interacting pair is remarkably stable, occurring at approximately 50% of race distance. This point coincides with the intersection of the payoff curves discussed above and therefore plays a central role in determining equilibrium behavior.

As a robustness check, we also compute equilibrium predictions using alternative data samples. Solving the model using all pit stops and, separately, only the last stops in strategies yields very similar equilibrium patterns, as shown in Online Appendix Figures A4 and A5. These results further support the robustness of our main findings.

To further assess the robustness of our results, we allow for alternative pit windows that depend on the tire compound used in the first stint. In the baseline analysis, we pooled all pit stops irrespective of tire compounds. However, Formula 1 regulations allow drivers to choose among three compounds that differ in pace and durability, which naturally generates compound-specific pit windows.

Online Appendix Figure A6 illustrates these differences by plotting the cumulative distribution of pit stops separately for each starting compound. For each compound, we identify the shortest pit window that contains 70% of observed stops. We find that drivers starting on the soft compound typically pit between 21% and 49% of race distance. For the medium compound, the corresponding window is from 29% to 59%, while for the hard compound it shifts later, spanning from 49% to 73% of the race. We use these compound-specific pit windows to recompute predicted best responses and equilibrium behavior.

We first consider the case in which both drivers start on soft tires and are therefore restricted to pitting between laps 21 and 49. The resulting predictions are shown in Online Appendix Figure A7. Qualitatively, the model delivers patterns similar to the baseline case: if the pursuer stops first, the leader's best response is to pit immediately, whereas if the leader stops first, the pursuer optimally delays his stop as long as possible, except for very late leader stops, in which case stopping immediately becomes optimal. The main difference relative to the baseline is that when the pursuer stops very early, the leader now prefers to respond immediately rather than postponing the stop. This change reflects the narrower pit window: with less scope to delay the stop, the leader cannot generate a sufficiently large tire advantage for the final stint, making immediate response optimal. Solving the game under these best responses yields equilibrium outcomes in which the first stop occurs earlier, at approximately 34% of race distance, compared to about 50% in the baseline case.

We next analyze the scenario in which both drivers start on medium tires. The predicted best responses and equilibrium outcomes, reported in Online Appendix Figure A8, are very similar to those obtained for the soft compound. Finally, when both drivers start on hard tires, we again find qualitatively comparable patterns, as shown in Online Appendix Figure A9.

Overall, while the precise SPNE depend on the assumed pit window, the structure of best responses is remarkably stable across specifications. Since best-response functions require assuming rationality only for the responding driver – rather than for both drivers simultaneously, as in equilibrium analysis – we focus our empirical tests on these best-response predictions. This approach allows us to separately assess the rationality of leaders and pursuers and provides a more transparent link between the theoretical model and observed behavior.

7 Discussion of Drivers' Rationality

7.1 Do Drivers Respond to Competitors' Strategies?

A natural first step in assessing whether drivers behave strategically is to test whether their pit-stop decisions respond to the presence of nearby competitors. If drivers do not take opponents' actions into account, the timing of pit-stops should be largely invariant to whether a driver is engaged in a close on-track battle or is effectively "alone" on the circuit. Conversely, systematic differences in the distribution of pit-stop timing between competitive and non-competitive situations would indicate that drivers condition their choices on their rivals' behavior.

To investigate this, we examine all race situations that satisfy the data-cleaning criteria in Section 4.3, with two modifications: (i) we restrict attention to single-stop strategies, and (ii) we do not exclude observations in which drivers are far apart. We group all ordered pairs of drivers into two categories based on pre-pit-stop proximity: close pairs (gap ≤ 3 seconds) and distant pairs (gap > 3 seconds). For each pair, we identify the first pit-stop (the earlier of the two stops) and the second pit-stop (the later one), and we construct cumulative distribution functions (CDFs) of these timings separately for close and distant pairs. Figure 9 displays the resulting four CDFs.

Two preliminary observations are worth noting. First, by construction, the CDF of first stops lies strictly to the left of the CDF of second stops. Drivers rarely pit on the same lap even when they are not engaged in direct competition, due to differences in tire management, fuel usage, and team strategy. This ordering is evident for both close and distant pairs. The key finding, however, emerges when we compare CDFs across proximity groups. If drivers ignored the presence of competitors, then the distribution of pit-stop laps should be similar whether a driver is running close to a rival or not. Instead, the data show substantial and systematic differences.

For first stops, the CDF for close pairs lies noticeably to the left of that for distant pairs. In other words, when drivers are engaged in a close on-track battle, the first stop in the pair occurs earlier on average than when the same drivers are running far apart. This pattern is exactly what one would expect if drivers attempt to use an early stop either to avoid being undercut or to launch an undercut of their own.

The pattern for second stops is more nuanced but equally informative. In the early part of the race, the CDF for second stops among close pairs closely resembles the CDF for first stops among distant pairs. That is, when two drivers are fighting, the later-stopping driver tends to delay only as much as a non-competitive driver who prefers an early stop. As the race progresses, however, the CDF for second stops among close pairs gradually converges toward the CDF for second stops of distant pairs. By the end of the race, the timing of the second stop in a close-pair

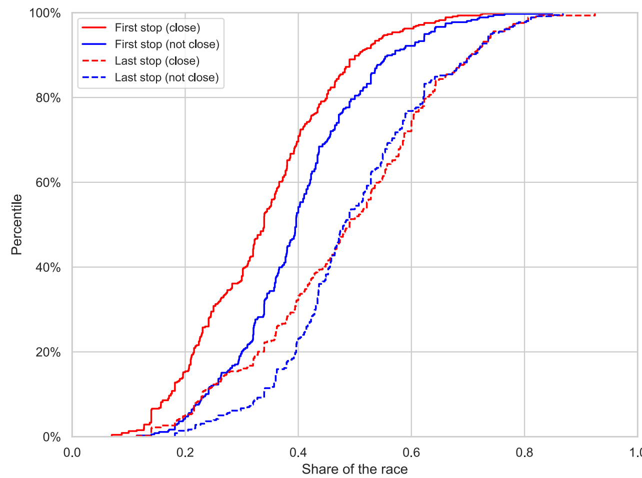

Notes: The figure plots empirical cumulative distribution functions (CDFs) of pit-stop timing, measured as the share of race distance completed at the time of the stop. The unit of observation is a pair of drivers who each make exactly one pit stop in a race. Within each pair, the *first stop* refers to the driver who pits earlier, while the *last stop* refers to the driver who pits later. Solid lines correspond to first pit stops, and dashed lines correspond to last pit stops. Red lines restrict the sample to close pit-stop interactions, defined as pairs of drivers who are within three seconds of each other on track at the time of the first stop, while blue lines include pairs that are not classified as close. The sample follows the data-cleaning procedure described in Section 4.3.

Figure 9: Cumulative Distribution of First and Last Pit Stops within One-Stop Driver Pairs

battle resembles that of a naturally late-stopping driver who is not engaged in direct competition.

In additional analyses, we disaggregate all pit-stop pairs into potential undercuts – situations in which the pursuing driver pits first – and potential overcuts, in which the leading driver pits first. Panels (a) and (b) of Online Appendix Figure A10 display the corresponding CDFs. The patterns mirror those in Figure 9. For both undercut and overcut scenarios, the first pit-stop in the pair occurs systematically earlier when the drivers are close to each other than when they are far apart, regardless of whether it is the leader or the pursuer who stops first. The same holds for second pit-stops: in the early part of the stint, second stops among close pairs occur noticeably earlier than those among distant pairs, while toward the end of the stint the two distributions converge. The only meaningful asymmetry is that the early-race gap between close and distant observations is more pronounced in potential undercuts.

We corroborate these graphical findings with a regression-based analysis. Specifically, we regress the pit-stop lap – expressed as a fraction of total race distance – on an indicator for pre-pit-stop proximity, controlling for circuit and year fixed effects and a rich set of covariates that include tire-compound information, each driver's proximity to other cars ahead and behind, and the position being contested. Online Appendix Table A12 reports the estimated coefficients. For first stops, the proximity dummy is negative and statistically significant across both undercut and overcut subsamples, with nearly identical magnitudes of approximately 0.031-0.035. In a 60-lap race, this corresponds to roughly two laps earlier. Second stops also occur earlier when drivers are close to each other, though the effect is smaller; among potential overcuts, it becomes statistically insignificant, consistent with the flattening of the CDFs observed in the panel (b) of Online Appendix Figure A10.

Taken together, these results provide clear evidence that drivers systematically adjust the timing of their pit-stops in response to nearby rivals. Their choices differ meaningfully from the timing patterns observed when no competitor is within strategic range, indicating that pit-stop decisions in Formula 1 are indeed shaped by the competitive environment.

7.2 Do Drivers Respond Optimally to Competitors' Strategies?

The previous subsection demonstrates that drivers systematically adjust their pit-stop timing in response to nearby competitors. We now turn to a more subtle question: how do they adjust their behavior, and to what extent do their observed decisions align with the optimal responses predicted by our model? Section 6.2 characterizes the equilibrium logic governing pit-stop timing in strategic duels. In particular, when the pursuing driver pits first (a potential undercut), the leading driver's optimal response is discontinuous: if the pursuer pits early, the leader should delay his own stop until the end of the pit window; if the pursuer pits late, the leader should instead pit immediately. For potential overcuts – when the leader pits first – the structure is analogous, but the switching point at which the pursuer should optimally stop immediately occurs substantially later in the stint.

Directly testing these sharp switching predictions in the data is challenging. The optimal response threshold depends on track characteristics, car-specific degradation profiles, and driver-specific performance. With the available data, we cannot identify the exact optimal switching lap for each driver-track pair. Nevertheless, the model yields robust comparative-static predictions that do not rely on knowing the precise threshold. First, the optimal gap between the two pit-stops is non-increasing in the lap of the first stop: later initiating stops should be followed by a shorter delay before the opponent responds. Second, the switching point should occur earlier for potential undercuts than for potential overcuts. Consequently, in the data we expect to observe that: (i) among undercuts, later pit-stops by the pursuer are followed by shorter intervals before the leader pits; (ii) among overcuts, later pit-stops by the leader are similarly followed by shorter intervals before the pursuer responds. These predictions correspond directly to the comparative statics illustrated in Figure 6.

We now examine whether real-world behavior is consistent with these patterns. For this analysis, we use the same cleaned sample as in the probability estimations, restricting attention to stints in which both drivers execute a single scheduled stop. Using the same data for both parts of the analysis does not pose a concern. The probability estimations relied on all pit-stops and focused on mapping how outcome probabilities vary with strategic timing choices. In contrast, the present exercise does not re-estimate those probabilities and does not use the data on the

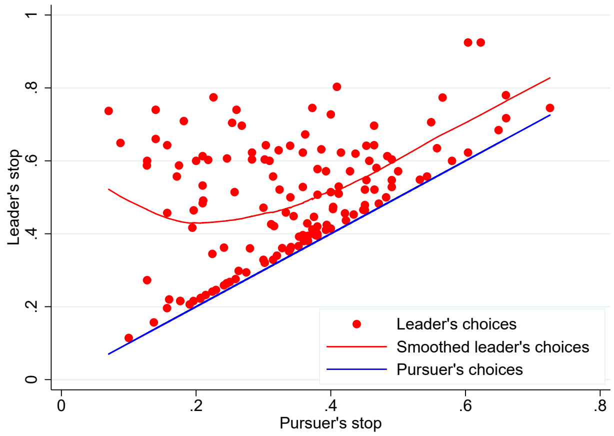

(a) Responses of leaders

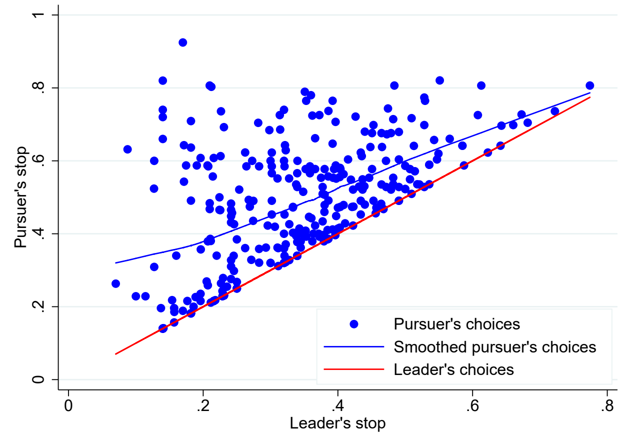

(b) Responses of pursuers

Notes: This figure plots empirical response relationships in pit-stop timing for pairs of drivers involved in close pit-stop interactions. The unit of observation is a pair of drivers who each make exactly one pit stop in a race. Pit-stop timing is measured as the share of race distance completed at the time of the stop. Panel (a) shows leaders' pit-stop choices as a function of pursuers' pit-stop timing, while Panel (b) shows pursuers' pit-stop choices as a function of leaders' pit-stop timing. Dots represent observed pit-stop choices. Solid lines show smoothed response functions estimated using local polynomial regressions. Throughout the figure, leaders' choices and responses are shown in red, while pursuers' choices and responses are shown in blue. The sample is restricted to close pit-stop interactions, defined as pairs of drivers who are within three seconds of each other on track at the time of the first pit stop. The sample follows the data-cleaning procedure described in Section 4.3.

Figure 10: Empirical responses of pursuers and leaders

final outcomes of the drivers; it simply examines the distribution of actual timing decisions made by drivers. Because drivers exhibited wide variation in pit-stop strategies and because our tests rest only on pit stops' choices, using the same underlying races does not cause any problems for interpretation of our results.

Figure 10 summarizes the relationship between the timing of the first stop and the response timing of the opponent. Panel (a) focuses on potential undercuts, plotting the leader's response times on the vertical axis against the pursuing driver's pit-stop lap on the horizontal axis. For early initiating stops, leaders display substantial heterogeneity: some pit almost immediately, while many delay their response deep into the stint. As the pursuing driver's stop occurs later, however, leaders respond more quickly, with the smoothed average showing a clear convergence to the pursuer's choices. This pattern is precisely in line with the model's predictions: later initiating stops compress the optimal response window.

Panel (b) presents the analogous relationship for potential overcuts. Here too the response interval decreases as the leader pits later, confirming the predicted monotonicity. However, the effect is less pronounced than in undercuts, and the slope of the smoothed relationship is shallower. Notably, even when leaders pit relatively early, there remains a substantial mass of observations in which the trailing driver responds immediately. Under the model, such behavior is generally suboptimal, suggesting that some drivers deviate from the theoretically efficient response in these scenarios.

We formalize these patterns using linear regressions in which the dependent variable is the fraction of the race between the two pit-stops in a pair, and the key explanatory variable is the race fraction at which the initiating driver pits (the pursuer for undercuts; the leader for overcuts). Online Appendix Table A13 reports the results. Columns 1 and 4 present simple OLS specifications; columns 2 and 5 add circuit and year fixed effects; and columns 3 and 6 further include controls for tire compound, each driver's proximity to other cars, and the position being contested. Across all specifications, the coefficient on the initiating driver's pit-stop lap is negative, indicating that later initiating stops are followed by shorter response intervals. Moreover, the coefficient is consistently larger in absolute value for undercuts, in line with the graphical evidence in Figure 10, and suggests that some drivers' decisions are not optimal when they try to overcut.

A potential concern with the preceding analysis is that using the fraction of the race between pit-stops as the dependent variable may mechanically generate a negative relationship with the lap of the initiating stop. Toward the end of a stint, tires are more heavily degraded, limiting the feasible range of viable stop laps. As a result, even in the absence of strategic adjustment, later initiating stops would tend to be followed by shorter feasible response windows. To verify that our results are not driven solely by such mechanical constraints, we employ alternative dependent variables that are less sensitive to end-of-stint tire limitations.

