Country Roads, Take Me Home: Evaluating Home Effect in the Tour de France

Eren Celik* Xiaoyi Guo† Angelo Mimmo ‡

1 Introduction

Home-field advantage is a long-standing topic of interest in team-based sports (for a comprehensive review, see Courneya and Carron 1992; Jones 2013). In team sports (e.g. football), under a balanced schedule of home and away matches in a competition, if venue had no effect on performance, one would expect a random distribution of wins between home and away games – approximately a 50/50 split. Therefore, a home-field advantage is typically identified when home teams consistently win more than half of their games under such balanced conditions. Research has shown that the size of home advantage is heterogeneous across sports, depending on the sport type and whether there is team cooperation (R. Pollard and G. Pollard 2005). While most of the literature is focusing on the assessment of home-advantageous effect in team sports, such as basketball, American football, and soccer, work on individual sports is rare. Nevill et al. (1997) carried out an analysis on the major professional golf and tennis tournaments, and found mixed results in terms of home advantage. By regressing tournament ranking on world ranking (both in logs) separately for home and away players, they found that home advantage does not seem a relevant feature in the 1993 US, Australian and French Open tennis tournaments, as neither the slope nor the intercept significantly differed between the two groups. For Wimbledon, a difference in the intercept was found, but this was attributed to the low mean world ranking of the British players in that period. For golf, no significant differences were found in either the slope or intercept of the two regression lines for the 1993 British Open, U.S. Masters, and U.S. PGA. A difference was found in the slope for the U.S. Open, but this was again attributed to large number of U.S. players invited and their relatively inferior rankings on average. Conversely, Koning (2011) used as measure of performance in tennis the win a single match, instead of tournament ranking. The author focused on all matches in Grand Slam, Masters, and International tournaments between 2000 and 2008 for men and between 2007 and 2008 for women. By applying a logit model, the author found that significant home advantage exists for men, but not for women. Ramchandani and Wilson (2020) focused on track and field athletics and analyzed the IAAF World Championships from 1991 to 2017 and the IAAF World Indoor Championships from 1987 to 2018. They considered 3 different KPIs related to both the share of gold medals and the medals points obtained by each country in each edition. By means of non-parametric statistical tests, they checked whether there was a difference in nations' performance under host and non-host conditions. For the indoor championships, they found a statistically significant home advantage effect for all the KPIs examined, but not for the track championships. To best of our knowledge, no research so far has specifically focused on home effect in cycling.

This paper tries to fill this gap by answering to the following question: *in the Tour de France, does riders' performance improve, decline, or remain unchanged when racing in their home country?* We use modified causal forests (MCF, Lechner 2019) to comprehensively understand the home effect in teamed-up individual sport competition —- professional cycling. Overall, we found a negative effect of being a home racer, i.e., the stage of the race takes place in riders' home country. We tend to credit this negative effect to the social and public pressure and other cultural factors, since we had controlled the non-cultural and -psychological factors well. We also identified heterogeneity of this effect in different riders with various characteristics. This may be particularly useful for professional coaches and other researchers trying to further work in this field.

*Master Student in Economics. Email: eren.celik@students.unisg.ch. Name of authors are ordered by their surnames.
†PhD Student in Economics and Research Associate at Swiss Intitute of Banking and Finance. Email: xiaoyi.guo@unisg.ch.
‡PhD Student in Econometrics and Research Associate at Swiss Institute of Banking and Finance. Email: angelo.mimmo@unisg.ch.