Online Appendix Table A14 reports estimates from a set of specifications in which the dependent variable is an indicator variables equal to one if the responding driver pits within X laps of the initiating driver's pit-stop. We consider thresholds of $X = 1, 3, 5$, and 10 laps in columns (1) through (4), respectively. The idea behind these restricted windows is straightforward. For sufficiently small values of X, drivers typically retain the option of pitting both within and beyond the window regardless of stint position; tire wear is rarely so severe as to force a stop within, for example, one or three laps after the opponent. Consequently, if later initiating stops predict higher probabilities of responding within X laps even for small X, this would indicate genuine strategic behavior rather than mere mechanical compression of feasible choices.

Panel A of Online Appendix Table A14 focuses on leaders' best responses in potential undercuts, while Panel B examines pursuers' responses in potential overcuts. Across both panels, all coefficients are positive, indicating that later initiating stops are associated with a higher probability that the responding driver pits within X laps. As expected, the effects become larger and more precisely estimated for larger windows. For example, in Panel A, a 10-percentage-point increase in the pursuer's pit-stop timing raises the probability that the leader responds within three laps by 7.76 percentage points – a sizeable and statistically significant effect. Crucially, a three-lap window is sufficiently short that drivers typically retain the ability to delay their response even late in the stint; thus, this pattern cannot be explained purely by mechanical end-of-stint constraints.

Taken together, the results in Online Appendix Table A14 reinforce our earlier findings. The relationship between the timing of the initiating stop and the timing of the response is not an artifact of diminishing late-race flexibility. Instead, drivers systematically compress the response window when strategic pressure is higher, and they do so more aggressively in potential undercuts than in potential overcuts. These patterns closely match the comparative-static predictions of our theoretical model, but also suggest that drivers' responses in overcuts are farther from the model-implied optimum.

7.3 Who Responds Optimally to Competitors' Strategies?

Thus far we have shown that, on average, drivers' pit-stop timing is broadly consistent with the qualitative prescriptions of our model. In the case of overcuts, however, the evidence suggests that observed decisions are, on average, somewhat less tightly aligned with the model's benchmark. A natural follow-up question is whether some teams or drivers make systematically better strategic choices than others. Answering this question requires us to exploit cross-team variation in observed responses while acknowledging a key identifying assumption: because we do not observe internal measures of a team's "strategic quality," we assume that a team's strategic capability is roughly stable over the relatively short window covered by our sample and therefore that cross-team differences in observed behavior reflect persistent differences in decision rules.

To study heterogeneity, we estimate the same slope relationship between the inter-stop gap and the initiating driver's pit-lap separately for each constructor. Intuitively, the slope measures how aggressively a team shortens or lengthens the response interval as the initiating stop moves later in the stint. Figure 11 summarizes these per-team estimates and associates them with realized success rates for the corresponding strategic situations. Panels (a) and (c) focus on leaders' best responses to potential undercuts (the leader is the player who may be undercutted), while Panels (b) and (d) report pursuers' best responses in potential overcuts (the pursuer is the trailing driver who might attempt an overcut). Panels (a)-(b) plot the estimated slopes for each team; Panels (c)-(d) plot the teams' empirical success rates in the corresponding role.

Two patterns stand out. First, teams differ substantially in their measured response slopes. In Panel (a) (leaders responding to potential undercuts), all teams exhibit the theoretically predicted negative slope: the later the pursuer pits, the sooner the leader tends to respond. In particular, Haas, Mercedes, Red Bull, and Williams display strongly negative slopes and therefore behave very close to the model-predicted optimal best response. By contrast, Alpine's estimated slope is the lowest, placing it farthest from the theoretical benchmark.

Second, alignment with the theoretical best response correlates with empirical effectiveness. Panel (c) shows team-level success rates for leaders who face a potential undercut; two of the five teams closest to the theoretical prediction are among the highest-performing in terms of keeping position after the interaction. Conversely, the team that departs most from the theory (Alpine) ranks only eighth by success rate. There is an important exception: Haas, despite following the theoretical best-response closely, appears relatively weak in realized success. This discrepancy is

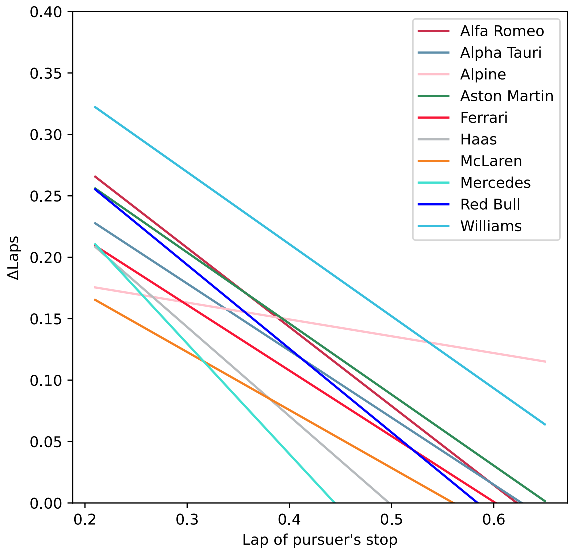

(a) Responses of leaders by teams

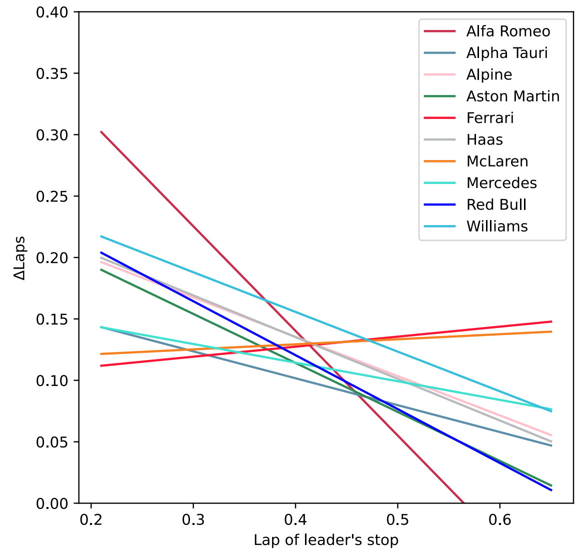

(b) Responses of pursuers by teams

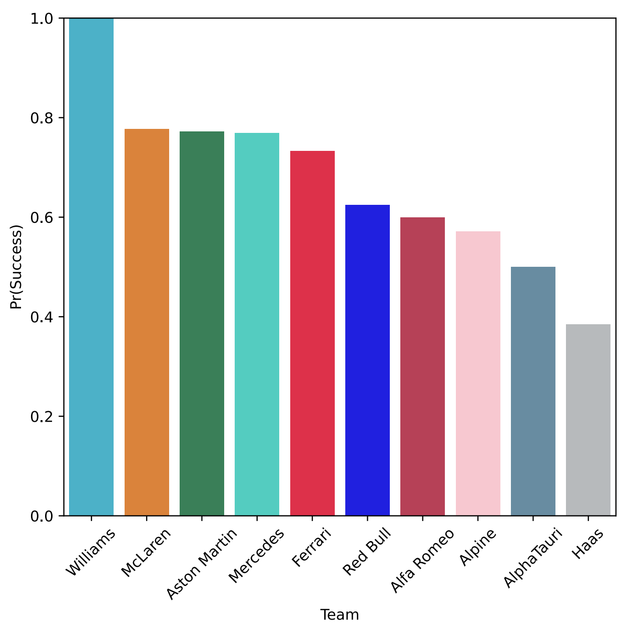

(c) Success rates of leaders' decisions by teams

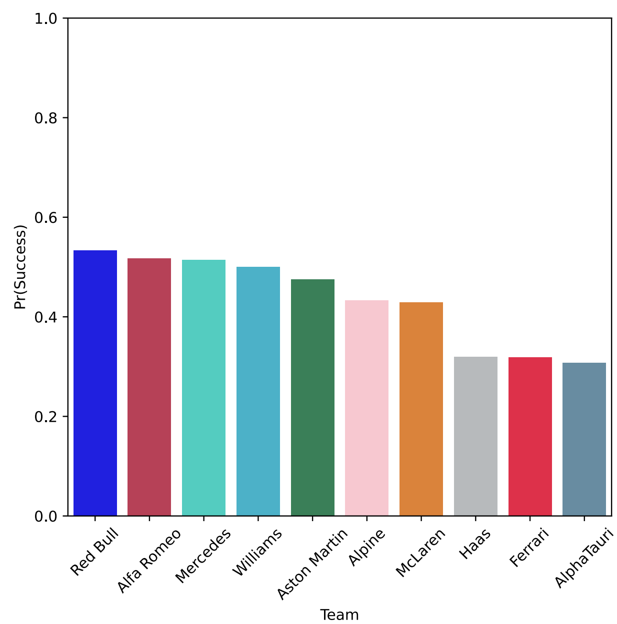

(d) Success rates of pursuers' decisions by teams

Notes: This figure reports empirical responses in pit-stop timing and associated success rates, disaggregated by team. Panels (a) and (b) plot estimated responses in pit-stop timing separately for leaders and pursuers. Each line corresponds to a team-specific linear fit relating the driver's pit-stop timing to the opponent's pit-stop decision, with pit-stop laps expressed as the share of race distance completed at the time of the stop. Panel (a) shows leaders' responses to the pursuer's pit-stop timing, while Panel (b) shows pursuers' responses to the leader's pit-stop timing. Panels (c) and (d) report team-level success rates of pit-stop decisions for leaders and pursuers, respectively. Success is defined as an overtake by the pursuer by the end of the race, consistent with the definition used in the main analysis. The sample consists of close pit-stop interactions within one-stop strategies, following the data-cleaning procedure described in Section 4.3. Throughout the figure, teams are identified by color as indicated in the legend.

Figure 11: Responses and Success Rates by Teams

plausibly explained by non-strategic factors outside the scope of our model – most notably car performance. If a team's car is comparatively slow, a theoretically correct timing decision may nevertheless fail to secure the position by race end because the rival simply has superior pace.

The story is similar from the pursuers' perspective. Panel (b) shows per-team slopes for the pursuer's best responses after a leader-stopping initiating action. Here, Alfa Romeo and Red Bull appear closest to the theoretical optimum, while Ferrari and McLaren lie furthest away and even have positive slopes. Panel (d) reports effectiveness when the leader stops first: the teams closest to the theory (again, Alfa Romeo and Red Bull) achieve the highest success rates, while McLaren and Ferrari – which depart most from the optimal response – rank only seventh and ninth out of ten.

Taken together, Figure 11 conveys two lessons. First, teams are heterogeneous in how they implement pit-stop responses: some systematically follow rules that are close to the model's optimal best responses while others do not. Importantly, deviations from optimal behavior are more pronounced in the case of overcuts, which helps explain our earlier finding that drivers, on average, respond less optimally in overcuts than in undercuts. Second, closer adherence to the model's prescriptions tends to translate into higher success in the on-track duel, which provides an additional, independent validation of the model. At the same time, strategic correctness is not perfectly collinear with overall team strength. Well-resourced teams with strong cars (e.g., Ferrari) can nevertheless make systematically suboptimal timing choices in certain strategic regimes (notably overcuts), undermining potential points gains; conversely, smaller teams (e.g., Haas, Alfa Romeo) sometimes exhibit superior decision rules even when their cars are not class-leading.

Building on the observation that teams differ markedly in how closely they follow the model's predicted best responses, we next examine whether driver characteristics also shape strategic quality. Although pit-stop calls are typically issued by the team, they arise from joint discussion between the pit wall and the driver. Teams often provide the final instruction, but drivers contribute information about tire condition, grip, and race context, and in some cases push for or resist a call. It is therefore plausible that driver-specific attributes – experience, skill, composure under pressure – may influence how effectively the team–driver unit executes strategic pit-stop responses.

To study this, we collect detailed career statistics for all drivers in our sample from *statsf1.com*, including their number of races, podiums, wins, points, and championships. For each season, we assign drivers the most up-to-date cumulative statistics available at its start, ensuring that improvements in skill or experience over time are reflected in the measures. Since no single observable fully captures "driver quality," we construct a composite driver-strength index using principal components analysis (PCA) applied to the above metrics along with average points per race, win share, and podium share. The first principal component loads positively and with comparable magnitude on all inputs (Panel A of Online Appendix Figure A11), making it a natural summary index of driver quality. The second component, by contrast, mixes positive and negative loadings

Table 9: Empirical responses with driver and team controls

Panel A: Responses of Leaders					
	Dependent variable: ΔLaps				
	(1) $X = \log(1 + \#\text{ of races})$	(2) $X = \mathbb{I}(\text{Winner})$	(3) $X = \mathbb{I}(\text{Champion})$	(4) $X = \mathbb{I}(\text{Top Driver})$	(5) $X = \mathbb{I}(\text{Top Team})$
X	0.027	−0.073	−0.104	−0.001	−0.004
	(0.026)	(0.063)	(0.063)	(0.106)	(0.139)
Lap × X	−0.071	0.086	0.127	−0.043	−0.066
	(0.071)	(0.156)	(0.154)	(0.280)	(0.329)
Circuit & Year FE	No	No	No	No	No
Other Controls	Yes	Yes	Yes	Yes	Yes
Observations	139	139	139	139	139
Mean of Dep. Variable	0.158	0.158	0.158	0.158	0.158
Adj. R-squared	0.572	0.581	0.586	0.569	0.569
Panel B: Responses of Pursuers					
	Dependent variable: ΔLaps				
	(1) $X = \log(1 + \#\text{ of races})$	(2) $X = \mathbb{I}(\text{Winner})$	(3) $X = \mathbb{I}(\text{Champion})$	(4) $X = \mathbb{I}(\text{Top Driver})$	(5) $X = \mathbb{I}(\text{Top Team})$
X	0.017	−0.019	−0.061	0.069	−0.114
	(0.011)	(0.058)	(0.058)	(0.052)	(0.070)
Lap × X	−0.030	0.030	0.054	−0.148	0.285
	(0.026)	(0.145)	(0.168)	(0.136)	(0.169)
Circuit & Year FE	No	No	No	No	No
Other Controls	Yes	Yes	Yes	Yes	Yes
Observations	286	288	288	286	288
Mean of Dep. Variable	0.144	0.143	0.143	0.144	0.143
Adj. R-squared	0.449	0.448	0.457	0.451	0.455

Notes: This table examines whether leaders' and pursuers' pit-stop responses vary with driver or team characteristics. The dependent variable is ΔLaps, defined as the difference in pit-stop timing between the two drivers in a pair, measured in laps. Panel A reports responses of leaders, while Panel B reports responses of pursuers. Each column introduces a different measure of driver or team quality, denoted by X: the total number of races, an indicator for having at least one race win, an indicator for being a world champion, an indicator for being classified as a top driver based on the principal component analysis (PCA), and an indicator for belonging to a top team. The variable Lap denotes the opponent's pit-stop lap, and the interaction term Lap × X captures whether pit-stop responses vary systematically with driver or team characteristics. All specifications include team fixed effects as well as team-specific responses to Lap, and a common set of additional control variables. Controls include tire compounds, distances to the nearest drivers ahead of and behind the pair, and the position contested by the two drivers. The sample consists of close pit-stop interactions within one-stop strategies, following the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

and is difficult to interpret. The resulting driver-strength index correlates strongly and positively with each underlying metric (Panel B of Online Appendix Figure A11), further validating its interpretation as a general measure of driver ability.

We then estimate how these driver characteristics relate to the optimality of strategic responses, regressing the inter-stop gap on the initiating driver's pit-lap, various driver-level measures, and an interaction between the two. As before, we control for tire compound, proximity to other cars, and the contested position, but now we also include team-specific intercepts and team-specific slopes to control for hetorogeneity in teams' decisions. Because we include team-specific coefficients and because the resulting sample is relatively small, we omit circuit and season fixed effects. Table 9 reports results for five driver characteristics: (1) the logarithm of one plus the number of races, (2) an indicator for having ever won a race, (3) an indicator for having won a championship, (4) an indicator for being above the median of our driver-strength index, and (5) an indicator for racing for a top-five team in the constructor standings.

Panel A examines leaders' best responses to undercuts. None of the driver characteristics are statistically significant. However, for both the composite strength measure and the measure of experience, the coefficients have the expected signs—stronger and more experienced drivers exhibit more negative slopes, i.e., more optimal responses—but the standard errors are large. Panel B presents pursuers' responses to leaders' pit stops. Here, all coefficients also remain statistically insignificant, with negative signs for the measures of experience and strength.

Given that we are primarily interested in the role of driver ability, we further explore this relationship using alternative specifications based on the driver-strength index (Online Appendix Table A15). Column 1 replicates the median-split indicator. Column 2 replaces it with the continuous driver-strength quantile, and Column 3 uses the logarithm of one plus standardized strength, which accommodates skewness and zero values in the distribution. Columns 4–6 repeat these three specifications using a season-normalized index in which each underlying feature is standardized within season rather than across the full sample. Across all six specifications, although the signs of the coefficients are consistent with the notion that better drivers make more optimal decisions, the estimates remain imprecise and statistically insignificant.

The limited statistical precision is unsurprising. These specifications rely on within-team variation in driver quality after controlling for team-specific strategies, but top drivers tend to cluster in top teams, and variation in driver skill within a given team is often small. Moreover, the specification includes demanding controls—team fixed intercepts and team-specific response slopes—which substantially reduce the usable variation. Given these constraints, it is nevertheless notable that even after absorbing team-level decision rules, the estimated effects remain directionally consistent.

Taken together, these results indicate that teams are the primary contributors to the strategic performance of team–driver pairs. Teams exhibit persistent differences in their pit-stop decision rules, and those whose behavior more closely follows the model's best responses tend to achieve higher success rates in strategic duels. Within teams, although the coefficients on drivers' experience and performance measures have the expected signs, the lack of statistical precision does not allow us to conclude that drivers' characteristics play an important role in the optimality of their decisions.

8 Conclusion

This paper studies the centipede game in a natural environment in which the game tree is not explicitly specified and must be inferred from experience. While the centipede game has played a central role in experimental research on sequential rationality, existing evidence is largely based on laboratory settings in which the structure of the game and the associated payoffs are made fully transparent to participants. We show that a centipede-like strategic interaction arises in real-world sequential decision making and can be analyzed using standard game-theoretic tools once the relevant payoff components are empirically identified.

Using detailed data from Formula 1 races, we model pit-stop timing decisions between closely competing drivers as an implicitly formulated centipede game. We estimate the payoff-relevant probabilities governing overtaking incentives and use them to compute equilibrium strategies. We then test the model's implications in the data. First, we show that drivers do respond strategically to each other: pit-stop timing shifts systematically when a rival pits nearby, even after controlling for tire degradation and race conditions. Second, we find that, on average, drivers' reactions follow the directional patterns predicted by the model, but responses to leaders' stops are somewhat farther from the optimum. Third, we document substantial heterogeneity across teams that helps explain observed differences in the optimality of decisions. Some teams consistently choose pit-stop timings that are extremely close to the theoretical optimum, while others deviate substantially – especially in the case of overcuts – patterns that translate directly into differences in strategic success rates on track. Finally, after absorbing team-specific strategic tendencies, we do not find evidence that drivers' observable characteristics contribute significantly to the optimality of team–driver decisions.

Our results contribute to the literature on the centipede game by providing evidence on equilibrium behavior in a complex sequential interaction outside the laboratory. They suggest that, even when the strategic structure of a centipede game is not explicitly presented, equilibrium reasoning can emerge in environments characterized by repeated interaction, high stakes, and substantial heterogeneity in agents' abilities.

The results of the paper should be interpreted carefully. We do not compare behavior in implicit and explicit formulations of the same game, and therefore do not isolate the causal role of implicitness per se. Rather, the paper documents how a canonical sequential game can be identified and studied in a natural setting where players face centipede-like incentives without observing the game tree.

Funding

The authors gratefully acknowledge support from the Basic Research Program of the HSE University (HSE-BR-2025-082). The funding source had no involvement in the research.

References

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

Anderson, A., Rosen, J., Rust, J., and Wong, K. P. (2025). Disequilibrium play in tennis. *Journal of Political Economy*, 133(1), 190–251.

Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavio*r, 8(1), 6-19.

Aumann, R. J. (1998). On the centipede game. *Games and Economic Behavior*, 23(1), 97-105.

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

Bornstein, G., Kugler, T., and Ziegelmeyer, A. (2004). Individual and group decisions in the centipede game: Are groups more "rational" players? *Journal of Experimental Social Psychology*, 40(5), 599-605.

Chiappori, P. A., Levitt, S., and Groseclose, T. (2002). Testing mixed-strategy equilibria when players are heterogeneous: The case of penalty kicks in soccer. *American Economic Review*, 92(4), 1138-1151.

Deck, A., Deck, C., & Zhu, Z. (2014). Decision Making in a Sequential Game: The Case of Pitting in NASCAR. *Journal of Sports Economics*, 15(2), 132–149.

García-Pola, B., Iriberri, N., and Kovářík, J. (2020). Non-equilibrium play in centipede games. *Games and Economic Behavior*, 120, 391-433.

Heilmeier, A., Graf, M., and Lienkamp, M. (2018). A race simulation for strategy decisions in circuit motorsports. In 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2986–2993. IEEE.

Heilmeier, A., Graf, M., Betz, J., & Lienkamp, M. (2020). Application of Monte Carlo Methods to Consider Probabilistic Effects in a Race Simulation for Circuit Motorsport. *Applied Sciences*, 10(12), Article 12.

Heilmeier, A., Thomaser, A., Graf, M., and Betz, J. (2020). Virtual strategy engineer: Using artificial neural networks for making race strategy decisions in circuit motorsport. *Applied Sciences*, 10(21), 7805.

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

Kawagoe, T. and Takizawa, H. (2012). Level-k analysis of experimental centipede games. *Journal of Economic Behavior & Organization*, 82(2-3), 548-566.

Kovash, K. and Levitt, S. D. (2009). Professionals do not play minimax: Evidence from major league baseball and the national football league (No. w15347). National Bureau of Economic Research.

Krockow, E. M., Colman, A. M., and Pulford, B. D. (2016). Cooperation in repeated interactions: A systematic review of Centipede game experiments, 1992–2016. *European Review of Social Psychology*, 27(1), 231-282.

Levitt, S. D., List, J. A., and Reiley, D. H. (2010). What happens in the field stays in the field: Exploring whether professionals play minimax in laboratory experiments. *Econometrica*, 78(4), 1413-1434.

Levitt, S. D., List, J. A., and Sadoff, S. E. (2011). Checkmate: Exploring backward induction among chess players. *American Economic Review*, 101(2), 975-990.

McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica: Journal of the Econometric Society*, 803-836.

Mourão, P. (2017). The economics of motorsports: The case of Formula One. Springer.

Nagel, R. and Tang, F. F. (1998). Experimental results on the centipede game in normal form: an investigation on learning. *Journal of Mathematical Psychology*, 42(2-3), 356-384.

Palacios-Huerta, I. (2003). Professionals play minimax. *The Review of Economic Studies*, 70(2), 395-415.

Palacios-Huerta, I. and Volij, O. (2008). Experientia docet: Professionals play minimax in laboratory experiments. *Econometrica*, 76(1), 71-115.

Palacios-Huerta, I. and Volij, O. (2009). Field centipedes. *American Economic Review*, 99(4), 1619-1635.

Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and the chain-store paradox. *Journal of Economic Theory*, 25(1), 92-100.

Thompson, K. (2014). Pacing: Individual strategies for optimal performance. Human Kinetics.

Tulabandhula, T. and Rudin, C. (2014). Tire changes, fresh air, and yellow flags: challenges in predictive analytics for professional racing. *Big Data*, 2(2), 97-112.

Walker, M. and Wooders, J. (2001). Minimax play at Wimbledon. *American Economic Review*, 91(5), 1521-1538.

A Online Appendix

Table A1: Descriptive Statistics by Pit-Stop Interaction Type: alternative samples

Panel A: All strategies						
	Undercut		***Overcut***		***Same lap***	
	Mean	SD	Mean	SD	Mean	SD
$\mathbb{I}$***(Overtake):***						
– right after the stops	0.548	0.498	0.187	0.390	0.082	0.276
– in the end of the stint/race	0.420	0.494	0.403	0.491	0.218	0.414
Δ***Laps:***						
– between the stops	0.113	0.124	0.131	0.120	0.000	0.000
– after the stops	0.355	0.172	0.352	0.161	0.464	0.219
Δ***Tire:***						
– between the stops	0.285	0.128	0.290	0.118	0.009	0.022
– after the stops	0.113	0.123	0.131	0.120	0.009	0.022
Other characteristics:						
– Average distance (meters)	99.592	48.545	82.356	46.476	100.707	51.964
– Number of stops	1.846	0.759	1.792	0.791	1.924	0.870
– Stint	1.351	0.570	1.353	0.597	1.506	0.732
– ΔTeam's strength	−0.532	2.879	0.418	2.642	−0.394	2.550
Observations	436		662		170	
Panel B: Last stops						
	Undercut		***Overcut***		***Same lap***	
	Mean	SD	Mean	SD	Mean	SD
$\mathbb{I}$***(Overtake):***						
– right after the stops	0.502	0.501	0.217	0.413	0.093	0.291
– in the end of the stint/race	0.405	0.492	0.453	0.498	0.259	0.440
Δ***Laps:***						
– between the stops	0.133	0.146	0.153	0.131	0.000	0.000
– after the stops	0.442	0.166	0.417	0.155	0.561	0.199
Δ***Tire:***						
– between the stops	0.332	0.131	0.324	0.122	0.009	0.018
– after the stops	0.133	0.144	0.153	0.130	0.009	0.018
Other characteristics:						
– Average distance (meters)	102.541	47.595	84.961	47.338	105.187	51.383
– ΔTeam's strength	−0.445	2.749	0.422	2.577	−0.352	2.365
Observations	247		419		108	

Notes: This table reports descriptive statistics for key race-state variables and outcomes, separately for undercuts, overcuts, and same-lap pit stops. The unit of observation is a pair of drivers involved in a close pit-stop interaction, conditional on both drivers finishing the race and satisfying the data-cleaning procedure described in Section 4.3. Panel A reports statistics for the full sample of pit stops that satisfy these criteria. Panel B restricts attention to the last pit stops within each strategic interaction. An *undercut* occurs when the pursuer pits before the leader, an *overcut* when the pursuer pits after the leader, and *same-lap* when both drivers pit on the same lap. $\mathbb{I}$(Overtake) indicates whether the pursuer overtakes the leader either immediately after the pit stops or by the end of the stint (or race). ΔLaps measures the normalized difference in laps either between pit stops or after both drivers have stopped, while ΔTire captures the corresponding normalized tire age advantage. Additional characteristics include the average on-track distance between drivers, the number of pit stops, stint indicators, and the difference in team strength. Means and standard deviations are reported.

Table A2: Summary of Overtaking Probability Estimates across Strategies

	Undercuts			Overcuts			Same Lap
	(1) $Pr_u^1(l_L, l_P)$	(2) $Pr_u^2(l_L, l_P \mid O=1)$	(3) $Pr_u^2(l_L, l_P \mid O=0)$	(4) $Pr_o^1(l_L, l_P)$	(5) $Pr_o^2(l_L, l_P \mid O=1)$	(6) $Pr_o^2(l_L, l_P \mid O=0)$	(7) $Pr_s^2(l_L, l_P)$
ΔLaps	2.894*	−0.935	1.896	−1.317	0.725	3.978**	−0.516
	(1.483)	(1.987)	(1.917)	(1.576)	(0.809)	(1.837)	(4.051)
ΔTire	0.851	−2.203	1.119	−3.638*	0.379	4.113*	
	(1.199)	(2.805)	(1.426)	(1.899)	(0.416)	(2.306)	
ΔLaps^2	−4.643**	0.436	−1.464	0.756	−0.752	−3.316**	−0.333
	(2.087)	(1.837)	(1.538)	(1.338)	(0.833)	(1.374)	(2.962)
ΔTire^2	−0.304	0.149	−0.633	3.661*	0.135	−3.279	
	(1.426)	(3.345)	(0.701)	(2.093)	(0.221)	(2.085)	
ΔTire × ΔLaps	−0.383	3.033	−1.936	3.999	−1.579	−5.652	
	(2.427)	(3.950)	(2.662)	(3.637)	(1.627)	(3.850)	
Constant	0.035	1.154*	−0.479	0.989**	0.874***	−0.954	0.844
	(0.263)	(0.594)	(0.555)	(0.416)	(0.147)	(0.584)	(1.346)
Circuit	RE	RE	RE	RE	RE	RE	RE
Other Controls	No	No	No	No	No	No	No
Observations	143	68	75	254	63	191	57
Mean of Dep. Variable	0.476	0.662	0.067	0.248	0.984	0.288	0.316
Within R-squared	0.145	0.312	0.011	0.020	0.117	0.079	0.098

Notes: This table summarizes the main regression results reported in Tables 2–8 across different strategic scenarios. Columns (1)–(3) correspond to undercut strategies, columns (4)–(6) to overcut strategies, and column (7) to same-lap pit stops. For undercuts and overcuts, column (1) and (4) report estimates of the probability of an intermediate overtake immediately after the pit stops. Columns (2) and (5) report the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate overtake, while columns (3) and (6) report the probability of a late overtake conditional on no intermediate overtake. Column (7) reports the probability that the pursuer overtakes the leader by the end of the interaction following same-lap pit stops. All specifications include circuit-level random effects and a control for the difference in team strength between the two drivers. Each column corresponds to the preferred specification from the respective baseline table. Sample restrictions and data-cleaning procedures follow those described in Section 4.3 and in the notes to the corresponding tables. Standard errors clustered at the driver level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A3: Intermediate overtake probabilities for undercuts: alternative samples

Panel A: All strategies						
	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.419	0.565**	0.851***	1.110	0.858	0.910
	(0.267)	(0.241)	(0.262)	(1.218)	(1.112)	(0.998)
ΔTire	−0.073	0.105	0.385	−0.786	−0.793	−0.458
	(0.217)	(0.229)	(0.293)	(0.885)	(0.814)	(0.740)
ΔLaps^2				−3.652**	−3.091*	−2.700*
				(1.773)	(1.665)	(1.460)
ΔTire^2				0.445	0.633	0.454
				(1.018)	(0.928)	(0.899)
ΔTire × ΔLaps				4.013*	4.535**	4.927**
				(2.283)	(2.132)	(1.866)
Constant	0.541***	0.448***	0.379***	0.612***	0.570***	0.506***
	(0.091)	(0.091)	(0.112)	(0.201)	(0.184)	(0.172)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	435	435	432	435	435	432
Mean of Dep. Variable	0.549	0.549	0.549	0.549	0.549	0.549
Within R-squared	0.044	0.058	0.068	0.103	0.109	0.119
Panel B: Last stops						
ΔLaps	0.652**	0.657***	0.800***	1.872	1.072	1.167
	(0.255)	(0.229)	(0.245)	(1.574)	(1.374)	(1.296)
ΔTire	0.733***	0.854***	0.858***	2.455*	2.242*	2.567**
	(0.218)	(0.222)	(0.306)	(1.439)	(1.257)	(1.230)
ΔLaps^2				−2.972	−2.025	−1.880
				(1.969)	(1.702)	(1.459)
ΔTire^2				−2.626	−2.395*	−2.851**
				(1.553)	(1.328)	(1.198)
ΔTire × ΔLaps				1.322	2.560	2.710
				(2.762)	(2.468)	(2.268)
Constant	0.186*	0.142	0.166	−0.149	−0.086	−0.112
	(0.095)	(0.086)	(0.133)	(0.322)	(0.289)	(0.316)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	246	246	244	246	246	244
Mean of Dep. Variable	0.504	0.504	0.500	0.504	0.504	0.500
Within R-squared	0.071	0.101	0.131	0.129	0.155	0.193