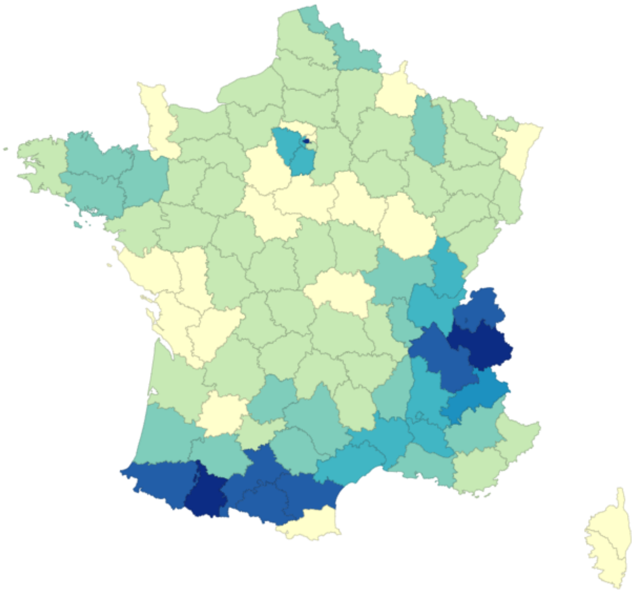

Figure 1: **French regions holding Tour de France from 2000 - 2024** (Breteau 2024)

Section 2 introduces the background of the race itself and related characteristics in this game. Section 3 introduces all data we used and how we deal with different variables carefully. In section 4, we specifically argue our empirical strategies and estimation methods, we further present the expected results that we are looking for. Subsequently, we discuss the empirical results we acquired in section 5, and we finally conclude in section 6.

2 Background

2.1 Overview of Tour de France

Tour de France (or "*Le Tour*") is one of the most important annual men multiple-stage cycling race. First organized in 1903, it is the oldest and most prestigious bicycle tour in the world and considered as one of the three major European professional cycling stage races, the *Grand Tours*. It covers around 3,500 km of distance and is completed in around 21 - 24 days. The route of the tour changes from year to year and, in general, it is held in June or July. Regardless of the starting point, riders have to pass the mountainous southern France – with different stages on the Pyrenees and the Alps – and finally reach a French city – in general, the final stage takes place in Paris. The race route is publicly announced about 8-9 months before. Most of the stages of the tour take place in France, but sometimes they start from other places, such as Brussels (2019), Copenhagen (2022), and Florence (2024). Figure 1 illustrates the number of times that the different French counties have been passed through between 2000 and 2025. The darkest blue refers to counties crossed 25 or more times in the period considered, while the lightest yellow refers to counties crossed lower than 5 times.

2.2 Teams, Riders, and Stages

The tour is organized into multiple stages – typically 21. Each stage is carried out for one whole day. Generally, there are two rest days throughout the tour. Each stage has its individual rankings, while the general classification is given by the cumulative time across all stages. This means that the rider with the lowest total time across all the stage is the overall leader and wears the "yellow jersey". There are also other kinds of jerseys, such as "polka dot" (for the best climber), "green" (for the best sprinter) and "white" (for the best young rider).

Although the riders are competing against each other and what matters is the individual result, the entire tour is organized in teams. Every year, there are around 20 - 22 teams, each composed by 8 riders, taking part in the race. The teams are selected following the regulations of International Cyclists Union (Union Cycliste Internationale, UCI), with the 18 UCI World Tour teams (i.e. the best ones) plus the "wild card" teams invited directly by the organizers. In general, the competitive level is very high. However, among all riders taking part into the tour, not everyone finishes it[1]. This could be due not only to random circumstances (such as accidents or illnesses), but also due to specific strategies (Torgler 2007; Phillips and Hopkins 2020), for example to save energies and resources for the coming events of the season. Team leaders (also known as captains) are helped and protected from other teammates to earn higher final rankings; therefore, it may appear that some of the riders may strategically preserve their efforts or “sacrifice” themselves to help the team’s captain getting higher rankings and make the team higher ranked as well.

3 Data and Variables Selection

Our primary data source is a random sample of men’s stages of Tour de France from 2020 to 2024. This dataset includes comprehensive information on individual riders, their affiliated teams, and stage-level race details. The data were initially obtained from the professional cycling statistics platform Pro Cycling Stats (PCS, `procyclingstats.com`). However, due to limitations in data quality and the lack of key control variables, we conducted supplementary data collection from the same source. Additionally, we employed the GeoPy geocoding service to enhance the accuracy of geographical information.