Notes: This table reports estimates of the probability that the pursuer overtakes the leader immediately after pit stops when employing an undercut strategy. The dependent variable is a dummy equal to one if the pursuer is ahead one lap after the leader's stop. ΔLaps measures the share of the race distance between the pursuer's and the leader's pit stops, while ΔTire captures the pursuer's tire-age advantage during this interval. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential undercut opportunities, while Panel B restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A4: Final overtake probability conditional on a successful undercut: alternative samples

Panel A: All strategies						
	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.118	−0.164	−0.169	−0.269	−0.215	−0.393
	(0.196)	(0.195)	(0.238)	(0.830)	(0.825)	(0.981)
ΔTire	−0.305	−0.407	−0.477	0.523	0.728	1.009
	(0.318)	(0.311)	(0.316)	(0.952)	(0.878)	(1.036)
ΔLaps^2				0.613	0.523	0.773
				(1.114)	(1.120)	(1.233)
ΔTire^2				−0.208	−0.673	−1.573
				(1.347)	(1.281)	(1.369)
ΔTire × ΔLaps				−2.639	−2.906*	−2.776
				(1.731)	(1.711)	(2.110)
Constant	0.780***	0.823***	0.835***	0.748***	0.749***	0.768***
	(0.092)	(0.091)	(0.103)	(0.144)	(0.135)	(0.192)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	238	238	235	238	238	235
Mean of Dep. Variable	0.702	0.702	0.702	0.702	0.702	0.702
Within R-squared	0.085	0.099	0.104	0.101	0.115	0.120
Panel B: Last stops						
ΔLaps	−0.274	−0.211	0.113	0.369	0.383	−0.332
	(0.330)	(0.320)	(0.512)	(1.235)	(1.243)	(1.640)
ΔTire	−0.686*	−0.681**	−0.710	1.301	1.503	1.265
	(0.343)	(0.343)	(0.418)	(1.435)	(1.395)	(1.659)
ΔLaps^2				−0.115	−0.084	0.990
				(1.465)	(1.470)	(1.870)
ΔTire^2				−1.445	−1.743	−1.948
				(1.887)	(1.838)	(1.934)
ΔTire × ΔLaps				−4.046*	−4.199**	−3.043
				(2.004)	(2.001)	(2.768)
Constant	0.927***	0.908***	0.782***	0.655**	0.635**	0.714*
	(0.169)	(0.167)	(0.241)	(0.292)	(0.295)	(0.399)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	123	123	118	123	123	118
Mean of Dep. Variable	0.715	0.715	0.712	0.715	0.715	0.712
Within R-squared	0.067	0.135	0.135	0.088	0.155	0.155

Notes: This table reports estimates of the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate undercut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the former leader. ΔLaps measures the share of the race distance from the leader's stop to the end of the interaction, while ΔTire captures the leader's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential undercut opportunities, conditional on the pursuer overtaking the leader immediately after the stops, while Panel B additionally restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Final overtake probability conditional on an unsuccessful undercut: alternative samples

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
Panel A: All strategies						
ΔLaps	0.059	0.059	0.176	0.519**	0.519**	0.750
	(0.081)	(0.081)	(0.175)	(0.222)	(0.222)	(0.503)
ΔTire	−0.167**	−0.167***	−0.205	0.199	0.199	−0.328
	(0.065)	(0.065)	(0.178)	(0.427)	(0.427)	(0.871)
ΔLaps^2				−0.471*	−0.471*	−0.691
				(0.271)	(0.271)	(0.460)
ΔTire^2				−0.323	−0.323	0.249
				(0.622)	(0.622)	(0.974)
ΔTire × ΔLaps				−0.706	−0.706	−0.033
				(0.536)	(0.536)	(1.253)
Constant	0.100**	0.100***	0.101	0.003	0.003	0.008
	(0.038)	(0.038)	(0.084)	(0.051)	(0.051)	(0.152)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	196	196	192	196	196	192
Mean of Dep. Variable	0.077	0.077	0.078	0.077	0.077	0.078
Within R-squared	0.098	0.082	0.100	0.103	0.087	0.110
Panel B: Last stops						
ΔLaps	−0.077	−0.077	0.051	−0.216	−0.216	−1.616
	(0.170)	(0.170)	(0.355)	(1.364)	(1.364)	(2.761)
ΔTire	−0.338**	−0.338**	−0.465*	−1.163	−1.163	−4.187
	(0.133)	(0.133)	(0.269)	(1.317)	(1.317)	(2.704)
ΔLaps^2				0.024	0.024	1.095
				(1.064)	(1.064)	(2.134)
ΔTire^2				0.747	0.747	2.708
				(1.016)	(1.016)	(1.955)
ΔTire × ΔLaps				1.277	1.277	6.685
				(2.034)	(2.034)	(4.594)
Constant	0.185*	0.185*	0.155	0.260	0.260	0.703
	(0.100)	(0.100)	(0.187)	(0.419)	(0.419)	(0.820)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	122	122	119	122	122	119
Mean of Dep. Variable	0.090	0.090	0.092	0.090	0.090	0.092
Within R-squared	0.062	0.040	0.070	0.068	0.060	0.131

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction conditional on an unsuccessful intermediate undercut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. ΔLaps measures the share of the race distance from the leader's stop to the end of the interaction, while ΔTire captures the leader's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential undercut opportunities, conditional on the pursuer not overtaking the leader immediately after the stops, while Panel B additionally restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Intermediate overtake probabilities for overcuts: alternative samples

	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
Panel A: All strategies						
ΔLaps	0.408**	0.297*	0.085	0.105	0.224	0.598
	(0.159)	(0.171)	(0.212)	(0.827)	(0.802)	(0.783)
ΔTire	−0.056	−0.073	−0.161	−1.461*	−1.401*	−1.445*
	(0.172)	(0.173)	(0.227)	(0.850)	(0.819)	(0.783)
ΔLaps^2				−0.097	−0.451	−1.130
				(0.866)	(0.857)	(0.899)
ΔTire^2				1.884*	1.818*	1.980*
				(1.078)	(1.047)	(1.008)
ΔTire × ΔLaps				1.330	1.032	−0.041
				(2.244)	(2.141)	(2.006)
Constant	0.143**	0.180***	0.196**	0.362**	0.375**	0.337**
	(0.061)	(0.064)	(0.086)	(0.159)	(0.154)	(0.155)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	661	661	661	661	661	661
Mean of Dep. Variable	0.188	0.188	0.188	0.188	0.188	0.188
Within R-squared	0.034	0.019	0.026	0.043	0.031	0.041
Panel B: Last stops						
ΔLaps	0.235	0.099	−0.026	−1.633	−1.395	−0.908
	(0.224)	(0.256)	(0.262)	(1.112)	(1.094)	(1.064)
ΔTire	−0.374	−0.400	−0.475	−3.791***	−3.701***	−4.103***
	(0.240)	(0.263)	(0.293)	(1.190)	(1.215)	(1.230)
ΔLaps^2				1.247	0.884	0.170
				(1.058)	(1.056)	(1.050)
ΔTire^2				3.992***	3.942***	4.650***
				(1.406)	(1.427)	(1.401)
ΔTire × ΔLaps				4.588	4.041	2.971
				(2.770)	(2.703)	(2.597)
Constant	0.291***	0.329***	0.329***	0.951***	0.936***	0.932***
	(0.096)	(0.107)	(0.112)	(0.243)	(0.249)	(0.262)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	418	418	417	418	418	417
Mean of Dep. Variable	0.218	0.218	0.216	0.218	0.218	0.216
Within R-squared	0.054	0.050	0.062	0.093	0.087	0.113

Notes: This table reports estimates of the probability that the pursuer overtakes the leader immediately after pit stops when employing an overcut strategy. The dependent variable is a dummy equal to one if the pursuer is ahead one lap after his own stop. ΔLaps measures the share of the race distance between the leader's and the pursuer's pit stops, while ΔTire captures the leader's tire-age advantage during this interval. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential overcut opportunities, while Panel B restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Final overtake probability conditional on a successful overcut: alternative samples

Panel A: All strategies						
	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	−0.047	−0.047	−0.096	−0.067	−0.067	0.008
	(0.078)	(0.078)	(0.141)	(0.423)	(0.423)	(0.594)
ΔTire	0.075	0.075	−0.024	0.413	0.413	0.127
	(0.079)	(0.079)	(0.144)	(0.348)	(0.348)	(0.551)
ΔLaps^2				0.130	0.130	−0.004
				(0.479)	(0.479)	(0.529)
ΔTire^2				−0.117	−0.117	0.477
				(0.378)	(0.378)	(0.618)
ΔTire × ΔLaps				−0.833	−0.833	−1.151
				(1.119)	(1.119)	(1.646)
Constant	0.975***	0.975***	0.993***	0.956***	0.956***	0.968***
	(0.045)	(0.045)	(0.058)	(0.064)	(0.064)	(0.128)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	124	124	118	124	124	118
Mean of Dep. Variable	0.976	0.976	0.975	0.976	0.976	0.975
Within R-squared	0.017	0.009	0.048	0.029	0.024	0.070
Panel B: Last stops						
ΔLaps	−0.080	−0.080	−0.211	0.097	0.097	−0.650
	(0.125)	(0.125)	(0.204)	(0.933)	(0.933)	(1.208)
ΔTire	0.017	0.017	−0.199	0.471	0.471	0.149
	(0.120)	(0.120)	(0.165)	(0.447)	(0.447)	(0.686)
ΔLaps^2				−0.053	−0.053	0.527
				(0.965)	(0.965)	(1.158)
ΔTire^2				−0.069	−0.069	0.320
				(0.214)	(0.214)	(0.412)
ΔTire × ΔLaps				−1.196	−1.196	−1.435
				(1.624)	(1.624)	(2.383)
Constant	1.007***	1.007***	1.060***	0.944***	0.944***	1.146***
	(0.080)	(0.080)	(0.098)	(0.159)	(0.159)	(0.234)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	91	91	84	91	91	84
Mean of Dep. Variable	0.978	0.978	0.976	0.978	0.978	0.976
Within R-squared	0.012	0.010	0.073	0.029	0.040	0.120

Notes: This table reports estimates of the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate overcut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the former leader. ΔLaps measures the share of the race distance from the pursuer's stop to the end of the interaction, while ΔTire captures the pursuer's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential overcut opportunities, conditional on the pursuer overtaking the leader immediately after the stops, while Panel B additionally restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A8: Final overtake probability conditional on an unsuccessful overcut: alternative samples

Panel A: All strategies						
	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.149	0.178	0.307*	1.649***	1.651***	1.803**
	(0.125)	(0.121)	(0.153)	(0.495)	(0.506)	(0.676)
ΔTire	0.374*	0.406*	0.433*	1.940***	1.949***	2.003***
	(0.217)	(0.214)	(0.235)	(0.674)	(0.673)	(0.722)
ΔLaps^2				−1.690***	−1.666***	−1.676**
				(0.453)	(0.468)	(0.715)
ΔTire^2				−2.445***	−2.410***	−2.241***
				(0.831)	(0.813)	(0.750)
ΔTire × ΔLaps				−1.491	−1.531	−1.904
				(1.715)	(1.721)	(1.820)
Constant	0.163***	0.144***	0.119*	−0.180	−0.186*	−0.220
	(0.057)	(0.054)	(0.063)	(0.108)	(0.109)	(0.139)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	537	537	537	537	537	537
Mean of Dep. Variable	0.272	0.272	0.272	0.272	0.272	0.272
Within R-squared	0.076	0.066	0.071	0.098	0.083	0.088
Panel B: Last stops						
ΔLaps	−0.059	−0.071	−0.178	2.782***	2.940***	3.276***
	(0.164)	(0.171)	(0.269)	(0.822)	(0.838)	(1.005)
ΔTire	0.217	0.195	0.190	2.707**	2.872**	3.175***
	(0.297)	(0.300)	(0.324)	(1.293)	(1.252)	(1.146)
ΔLaps^2				−2.711***	−2.858***	−3.180***
				(0.677)	(0.719)	(0.945)
ΔTire^2				−2.570*	−2.700**	−2.697**
				(1.383)	(1.344)	(1.183)
ΔTire × ΔLaps				−3.743	−4.176*	−4.977**
				(2.436)	(2.338)	(2.193)
Constant	0.293***	0.303***	0.363***	−0.426*	−0.455**	−0.511**
	(0.087)	(0.095)	(0.130)	(0.225)	(0.224)	(0.244)
Circuit	No	RE	FE	No	RE	FE
Other Controls	No	No	Yes	No	No	Yes
Observations	327	327	327	327	327	327
Mean of Dep. Variable	0.309	0.309	0.309	0.309	0.309	0.309
Within R-squared	0.077	0.062	0.065	0.102	0.090	0.092

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction conditional on an unsuccessful intermediate overcut. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. ΔLaps measures the share of the race distance from the pursuer's stop to the end of the interaction, while ΔTire captures the pursuer's tire-age advantage during this period. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic and interaction terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all potential overcut opportunities, conditional on the pursuer not overtaking the leader immediately after the stops, while Panel B additionally restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A9: Final overtaking probability under same-lap pit stops: alternative samples

Panel A: All strategies						
	(1) Overtake	(2) Overtake	(3) Overtake	(4) Overtake	(5) Overtake	(6) Overtake
ΔLaps	0.207**	0.207**	0.408*	1.227***	1.227***	0.480
	(0.100)	(0.100)	(0.203)	(0.405)	(0.405)	(0.671)
ΔLaps^2				−1.055**	−1.055**	−0.083
				(0.412)	(0.412)	(0.768)
Constant	0.127**	0.127**	0.045	−0.068	−0.068	0.033
	(0.054)	(0.054)	(0.093)	(0.070)	(0.070)	(0.132)
Circuit	No	RE	FE	No	RE	FE
Other Control	No	No	Yes	No	No	Yes
Observations	170	170	168	170	170	168
Mean of Dep. Variable	0.218	0.218	0.220	0.218	0.218	0.220
Within R-squared	0.019	0.011	0.031	0.038	0.016	0.031
Panel B: Last stops						
ΔLaps	0.008	0.008	0.009	2.222**	2.222**	0.491
	(0.144)	(0.144)	(0.459)	(0.878)	(0.878)	(2.341)
ΔLaps^2				−1.992**	−1.992**	−0.453
				(0.781)	(0.781)	(2.080)
Constant	0.262**	0.262***	0.234	−0.274	−0.274	0.126
	(0.095)	(0.095)	(0.243)	(0.211)	(0.211)	(0.590)
Circuit	No	RE	FE	No	RE	FE
Other Control	No	No	Yes	No	No	Yes
Observations	108	108	106	108	108	106
Mean of Dep. Variable	0.259	0.259	0.264	0.259	0.259	0.264
Within R-squared	0.013	0.024	0.097	0.038	0.031	0.097

Notes: This table reports estimates of the probability that the pursuer overtakes the leader by the end of the interaction when both drivers pit on the same lap. The dependent variable is a dummy equal to one if the pursuer finishes the interaction (either the end of the race or the next pit stop by one of the drivers) ahead of the leader. Because pit stops occur on the same lap, ΔTire is equal to zero by construction, and ΔLaps measures the share of the race distance from the joint stop to the end of the interaction. The regressions therefore include ΔLaps and its square to allow for non-linear effects. Columns (1)–(3) report linear specifications, while columns (4)–(6) allow for quadratic terms. Columns (2) and (5) include circuit-level random effects; columns (3) and (6) include circuit fixed effects and additional controls for season and tire compounds. All specifications control for the difference in team strength between the two drivers. Panel A includes all same-lap pit stops, while Panel B additionally restricts the sample to the last pit stops of both drivers. Both panels follow the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A10: Logit and Probit Estimates of Overtaking Probabilities

	Undercuts			Overcuts			Same Lap
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$Pr_u^1(l_L, l_P)$	$Pr_u^2(l_L, l_P \mid O=1)$	$Pr_u^2(l_L, l_P \mid O=0)$	$Pr_o^1(l_L, l_P)$	$Pr_o^2(l_L, l_P \mid O=1)$	$Pr_o^2(l_L, l_P \mid O=0)$	$Pr_s^2(l_L, l_P)$
Panel A: Logit Model							
ΔLaps	19.219**	−8.008	2376.108***	−6.780	−20581.862	28.911**	13.146**
	(9.143)	(11.551)	(761.411)	(7.141)	(.)	(14.054)	(5.495)
ΔTire	8.243	−18.956	16517.757***	−19.405**	1669.206	23.286	
	(8.612)	(18.344)	(5126.495)	(9.400)	(.)	(16.938)	
ΔLaps^2	−28.319**	3.736	−1411.181***	3.477	15915.898	−25.302**	−11.697**
	(13.548)	(8.924)	(465.161)	(5.909)	(.)	(10.668)	(4.709)
ΔTire^2	−5.996	8.727	−67538.050***	19.383*	−5141.647	−16.631	
	(9.580)	(21.958)	(21727.844)	(10.347)	(.)	(14.131)	
ΔTire × ΔLaps	−11.107	24.792	−19786.223***	21.921	−2815.376	−34.299	
	(13.195)	(24.440)	(6251.490)	(17.093)	(.)	(28.354)	
Constant	−3.200	4.817	−996.988***	2.737	6760.095	−9.376**	−4.298***
	(1.982)	(3.989)	(311.321)	(1.959)	(.)	(4.468)	(1.445)
Circuit	No	No	No	No	No	No	No
Other Controls	No	No	No	No	No	No	No
Observations	143	68	75	254	63	191	108
Mean of Dep. Variable	0.476	0.662	0.067	0.248	0.984	0.288	0.259
Pseudo R-squared	0.113	0.184	0.555	0.053	1.000	0.088	0.035
Panel B: Probit Model							
ΔLaps	11.322**	−4.754	1439.323***	−4.124	−6171.590	16.909**	−0.173
	(5.091)	(7.301)	(481.784)	(4.297)	(.)	(7.940)	(11.336)
ΔTire	5.065	−11.686	10067.918***	−11.428**	9475.423	14.029	
	(5.266)	(11.455)	(3220.135)	(5.530)	(.)	(9.517)	
ΔLaps^2	−16.382**	2.105	−853.950***	2.224	7279.638	−14.750**	−1.974
	(7.081)	(5.715)	(293.722)	(3.652)	(.)	(6.016)	(8.305)
ΔTire^2	−3.703	5.948	−41319.849***	11.304*	−2309.870	−10.218	
	(5.900)	(13.294)	(13398.526)	(6.052)	(.)	(8.032)	
ΔTire × ΔLaps	−6.513	14.813	−12056.551***	12.963	−16601.931	−20.501	
	(7.853)	(15.346)	(3937.451)	(10.169)	(.)	(16.000)	
Constant	−1.947	2.939	−604.432***	1.619	1216.991	−5.528**	0.623
	(1.186)	(2.489)	(197.059)	(1.173)	(.)	(2.515)	(3.750)
Circuit	No	Yes	Yes	No	Yes	Yes	FE
Other Controls	No	No	No	No	No	No	No
Observations	143	68	75	254	63	191	57
Mean of Dep. Variable	0.476	0.662	0.067	0.248	0.984	0.288	0.316
Pseudo R-squared	0.111	0.188	0.562	0.053	1.000	0.088	0.074

Notes: This table reports logit and probit estimates of the overtaking probabilities analyzed in the main text. Panel A presents logit estimates, while Panel B presents probit estimates. Columns correspond to the same strategic scenarios and outcomes as in Table A2. For undercuts and overcuts, columns (1) and (4) report the probability of an intermediate overtake immediately after the pit stops. Columns (2) and (5) report the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate overtake, while columns (3) and (6) report the probability of a late overtake conditional on no intermediate overtake. Column (7) reports the probability that the pursuer overtakes the leader by the end of the interaction following same-lap pit stops. All specifications control for the difference in team strength between the two drivers. Sample restrictions and data-cleaning procedures follow those described in Section 4.3 and in the corresponding baseline tables. Standard errors clustered at the driver level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Robustness to Initial Distance between Drivers

	Undercuts			Overcuts			Same Lap
Panel A: Smaller distance (< 1.5 seconds)							
	(1) $Pr_u^1(l_L, l_P)$	(2) $Pr_u^2(l_L, l_P \mid O = 1)$	(3) $Pr_u^2(l_L, l_P \mid O = 0)$	(4) $Pr_o^1(l_L, l_P)$	(5) $Pr_o^2(l_L, l_P \mid O = 1)$	(6) $Pr_o^2(l_L, l_P \mid O = 0)$	(7) $Pr_s^2(l_L, l_P)$
ΔLaps	2.830	2.634	0.000	−0.007	0.751	3.245	2.331
	(2.262)	(6.882)	(.)	(1.989)	(0.890)	(2.894)	(6.101)
ΔTire	2.403	7.361	0.000	−2.999	0.479	3.368	
	(2.616)	(9.946)	(.)	(2.339)	(0.482)	(2.875)	
ΔLaps^2	−4.076	−1.012	0.000	−0.023	−0.788	−2.846	−2.246
	(2.935)	(4.618)	(.)	(1.554)	(0.933)	(2.286)	(4.635)
ΔTire^2	−1.620	−8.496	0.000	3.698	0.080	−2.445	
	(4.224)	(8.861)	(.)	(2.459)	(0.249)	(2.525)	
ΔTire × ΔLaps	−2.794	−10.080	0.000	0.991	−1.898	−4.783	
	(3.814)	(12.413)	(.)	(4.804)	(1.958)	(5.009)	
Constant	−0.116	−0.497	0.000	0.770	0.877***	−0.662	−0.099
	(0.385)	(2.556)	(.)	(0.523)	(0.152)	(0.866)	(1.910)
Circuit	RE	RE	RE	RE	RE	RE	RE
Other Controls	No	No	No	No	No	No	No
Observations	73	45	28	165	47	118	25
Mean of Dep. Variable	0.616	0.689	0.107	0.285	0.979	0.339	0.400
Within R-squared	0.150	0.439	0.000	0.032	0.151	0.106	0.214
Panel B: Larger distance (< 5 seconds)							
ΔLaps	4.041***	1.690	1.647*	−0.707	0.684	3.508**	−0.657
	(1.339)	(1.707)	(0.849)	(1.539)	(0.760)	(1.614)	(3.703)
ΔTire	−0.892	2.028	1.233**	−2.813	0.403	3.911*	
	(1.061)	(1.846)	(0.591)	(1.784)	(0.437)	(2.177)	
ΔLaps^2	−6.606***	−1.455	−1.188*	0.320	−0.708	−2.937**	−0.050
	(2.009)	(1.622)	(0.689)	(1.360)	(0.780)	(1.285)	(2.702)
ΔTire^2	1.665	−3.178	−0.756*	2.708	0.070	−3.239	
	(1.216)	(2.133)	(0.404)	(1.963)	(0.187)	(2.101)	
ΔTire × ΔLaps	−0.984	−3.214	−2.324*	2.328	−1.518	−5.269	
	(1.863)	(3.458)	(1.203)	(3.464)	(1.561)	(3.661)	
Constant	0.231	0.201	−0.447*	0.827**	0.877***	−0.818	0.753
	(0.231)	(0.488)	(0.243)	(0.392)	(0.143)	(0.498)	(1.230)
Circuit	RE	RE	RE	RE	RE	RE	RE
Other Controls	No	No	No	No	No	No	No
Observations	219	81	138	303	69	234	70
Mean of Dep. Variable	0.370	0.617	0.065	0.228	0.986	0.291	0.271
Within R-squared	0.182	0.262	0.027	0.020	0.104	0.108	0.083

Notes: This table examines the robustness of the baseline results to the initial time gap between the leader and the pursuer at the moment of the first pit stop. Panel A restricts the sample to interactions in which the time gap between the two drivers is smaller than 1.5 seconds, while Panel B considers interactions with a time gap smaller than 5 seconds. Columns correspond to the same strategic scenarios and outcome definitions as in Table A2. For undercuts and overcuts, columns (1) and (4) report the probability of an intermediate overtake immediately after the pit stops. Columns (2) and (5) report the probability that the pursuer remains ahead by the end of the interaction conditional on a successful intermediate overtake, while columns (3) and (6) report the probability of a late overtake conditional on no intermediate overtake. Column (7) reports the probability that the pursuer overtakes the leader by the end of the interaction following same-lap pit stops. All specifications include circuit-level random effects and control for the difference in team strength between the two drivers. Sample restrictions follow the data-cleaning procedure described in Section 4.3, with the additional restriction on the initial time gap between drivers. Standard errors clustered at the driver level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A12: Close Pit-Stop Interactions and Pit-Stop Timing

	Dependent variable: First stop			Dependent variable: Last stop		
	(1) All stops	(2) Undercuts	(3) Overcuts	(4) All stops	(5) Undercuts	(6) Overcuts
Distance < 3 seconds	−0.035***	−0.031**	−0.032**	−0.021**	−0.026**	−0.013
	(0.011)	(0.012)	(0.014)	(0.009)	(0.012)	(0.009)
Constant	0.501***	0.531***	0.496***	0.580***	0.561***	0.559***
	(0.016)	(0.024)	(0.020)	(0.015)	(0.020)	(0.019)
Circuit & Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	720	300	339	720	300	339
Mean of Dep. Variable	0.363	0.367	0.367	0.496	0.506	0.525
Within R-squared	0.348	0.461	0.374	0.383	0.529	0.423

Notes: This table reports panel regressions relating pit-stop timing to close pit-stop interactions. The unit of observation is a pair of drivers who each make exactly one pit stop in a race. The dependent variable is the pit-stop lap expressed as the share of race distance completed at the time of the stop. Columns (1)–(3) report results for the first pit stop within the pair, while Columns (4)–(6) report results for the last pit stop. Columns labeled All stops pool all interactions, while columns labeled Undercuts and Overcuts restrict the sample to the corresponding strategic interactions. The indicator "Distance < 3 seconds" equals one if the two drivers are within three seconds of each other on track at the time of the first pit stop. All specifications include circuit and year fixed effects as well as additional control variables, including tire compounds, distances to the nearest drivers ahead of and behind the pair, and the track position contested by the two drivers. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A13: Empirical Responses in Pit-Stop Timing

	Dependent variable: ΔLaps					
	(1) Leader	(2) Leader	(3) Leader	(4) Pursuer	(5) Pursuer	(6) Pursuer
Lap	−0.460***	−0.523***	−0.621***	−0.282***	−0.317***	−0.440***
	(0.157)	(0.151)	(0.108)	(0.079)	(0.097)	(0.088)
Constant	0.311***	0.333***	0.325***	0.236***	0.248***	0.260***
	(0.061)	(0.051)	(0.057)	(0.032)	(0.033)	(0.040)
Circuit & Year FE	No	Yes	Yes	No	Yes	Yes
Other Controls	No	No	Yes	No	No	Yes
Observations	143	136	132	311	310	288
Mean of Dep. Variable	0.155	0.156	0.159	0.141	0.141	0.143
Within R-squared	0.143	0.159	0.626	0.060	0.059	0.426

Notes: This table reports panel regressions describing empirical responses in pit-stop timing within pairs of drivers. The dependent variable, ΔLaps, measures the normalized difference in pit-stop laps between the two drivers in a pair. Columns (1)–(3) relate the leader's pit-stop timing to the pursuer's pit-stop decision, while Columns (4)–(6) relate the pursuer's pit-stop timing to the leader's decision. The variable Lap denotes the opponent's pit-stop lap, expressed as the share of race distance completed at the time of the stop. Specifications progressively add circuit and year fixed effects and additional control variables. Controls include tire compounds, distances to the nearest drivers ahead of and behind the pair, and the position contested by the two drivers. The sample consists of close pit-stop interactions within one-stop strategies, following the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A14: Empirical Responses Using Alternative Outcome Measures

Panel A: Responses of Leaders

	Dependent variable: ΔLaps			
	(1)	(2)	(3)	(4)
	$Pr(\Delta Laps \leq 1)$	$Pr(\Delta Laps \leq 3)$	$Pr(\Delta Laps \leq 5)$	$Pr(\Delta Laps \leq 10)$
Lap	0.592	0.776***	0.901***	1.201***
	(0.457)	(0.262)	(0.201)	(0.288)
Constant	−0.150	0.091	0.288	0.536***
	(0.343)	(0.180)	(0.168)	(0.167)
Circuit & Year FE	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes
Observations	132	132	132	132
Mean of Dep. Variable	0.288	0.424	0.492	0.652
Within R-squared	0.296	0.459	0.425	0.505

Panel B: Responses of Pursuers

	Dependent variable: ΔLaps			
	(1)	(2)	(3)	(4)
	$Pr(\Delta Laps \leq 1)$	$Pr(\Delta Laps \leq 3)$	$Pr(\Delta Laps \leq 5)$	$Pr(\Delta Laps \leq 10)$
Lap	0.134	0.724**	0.933**	0.941***
	(0.326)	(0.338)	(0.354)	(0.282)
Constant	0.362**	0.241	0.344*	0.512***
	(0.140)	(0.160)	(0.175)	(0.164)
Circuit & Year FE	Yes	Yes	Yes	Yes
Other Controls	Yes	Yes	Yes	Yes
Observations	288	288	288	288
Mean of Dep. Variable	0.260	0.410	0.517	0.656
Within R-squared	0.101	0.181	0.263	0.285

Notes: This table reports panel regressions describing empirical responses in pit-stop timing using alternative outcome measures. The unit of observation is a pair of drivers who each make exactly one pit stop in a race. The dependent variable is an indicator for whether the absolute difference in pit-stop laps between the two drivers, ΔLaps, does not exceed a given threshold. Columns (1)–(4) in each panel consider thresholds of one, three, five, and ten laps, respectively. Panel A reports responses of leaders to the pursuer's pit-stop timing, while Panel B reports responses of pursuers to the leader's pit-stop timing. The variable Lap denotes the opponent's pit-stop lap, expressed as the share of race distance completed at the time of the stop. All specifications include circuit and year fixed effects as well as additional control variables. Controls include tire compounds, distances to the nearest drivers ahead of and behind the pair, and the position contested by the two drivers. The sample consists of close pit-stop interactions within one-stop strategies, following the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