During the data cleaning process, we identified several issues: (i) the total number of stages (i.e. `number_of_stages_tour`) for 2022 was 24, which was wrong, so we imputed it to 21; (ii) weight and height were missing for some riders, thus we retrieved them directly from PCS and we computed the body mass index (`bmi`) for all the riders; (iii) we found two duplicated columns, `distance_km_tour` and `totaltdfdistance_tour`, thus we dropped the latter; (iv) some riders lacked publicly available body measurement data (namely, Thomas Gachignard, Sandy Dujardin and Mattéo Vercher, all riding for Team TotalEnergies in the 2024 Tour); (v) the results of the rider Nairo Quintana for the 2022 Tour were not available because, after testing positive for the UCI-prohibited substance Tramadol, he was retroactively disqualified; and (vi) we found observations for Maximilian Schachmann for the 2021 Tour, in which he never took place (moreover, assigned to a team for which he never rode for). For the last three cases reported, these observations were excluded from our final dataset. This way, from the orignal 14.744 observations, we were able to keep 14.639 observations, losing only 105 of them.

Our main dependent variable is stage-level performance defined as a binary indicator for whether a rider finishes within the top 15 in a given stage. This threshold was chosen for several reasons. In fact, approximately 170 - 180 riders take part in each stage; however, maintaining a low cumulative time is critical for winning the overall classification or obtaining a classification jersey. The UCI awards tiered points to the top 15 finishers in each stage, and these points play a pivotal role in assessing a riders professional standing and influencing future contract negotiations with teams. Indeed, contracts in cycling often have a very short term and riders need to deliver strong results in order to get a renewal (Rebeggiani 2016).

The major treatment is a dummy variable of whether the home advantage holds. It is 1 if the country of a given stage (`country_stage`) is equal to the nationality of this rider (`country_rider`). We do this because we have well-controlled other factors related to the country and nationality variations, which will be elaborated later. The home advantage here can be interpreted as the crowd support of “home” residents, which might provide an extra boost, improving a riders performance. This effect is controversial in football (Fischer and Haucap 2021; Correia-Oliveira and Anrade-Souza 2022) and confirmative in basketball (Courneya and Carron 1992). On the other side, competing in their own country might put higher pressure on these riders, negatively affecting their performance. Table 2 depicts the country distribution of riders, teams, and all matches. Therefore, the home effect that we will evaluate is strongly concentrated on French riders, which is a potential limitation that will be discussed in section 6.

Our control variables are classified in 5 major clusters: (i) Tour features, (ii) rider’s features, (iii) team features, (iv) stage features, and (v) combination of rider/team and stage features.

Tour features. The total length of the Tour is generally around 3,500 km, but there might be slight variations from year to year. For example, in the 5 years considered in this study, the total distance ranged from 3328 to 3492 km. Moreover, year by year, the general level and quality of the teams involved and the weather could differ. Therefore, we controlled for the year (directly in the MCF and with year dummies in the Logit model

[1]For example, among 176 riders in Tour 2024, only 141 finish the race.

[2]) and for the total distance of that year edition of the Tour.