Table A15: Empirical responses with alternative strength measures

Panel A: Responses of Leaders

	Dependent variable: ΔLaps					
	Strength, standardized features			Strength, standardized by year features		
	(1)	(2)	(3)	(4)	(5)	(6)
	$X = \mathbb{I}$(Top Driver)	$X =$ Percentile	$\log(1+\text{Strength})$	$X = \mathbb{I}$(Top Driver)	$X =$ Percentile	$\log(1+\text{Strength})$
X	−0.001	0.026	0.004	−0.001	0.044	0.004
	(0.106)	(0.175)	(0.038)	(0.106)	(0.181)	(0.045)
Lap × X	−0.043	−0.178	−0.041	−0.043	−0.247	−0.056
	(0.280)	(0.475)	(0.103)	(0.280)	(0.496)	(0.123)
Circuit & Year FE	No	No	No	No	No	No
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	139	139	139	139	139	139
Mean of Dep. Variable	0.158	0.158	0.158	0.158	0.158	0.158
Adj. R-squared	0.569	0.570	0.572	0.569	0.571	0.574

Panel B: Responses of Pursuers

	Dependent variable: ΔLaps					
	Strength, standardized features			Strength, standardized by year features		
	(1)	(2)	(3)	(4)	(5)	(6)
	$X = \mathbb{I}$(Top Driver)	$X =$ Percentile	$\log(1+\text{Strength})$	$X = \mathbb{I}$(Top Driver)	$X =$ Percentile	$\log(1+\text{Strength})$
X	0.069	0.048	0.005	0.077	0.041	0.004
	(0.052)	(0.091)	(0.019)	(0.062)	(0.086)	(0.022)
Lap × X	−0.148	−0.130	−0.019	−0.167	−0.113	−0.020
	(0.136)	(0.224)	(0.048)	(0.156)	(0.212)	(0.056)
Circuit & Year FE	No	No	No	No	No	No
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	286	286	286	286	286	286
Mean of Dep. Variable	0.144	0.144	0.144	0.144	0.144	0.144
Adj. R-squared	0.451	0.446	0.446	0.452	0.446	0.446

Notes: This table examines the robustness of leaders' and pursuers' pit-stop responses to alternative definitions of driver strength. The dependent variable is ΔLaps, defined as the difference in pit-stop timing between the two drivers in a pair, measured in laps. Panel A reports responses of leaders, while Panel B reports responses of pursuers. Each panel reports specifications using alternative measures of driver strength. Columns (1)–(3) use strength measures constructed from standardized features computed over the full sample, while columns (4)–(6) use strength measures based on features standardized separately by year. Across specifications, driver strength is measured using an indicator for being classified as a top driver, the percentile rank of the strength index, or $\log(1 + \text{Strength})$. The variable Lap denotes the opponent's pit-stop lap, and the interaction term Lap × X captures whether pit-stop responses vary with the driver's relative strength under alternative definitions. All specifications include team fixed effects, team-specific responses to Lap, and a common set of additional control variables. Controls include tire compounds, distances to the nearest drivers ahead of and behind the pair, and the position contested by the two drivers. The sample consists of close pit-stop interactions within one-stop strategies, following the data-cleaning procedure described in Section 4.3. Standard errors clustered at the circuit level are reported in parentheses. The p-values are denoted as follows: *** p<0.01, ** p<0.05, * p<0.1.

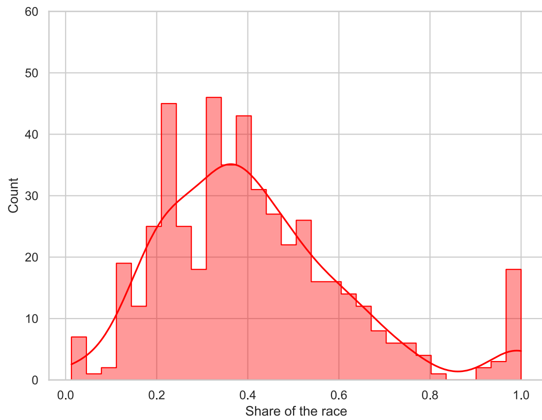

(a) Leader (1-stop races)

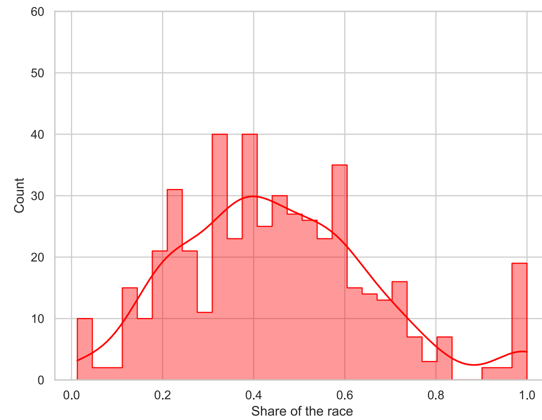

(b) Pursuer (1-stop races)

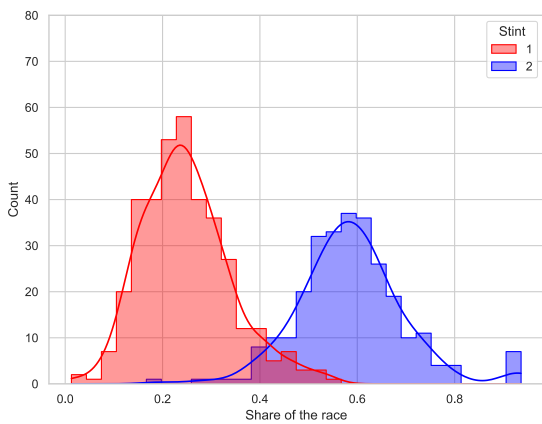

(c) Leader (2-stop races)

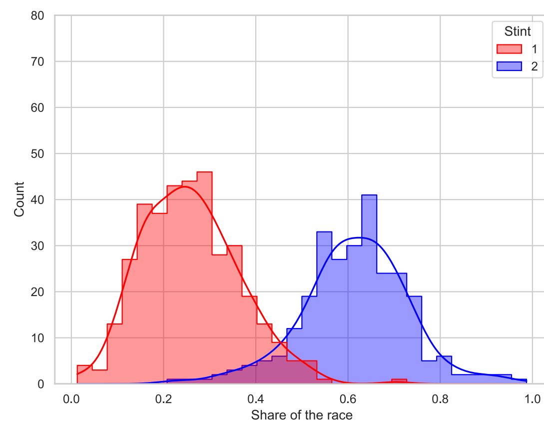

(d) Pursuer (2-stop races)

Notes: The figure reports the distributions of pit-stop timing, expressed as a share of race distance. Panels (a) and (b) show pit stops for leaders and pursuers in one-stop races. Panels (c) and (d) display the timing of the first and second stints in two-stop races, respectively. The sample includes all pit stops that satisfy the data-cleaning procedure described in Section 4.3 and conditions on drivers finishing the race.

Figure A1: Distributions of pit-stop timing by role and race type

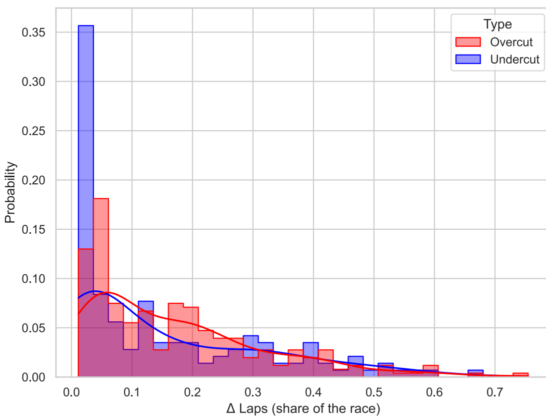

(a) Δ Laps (between stops)

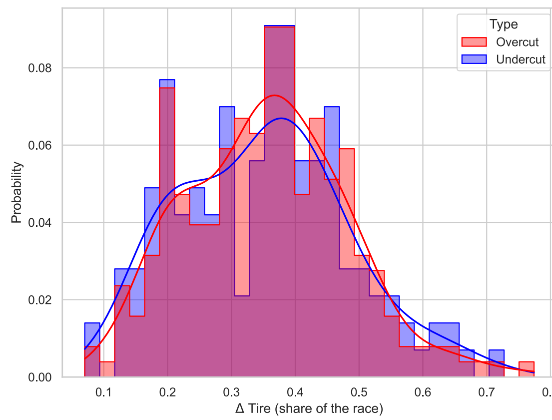

(b) Δ Tire (between stops)

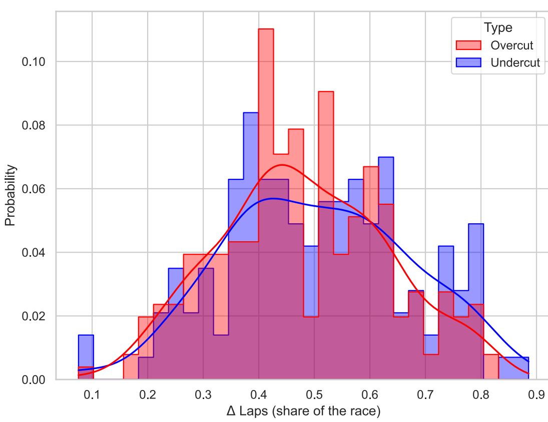

(c) Δ Laps (after stops)

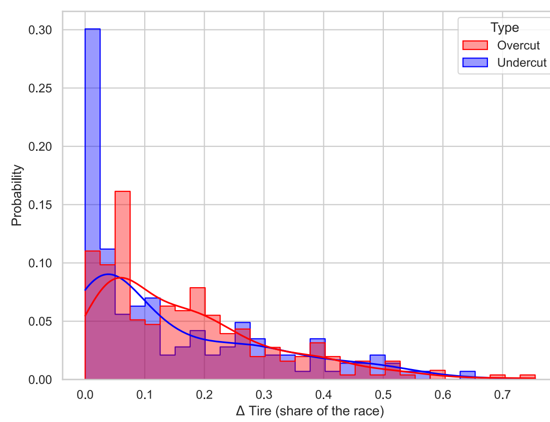

(d) Δ Tire (after stops)

Notes: The figure shows the distributions of the main explanatory variables used in the estimation, separately for undercut and overcut strategies. Panels (a) and (b) report the distributions of ΔLaps and ΔTire during the period between drivers' pit stops. Panels (c) and (d) show the corresponding distributions for the period after both drivers have completed their pit stops and until the end of their interaction. All variables are measured as shares of total race distance. The sample consists of single-stop strategies by both drivers and follows the data-cleaning procedures described in Section 4.3.

Figure A2: Distributions of main explanatory variables

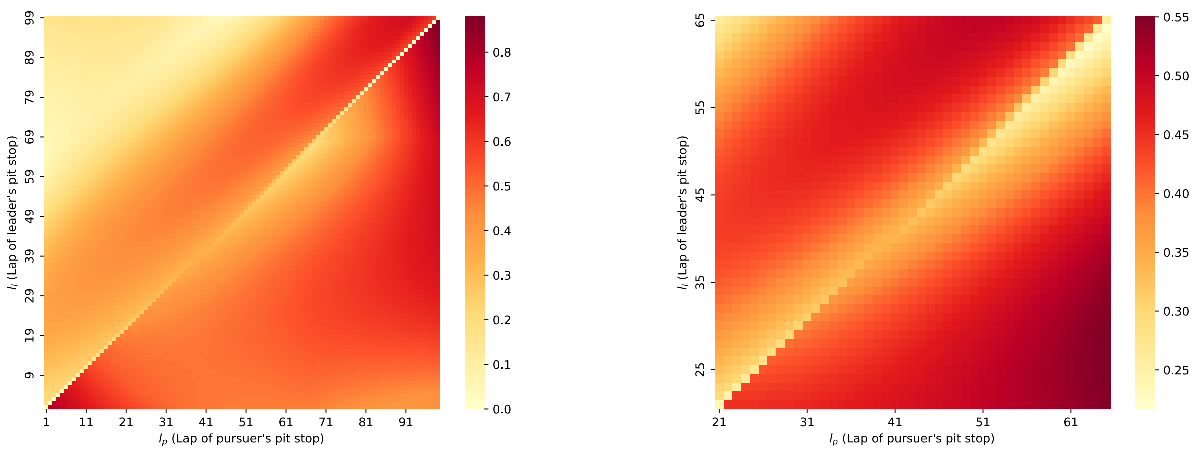

(a) All feasible pit-stop choices

(b) Empirically relevant pit window

Notes: The figure shows predicted probabilities that the pursuer overtakes the leader, $P(l_L, l_P)$, for all combinations of pit-stop laps chosen by the leader (l_L) and the pursuer (l_P). Panel (a) reports predictions for all feasible lap combinations in a 100-lap race. Panel (b) restricts attention to the empirically relevant pit window, defined as the interval containing 80% of observed pit stops. The diagonal corresponds to same-lap pit stops. Predictions are computed using the random-effects estimates based on single-stop strategies only. Warmer colors indicate a higher probability that the pursuer finishes ahead of the leader.

Figure A3: Predicted overtake probabilities across pit-stop choices

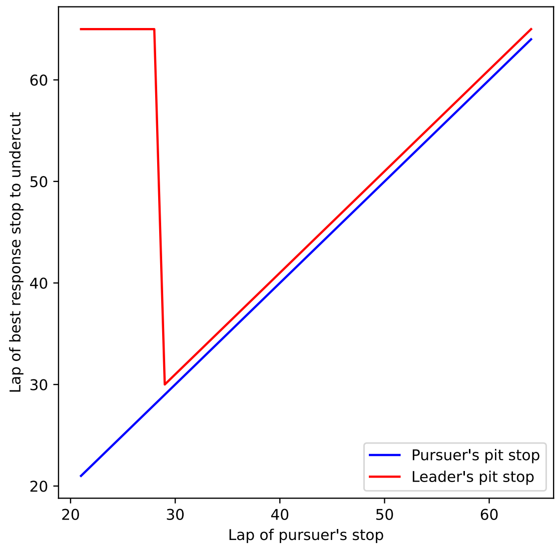

(a) Best responses of leaders

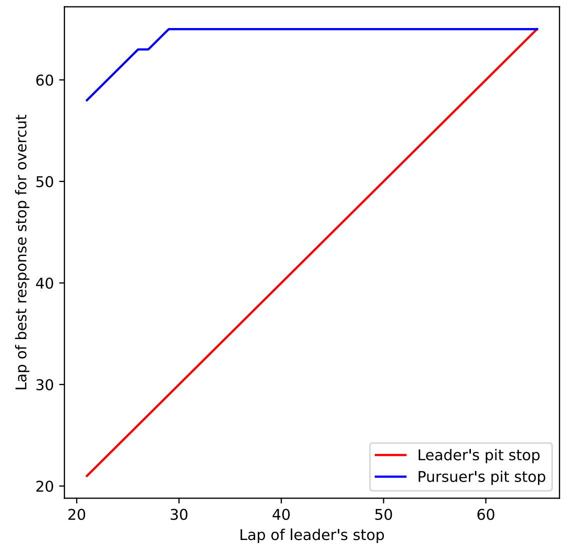

(b) Best responses of pursuers

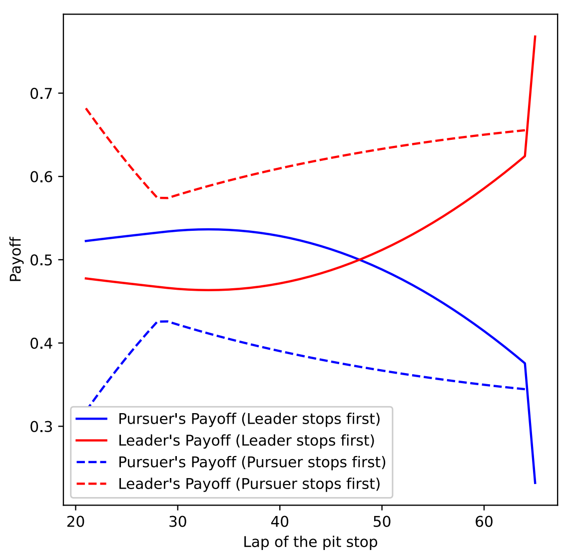

(c) Payoffs for pursuer and leader

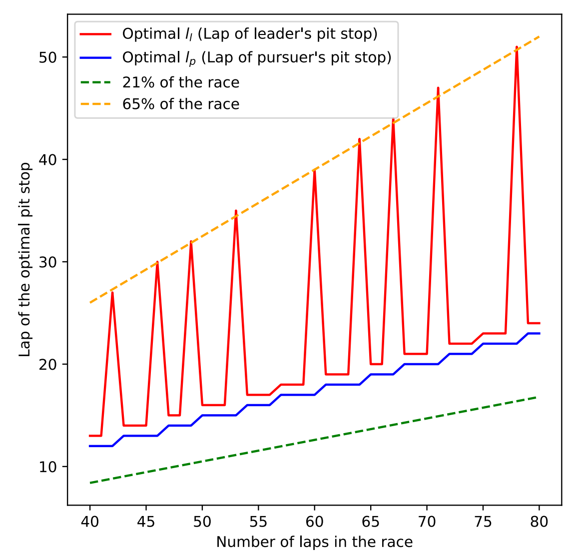

(d) SPNE

Notes: This figure reproduces the theoretical predictions presented in the main text using estimates based on all pit stops rather than restricting the sample to one-stop strategies. Overtake probabilities are estimated using random-effects specifications. Panel (a) shows the leader's best-response pit-stop timing as a function of the pursuer's pit-stop decision. Panel (b) shows the pursuer's best-response pit-stop timing as a function of the leader's pit-stop decision. Panel (c) plots the payoffs of the pursuer and the leader as a function of pit-stop timing, distinguishing whether the leader or the pursuer stops first. Panels (a)–(c) are computed for a 100-lap race. Panel (d) reports the SPNE pit-stop timing as a function of total race length. Throughout the figure, pursuer's decisions and payoffs are shown in blue, while leader's decisions and payoffs are shown in red. All predictions are restricted to the empirically relevant pit window.

Figure A4: Model predictions using all pit stops

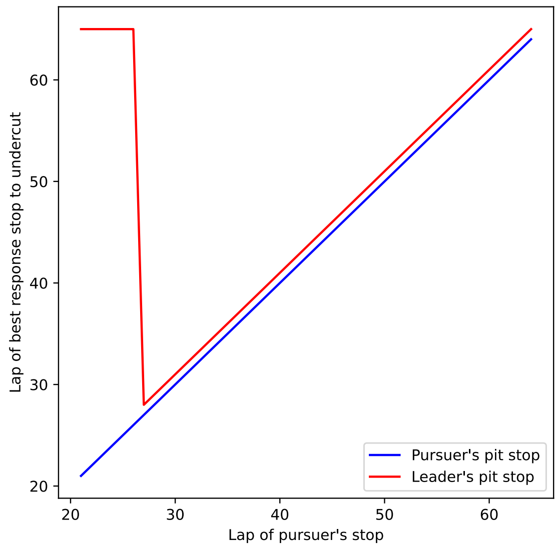

(a) Best responses of leaders

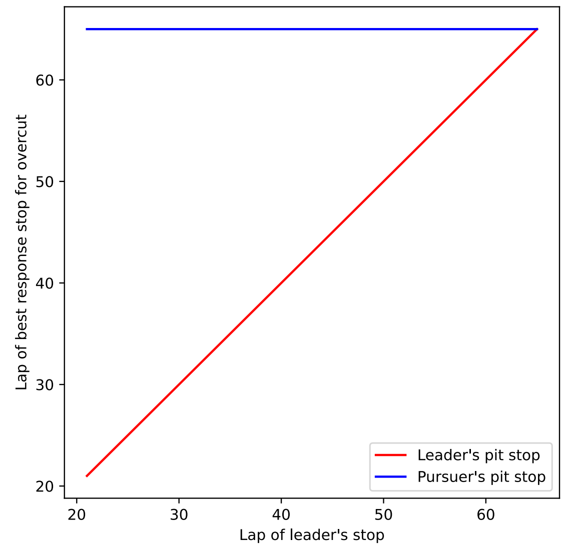

(b) Best responses of pursuers

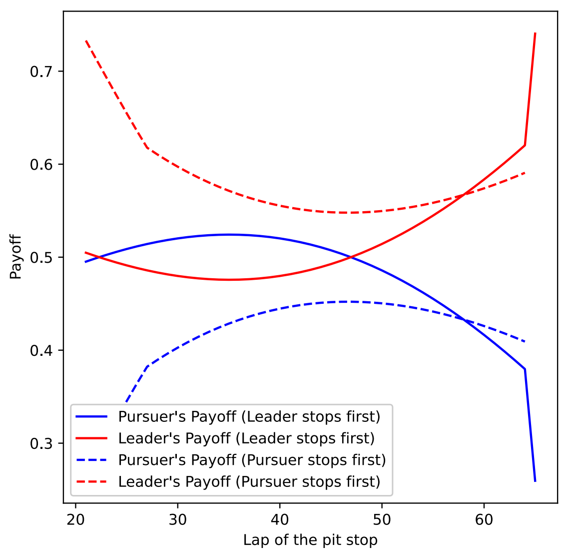

(c) Payoffs for pursuer and leader

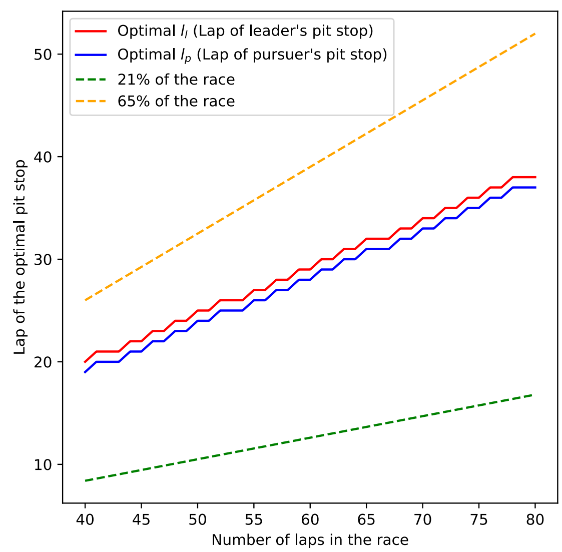

(d) SPNE

Notes: This figure reproduces the theoretical predictions presented in the main text using estimates based on last pit stops rather than restricting the sample to one-stop strategies. Overtake probabilities are estimated using random-effects specifications. Panel (a) shows the leader's best-response pit-stop timing as a function of the pursuer's pit-stop decision. Panel (b) shows the pursuer's best-response pit-stop timing as a function of the leader's pit-stop decision. Panel (c) plots the payoffs of the pursuer and the leader as a function of pit-stop timing, distinguishing whether the leader or the pursuer stops first. Panels (a)–(c) are computed for a 100-lap race. Panel (d) reports the SPNE pit-stop timing as a function of total race length. Throughout the figure, pursuer's decisions and payoffs are shown in blue, while leader's decisions and payoffs are shown in red. All predictions are restricted to the empirically relevant pit window.

Figure A5: Model predictions using last pit stops

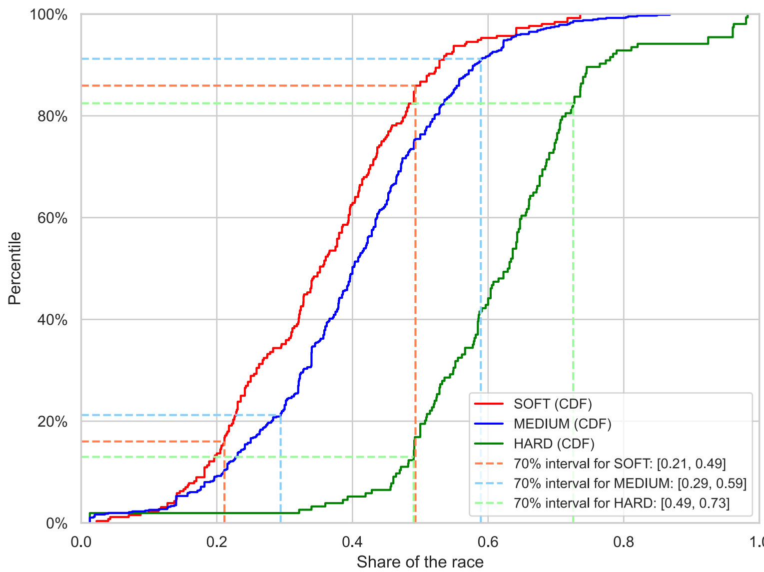

Notes: The figure plots the empirical cumulative distribution function (CDF) of pit-stop timing, measured as the share of race distance completed at the time of the stop, separately by starting tire compound (soft, medium, and hard). Vertical dashed lines indicate the shortest intervals that contain 70% of pit stops for each compound. These compound-specific intervals are used to define alternative pit windows in the theoretical analysis. The sample includes all single-stop strategies that satisfy the data-cleaning procedure described in Section 4.3.

Figure A6: CDF of pit stops and implied pit windows by starting tire compounds

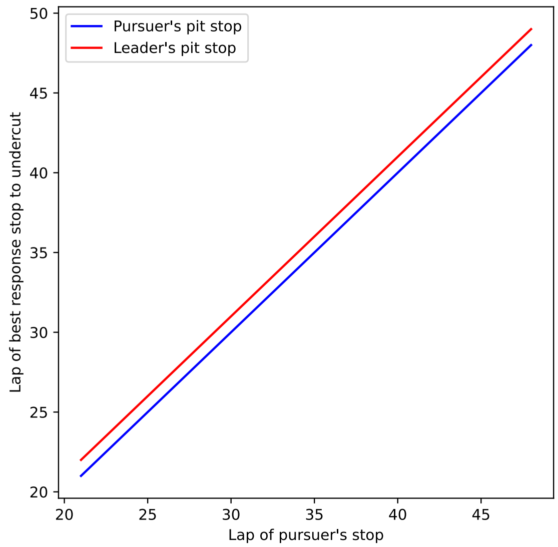

(a) Best responses of leaders

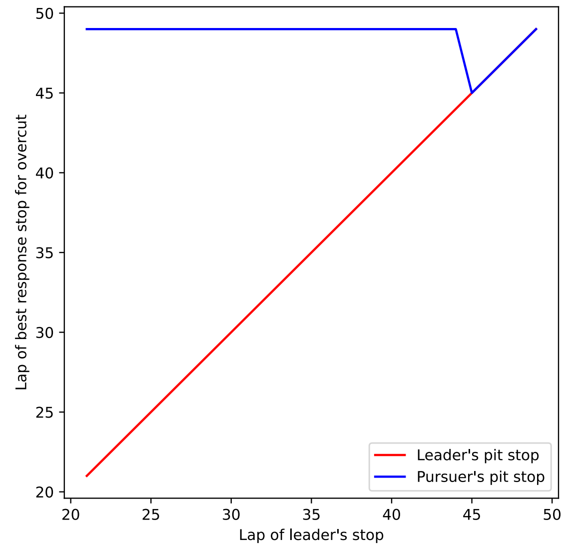

(b) Best responses of pursuers

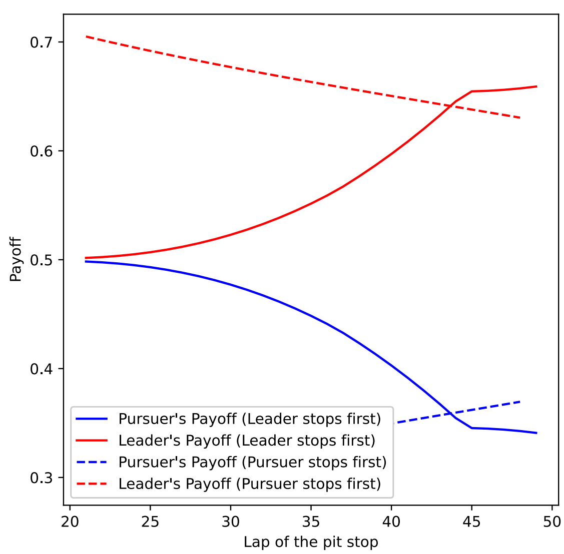

(c) Payoffs for pursuer and leader

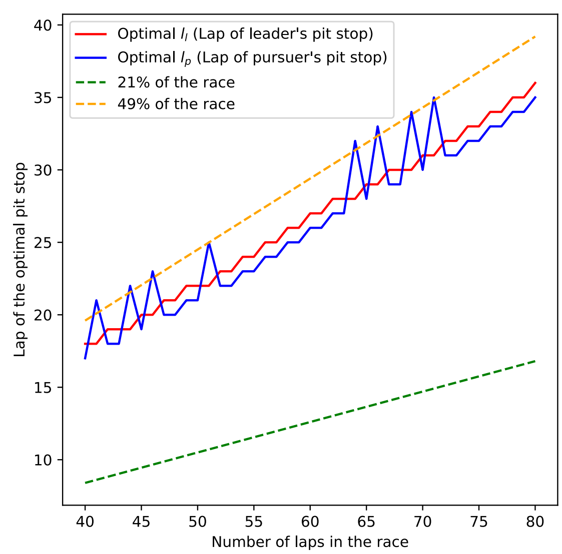

(d) SPNE

Notes: This figure reproduces the theoretical predictions presented in the main text using estimates based on one-stop strategies. Overtake probabilities are estimated using random-effects specifications. Panel (a) shows the leader's best-response pit-stop timing as a function of the pursuer's pit-stop decision. Panel (b) shows the pursuer's best-response pit-stop timing as a function of the leader's pit-stop decision. Panel (c) plots the payoffs of the pursuer and the leader as a function of pit-stop timing, distinguishing whether the leader or the pursuer stops first. Panels (a)–(c) are computed for a 100-lap race. Panel (d) reports the SPNE pit-stop timing as a function of total race length. Available pit-stop laps are restricted to the empirically relevant pit window defined as the shortest interval containing 70% of pit stops among drivers starting the race on the soft compound. Throughout the figure, pursuer's decisions and payoffs are shown in blue, while leader's decisions and payoffs are shown in red.