Rider's features. We include several control variables related to individual rider characteristics. All of these have been scraped from the PCS website. First, we use the riders age (defined as the age in years at the starting date of the tour) as control for experience, following Torgler (2007), and the body mass index (BMI), following Torgler (2007) and Phillips and Hopkins (2020). Indeed, prior research in exercise physiology (Lucía et al. 2001; Van Reeth and D. J. Larson 2016) suggests that body measurement data, such as BMI, can serve as a reliable proxy for physical characteristics. BMI is calculated from the height and weight of the rider. Second, we incorporate riders' PCS points from the previous calendar year as a measure of their quality and skills, similarly to Torgler (2007). PCS points are cumulative scores assigned to riders based on their competitive achievements. We opt for PCS points instead of UCI points because the latter are calculated on a rolling 52-week basis and are less readily available for historical reference. In contrast, PCS scores are accessible from 1976 up to today and provide a consistent measure over time. In fact, since its launch in 2013, PCS has always kept the same point scales. For riders with missing PCS scores - often due to injury-related absences in one year – we impute values using the mean of their available PCS scores over the previous three years[3]. We further account for rider specialization by including PCS-derived specialty scores that reflect historical performance across different types of *parcours* (also known as routes or profile score) and stage types (regular races vs. time trials). The rider's specialty is defined as the profile score specialty (namely, "sprint", "climber" and "hills") for which he has the largest number of PCS points. Similarly, we define the rider as a time trial (TT) specialist if his TT points account for at least 50% of the maximum points across all the other categories available on PCS (i.e. "one day races", "general classification", "sprint", "climber" and "hills"). We also control for whether the rider was the captain[4] during that year of the Tour. Generally, the captain is the one, within a team, with the highest chances of winning the Tour one or more stages of it. In fact, cycling is a sport with both individual and team dimensions and both cooperation and competition can occur both within and across teams. The so-called *coopetition* (Brandenburger and Nalebuff 2011) has been specifically observed and verified in the cycling sports (Matthes and Piazolo 2024). In this setting, the captain of the team benefits from the work of their teammates (Menaspa et al. 2013). Conversely, the performance of *domestiques* (i.e. the cyclists riding in support of their captain) suffers from this, as they are expected to sacrifice their own chances of success in order to improve the one of their captain (Torgler 2007).

Team features. The quality and strength of a team differs depending on its members. A cyclists overall performance is influenced by the skills and attributes of their teammates (Phillips and Hopkins 2020). Having more experienced or more skilled riders in the same team positively affects the performance of a rider (Torgler 2007). More experienced and more skilled riders can help each other more effectively. Moreover, this might also boost within-team competition, with an effect on individual results. Therefore, we include the sum of the PCS points obtained from each rider in a team in the previous calendar year as a control for the overall quality of a team.

Stage features. Since each stage is different, we incorporate as controls its distance, an indicator whether it was a regular road race or a time-trial (since these are two very different types of specialties), its profile score (i.e. "Flat", "Hills, flat finish", "Hills, uphill finish", "Mountains, flat finish" and "Mountains, uphill finish") and the percentage of tour completed including that stage (stage number divided by the total number of stages).

Combination of rider/team and stage features. We also control for series of indicators highlighting the match between rider' features and stage features. First, we control for whether the rider is a TT specialist and the stage considered is a time trial. Second, we control for whether the profile score of the current, the previous and the following stages matches with the specialty of the rider as defined above. In particular, the rider's specialty "Sprinter" was matched with profile score "Flat", "Hills" with "Hills, flat finish", "Hills, uphill finish" and "Climber" with "Mountains, flat finish" and "Mountains, uphill finish". We expect a positive effect for TT and specialty match indicators, as a current stage that matches rider's characteristics increases his possibility of a good result. Conversely, we expect a negative effect for specialty match indicators for the previous (as the rider might have been consumed more energies in it since he had higher chances of a good result) and the next stages (as the rider might preserve his energies for it). We also introduce a novel control for route familiarity. Riders may be familiar with a stage for several reasons: they may have been born and trained

[2]More details about the two methods are provided in section 4.1.

[3]Only two cases in the whole sample: Guillaume Van Keirsbulck in 2021 and Egan Bernal in 2022.

[4]The captain is the first rider in the official start list of his team and his assigned number ends with 1 (i.e. 1, 11, 21, etc.).