Figure A7: Model predictions with soft compound pit window

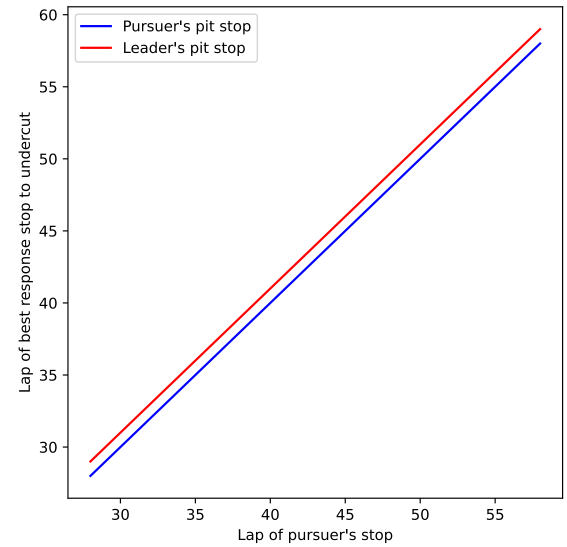

(a) Best responses of leaders

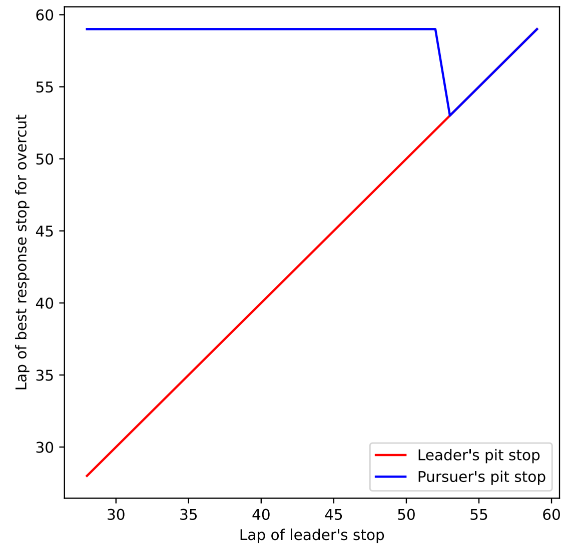

(b) Best responses of pursuers

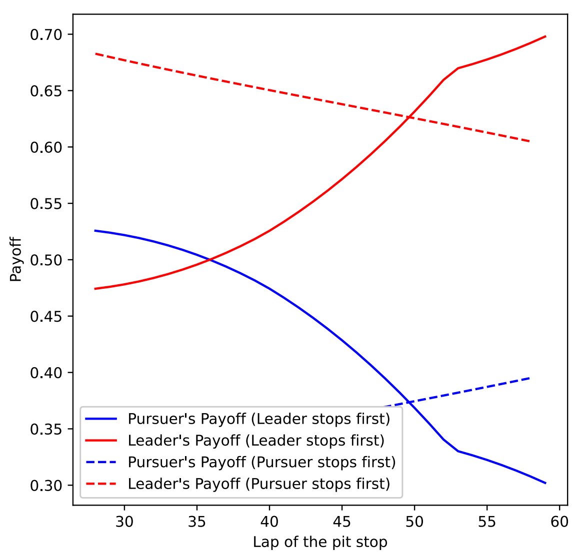

(c) Payoffs for pursuer and leader

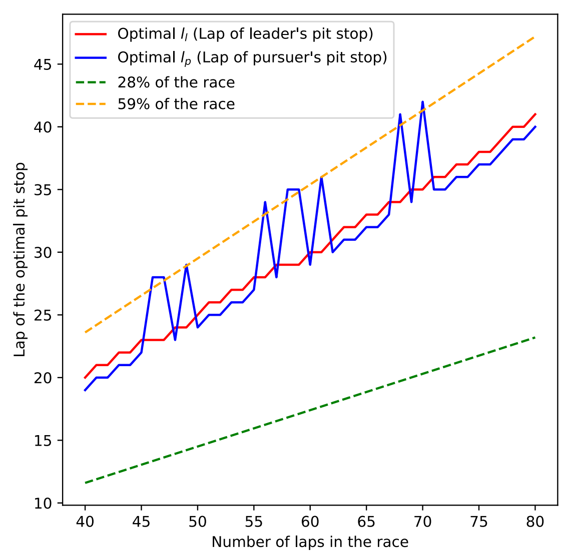

(d) SPNE

Notes: This figure reproduces the theoretical predictions presented in the main text using estimates based on one-stop strategies. Overtake probabilities are estimated using random-effects specifications. Panel (a) shows the leader's best-response pit-stop timing as a function of the pursuer's pit-stop decision. Panel (b) shows the pursuer's best-response pit-stop timing as a function of the leader's pit-stop decision. Panel (c) plots the payoffs of the pursuer and the leader as a function of pit-stop timing, distinguishing whether the leader or the pursuer stops first. Panels (a)–(c) are computed for a 100-lap race. Panel (d) reports the SPNE pit-stop timing as a function of total race length. Available pit-stop laps are restricted to the empirically relevant pit window defined as the shortest interval containing 70% of pit stops among drivers starting the race on the medium compound. Throughout the figure, pursuer's decisions and payoffs are shown in blue, while leader's decisions and payoffs are shown in red.

Figure A8: Model predictions with medium compound pit window

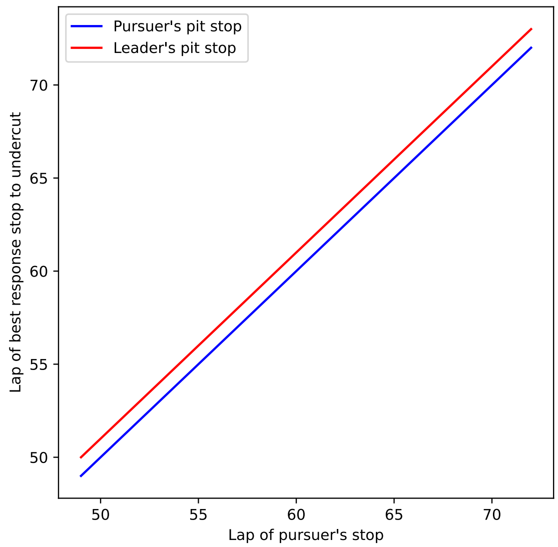

(a) Best responses of leaders

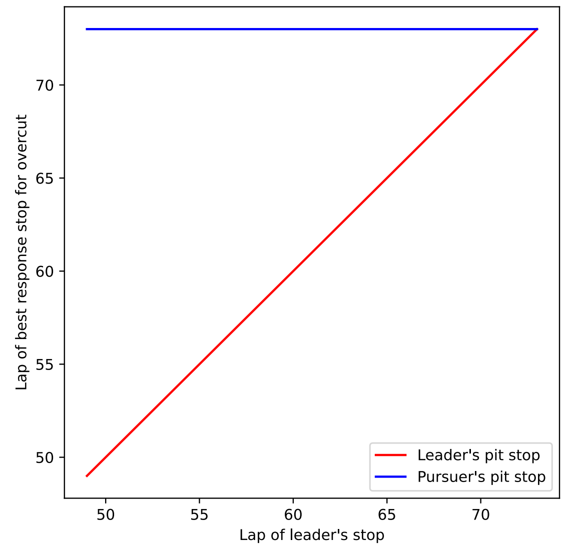

(b) Best responses of pursuers

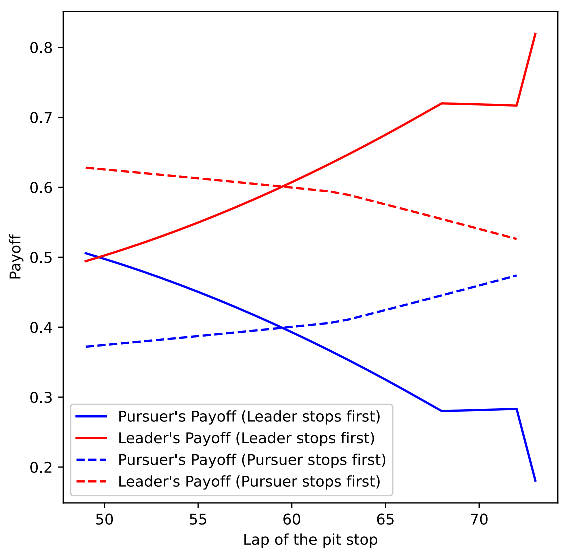

(c) Payoffs for pursuer and leader

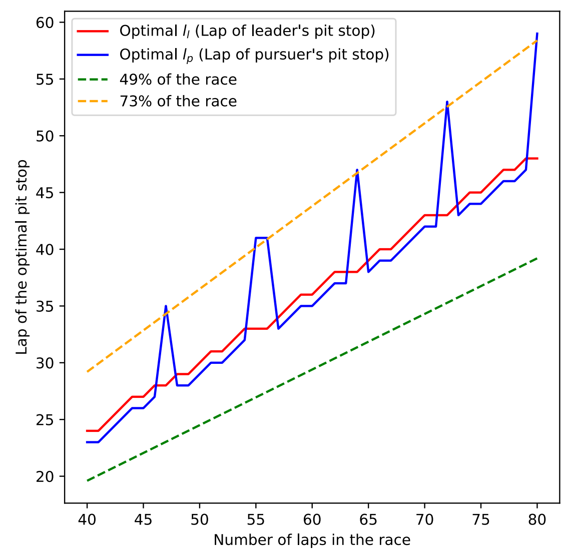

(d) SPNE

Notes: This figure reproduces the theoretical predictions presented in the main text using estimates based on one-stop strategies. Overtake probabilities are estimated using random-effects specifications. Panel (a) shows the leader's best-response pit-stop timing as a function of the pursuer's pit-stop decision. Panel (b) shows the pursuer's best-response pit-stop timing as a function of the leader's pit-stop decision. Panel (c) plots the payoffs of the pursuer and the leader as a function of pit-stop timing, distinguishing whether the leader or the pursuer stops first. Panels (a)–(c) are computed for a 100-lap race. Panel (d) reports the SPNE pit-stop timing as a function of total race length. Available pit-stop laps are restricted to the empirically relevant pit window defined as the shortest interval containing 70% of pit stops among drivers starting the race on the hard compound. Throughout the figure, pursuer's decisions and payoffs are shown in blue, while leader's decisions and payoffs are shown in red.

Figure A9: Model predictions with hard compound pit window

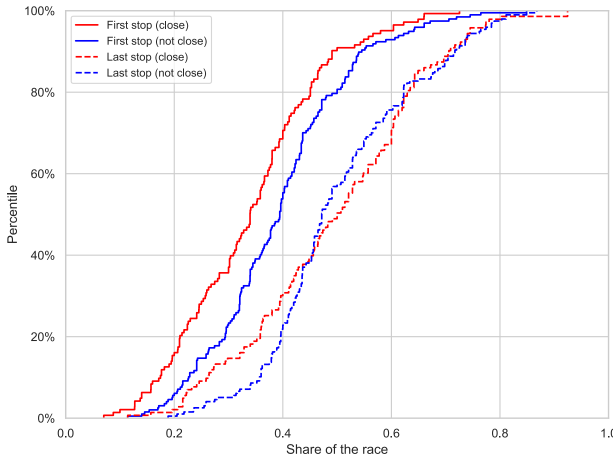

(a) Undercuts

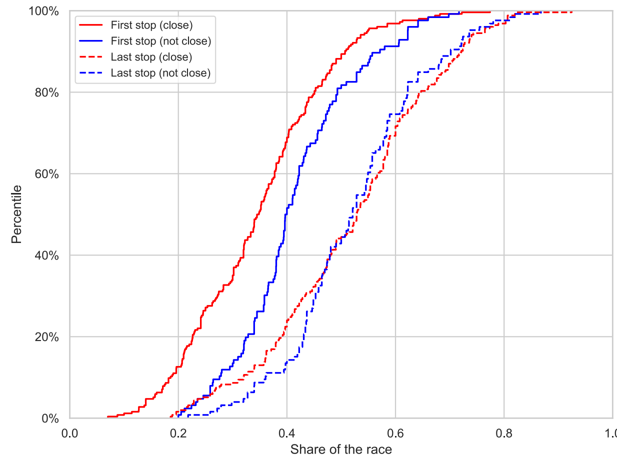

(b) Overcuts

Notes: The figure plots empirical cumulative distribution functions (CDFs) of pit-stop timing, measured as the share of race distance completed at the time of the stop. The unit of observation is a pair of drivers who each make exactly one pit stop in a race. Within each pair, the first stop refers to the driver who pits earlier, while the last stop refers to the driver who pits later. Panel (a) restricts attention to undercut interactions, defined as cases in which the pursuer is driver who pits first. Panel (b) shows overcut interactions, where the driver who pits first is the leader. Solid lines correspond to first pit stops, and dashed lines correspond to last pit stops. Red lines restrict the sample to close pit-stop interactions, defined as pairs of drivers who are within three seconds of each other on track at the time of the first stop, while blue lines include pairs that are not classified as close. The sample follows the data-cleaning procedure described in Section 4.3.

Figure A10: CDF of pit-stop timing in undercut and overcut interactions

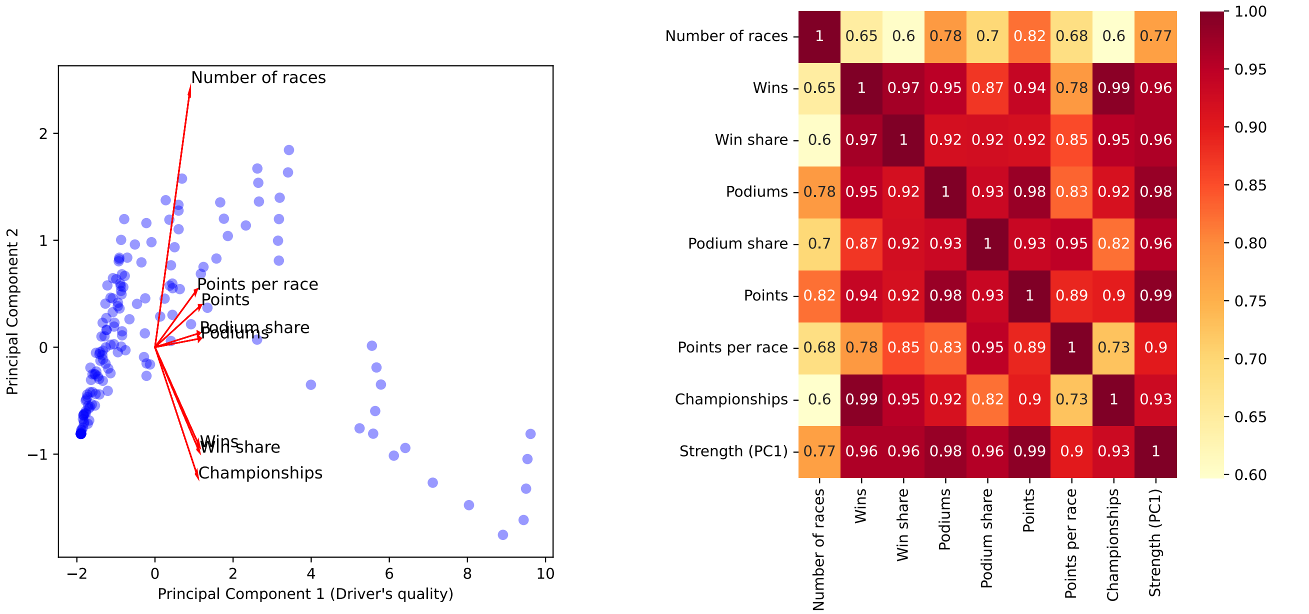

(a) Distribution of the first two components

(b) Correlation heatmap

Notes: This figure summarizes the construction and interpretation of the driver performance index based on principal component analysis (PCA). Panel (a) plots drivers in the space of the first two principal components. Arrows indicate variable loadings for the underlying performance measures, including wins, win share, podiums, podium share, points, points per race, championships, and the total number of races. The first principal component loads positively on all major performance outcomes and is interpreted as a summary measure of overall driver strength. Panel (b) reports pairwise correlations between the original performance measures and the first principal component. The heatmap illustrates that the first principal component is highly correlated with all underlying indicators of driver performance. The resulting first principal component is used as a composite measure of driver quality in the empirical analysis.

Figure A11: Principal Component Analysis of Driver Performance