in the region (with benefits from both familiarity with the area and the terrain and enhanced psychological or logistical support) or they may have prepared in advance based on the publicly released route, which is typically announced 8-9 months prior to the race and often revisits similar regions (see Figure 1). Assuming that the latter is homogeneous across riders and teams, to quantify the former, we geocoded the coordinates of each riders birthplace, as well as the starting and ending locations of each stage. We then computed the driving distance from the riders birthplace to both the start and end points of the stage. If either distance is less than or equal to 150 km[5] we assign a value of 1 to the area knowledge indicator, 0 otherwise. Finally, we also control for the whether the country of a given stage (`country_stage`) is equal to the country of the rider's team (`country_team`). The inclusion of this control has been carefully assessed. The main sponsors and the so-called "paying agent" determine who holds the UCI license, as well as the name and the nationality of the team (Rebeggiani 2016), since nowadays all cycling teams are financed almost exclusively through sponsorships (Van Reeth 2022b). In any case, our idea is that the country of the team affects both the rider's nationality (i.e. the treatment) and the rider's performance in a stage (i.e. the outcome). On the performance side, riders in a team from the same country as the stage considered might feel a stronger support from the local fans. Similarly, the focus and the pressure put on their riders by these teams might be higher. In both cases, we expect this "home team effect" on riders' performance to be positive. This is also consistent with the spirit of the 2005 UCI ProTour reform "to avoid the historical phenomenon of cyclists concentrating on competitions in their home countries (or in the sponsor's home country)" (Rebeggiani 2016, p.p.35). On the treatment side, instead, the assumption is that teams tend choose riders from their domestic country more than from other countries. Van Reeth 2022a reports that, in 2020, around one third of the riders in top-level cycling teams were in a team of their native country, although this share was much larger in the past 3 decades. Our idea behind this relation is that teams choose their riders and not vice versa. This is consistent with the fact that cyclists contracts are typically annual or biannual and the renewal really depends on the sport results (Philips and Hopkins 2020, Larson and Maxcy 2013; 2016) . This is also reflected by very short careers at the UCI WorldTour level for a very large share of riders: "more than one-third exit in the first two years, and less than half stay at the top level for more than 5 years" (D. J. Larson and J. G. Maxcy 2016, p.p.134). Therefore, we can conclude that at least the vast majority of the riders has very low bargaining power. We therefore assume that the the hiring by the team of a domestic rider is either direct (e.g. for reasons of culture, language and/or logistics) or determined by some unobserved factor that also determines the nationality of the team at the same time (namely, the sponsor which might want riders from the same county for visibility and commercial reasons). In both cases, this home country match at the team level represents a good control. Table 2 also shows the cross-count of teams, riders, and stages. As already mentioned, the team and riders matching results show that mostly French teams and riders are evaluated.

The following Table 1 depicts the summary statistics for numerical variables used.

Table 1: **Descriptive Statistics of Selected Variables**

Variable	Mean	Median	Variance	Min	Max	NAs
rider_points_previous_year	463.8	314.0	229178.8	6.0	3413.0	0
bmi	21.1	21.0	2.3	17.6	25.3	0
team_points_previous_year	3096.1	2752.0	2686469.3	352.0	9087.0	0
profile_icon_stage	2.9	3.0	2.4	1.0	5.0	0
perc_tour_completed	51.1	52.4	868.1	4.8	100.0	0
age	28.9	29.0	16.3	20.0	41.0	0

[5] Lucía et al. 2001 reports that professional riders train for around 35,000 km year. We divided this for 50 weeks and, then, for 5 days.

Table 2: **Match Counts by Country**

Country of Stage	Rider = Stage	Team = Stage	All Three Match
France	2583	3238	2124
Spain	30	12	6
Belgium	27	50	15
Italy	24	0	0
Denmark	12	0	0
Switzerland	4	0	0

4 Empirical Strategy

4.1 Estimation Method

So far, the literature discussing the home advantage and team performance lacks rigorous research and experimental design, econometric methods, and solid assumptions and controls. This paper will investigate the effects of home effects with three layers of control to our central question of whether it enhances the performance on average. Existing papers on this topic either use traditional OLS methods without serious identifications to exclude confounders and other assumptions(Correia-Oliveira and Anrade-Souza 2022; Rebeggiani 2016; Courneya and Carron 1992). This paper uses a novel developed causal machine learning estimator of modified causal forest (MCF, Lechner (2019) and Lechner and Mareckova (2022)) as our main estimator, and we also use the classical binary choice model (McCullagh 1980) as robustness check. We pick MCF because it has good properties in terms of its computational easiness and various aggregation levels of the causal effects. It is also powerful because it can estimate the grouped based treatment effects. This is important because traditional methods only estimate average treatment effects based on individualized ones, but the individualized effects are always distributed dispersed.

Denote by $Y_{ist} = \{0,1\}$ a binary indicator for high performance by rider $i \in I$ in stage $s \in \{1,\dots,21\}$ at time $t \in \{2020,\dots,2024\}$, $D_{ist} = \{0,1\}$ is a treatment indicator equal to one if the stage occurs in the riders home country. Let X to be the set of all covariates[6], where $X := \{\tilde{X}, Z\}$. $z \in Z$ represents the categorical variables that we want to estimate the group aggregated effects, and $\tilde{x} \in \tilde{X}$ represents the other covariates. Identification relies on the key assumptions that we will elaborate in the following section. MCF estimates three kinds of average causal effects. In this paper, the following formulations (Rubin 1972; Lechner and Mareckova 2022) to depict our effects of interest. Since we have exactly one treatment and three sets of features, the estimation of effects can be given by:

$$IATE(D;x) = \mathbb{E}[Y(D=1) - Y(D=0) \mid X = x], \tag{1}$$

$$GATE(D;z) = \int IATE(D;\tilde{x},z) f_{\tilde{X}|Z=z}(\tilde{x}) \mathrm{d}\tilde{x}, \tag{2}$$

$$ATE(D) = \int IATE(D;x) f_X(x) \mathrm{d}x. \tag{3}$$

We are also interested in investigating the group-based categories $z \in Z$, including the rider's specialty, the role of captain, the profile of the stage, the home indicator of the team, and the area knowledge. The space here does not allow us to further explain our logistic regression approach, but they are same as the classical ones.

4.2 Identification Strategy

Our goal is to identify the causal effect of home advantage on individual rider performance in each stage of the Tour de France. Let $Y_{ist} = \mathbb{1}[\text{Final Stage Rank} \leq 15]$ to be the rider performance, and the key treatment $D_{ist} = \mathbb{1}[\text{Rider's Country} = \text{Stage Country}]$, where $\mathbb{1}$ is the indicator function, $i \in I$ represents heterogeneous riders, $s \in S$ represents the stage number, and $t \in \{2020,\dots,2024\}$ represents the race year. To interpret the

[6]This is the same as the term features in machine learning literature.

estimated coefficient on D to be a causal effect, we rely on the following identification assumptions (Lechner 2001).

Assumption 1 ***Conditional Independence.*** *The potential outcomes are independent of treatment assignment conditional on observables and fixed effects:* $\{Y_{ist}(1), Y_{ist}(0)\} \perp\!\!\!\perp D_{ist} \mid x$. *This implies* $\mathbb{E}[Y_{ist}(1) - Y_{ist}(0) \mid D_{ist}, x] = \mathbb{E}[Y_{ist}(1) - Y_{ist}(0)]$.

Assumption 1 is assumed to hold conditional on rider-, team-, and stage-level controls (and their combinations), i.e., a selection on observable setting. As reported in section 3, our set of control variables is very broad and complete and can be classified in 5 major clusters: (i) Tour features, (ii) rider's features, (iii) team features, (iv) stage features, and (v) combination of rider/team and stage features. These controls include both well-established confounders from the existing literature and novel variables that have not yet been widely adopted (e.g. area knowledge). We acknowledge that there might other critical determinants of rider performance currently unobserved to us. For example, the positioning of the rest days (generally, 2 per Tour) might influence the riders' performance in the previous and in the following stages. Nonetheless, we do not expect this to heterogeneously differ across riders. Another possible unobserved confounder is doping. However, Vandeweghe (2022) argues that, thanks to improved testing, detecting and blood profiling in the past 20 years, nowadays doping in cycling is practiced by less than 10% of the riders. This might be confirmed, for example, by three observations. First, in the past few years, all professional cyclists suspected for doping came from lower-tier teams Vandeweghe (2022). Second, in 2020, 2022, 2023 and 2024, only one case per year was reported while no case was reported in 2021[7]. Third, the climbing times in the mountains are definitely slower than 20 or 30 years ago (Vandeweghe 2022). Therefore, we can affirm that, given the very broad and complete set of controls included, it is reasonable to argue that assumption 1 holds.

Assumption 2 ***Common support.*** *Given the treatment D and all control variables set* Ω, A, Λ*, there exists a strict positive probability* $\mathbb{P}(D = d \mid X = x) = p_d(x)$*,* $\forall d \in \{0, 1\}$.

Assumption 2 ensures that, conditional on observed covariates, there is a strictly positive probability of receiving either treatment status. This assumption is critical to avoid extrapolation and to ensure that comparisons between treated and control units are based on comparable observations. It underpins the validity of both our parametric model (discrete choice estimation) and our causal machine learning procedures, which require sufficient overlap in the covariate distributions across treatment groups.

To assess the plausibility of this assumption, we estimate the propensity scores (i.e. the conditional treatment probabilities) and verify that meaningful overlap exists across the covariate space. We restrict our analysis to the region of common support, ensuring that all reported treatment effect estimates are based on comparable units with observed data support.

Figure 2 illustrates the region of common support in our sample. We find that 83.64% of the estimated treatment propensity distributions overlap across the two groups. To satisfy the common support assumption, we exclude observations that lie outside this region. In total, 1,261 observations are removed, and the analysis proceeds on the remaining sample. In Table 3, we analyze the representativeness of the cut sample and its similarity with the full sample by means of balance check.

Assumption 3 ***Stable Unit Treatment Value Assumption (SUTVA).*** *For each unit i, the potential outcome depends only on its own treatment assignment:* $Y_i(\mathrm{D}) = Y_i(D_i) \quad$ *for all i. This implies that there is no interference between units and no hidden variation in treatment levels.*

SUTVA requires that the potential outcomes of each unit depend only on that units own treatment assignment and not on the treatment status of other units. This implies, for example, the absence of spillover effects between units. In literature so far, no direct evidence suggests that home advantage effects in individual competitive sports, such as cycling, exhibit spillovers. Furthermore, several aspects of the race structure allow us to rule out potential spillovers by construction. First, the majority of the audience in each stage is (country) local so the psychological or motivational benefits of a home crowd are unlikely to be transferred to non-home riders (except for the "home team effect" channel, for which we control for). Second, many mountain stages – often the most decisive ones in the competition – cross also more remote areas that, because of terrain and

[7] For more details, see Vandeweghe (2022) and Wikipedia (2025).

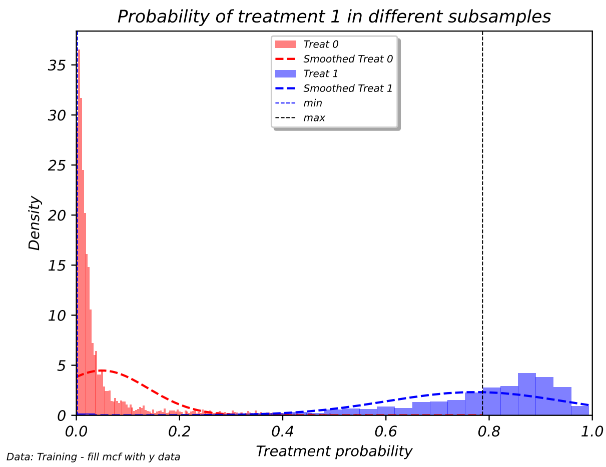

Figure 2: **Common Support Analysis for the Treatment**

NOTE: This figure illustrates the overlapped region of non-treated and treated samples. The vertical lines represent the overlapped region and the horizontal curves represent the smoothed distributions of the treatment distribution.

race speed, have little audience presence compared to stages in urbanized areas. Taken together, these factors support the plausibility of assumption 3 in our setting: while treatment effects may be heterogeneous across riders, the treatment status of one rider does not influence the outcome of another.

Assumption 4 ***Exogeneity of the confounders.*** *Given the treatment D and confounders X, it holds that $D \perp\!\!\!\perp X$.*

The exogeneity assumption holds if the observed values of the confounders are not effected by the treatment. We have carefully considered the listed controls in the previous section 3. Moreover, we removed from the original dataset all the potentially endogenous confounders (e.g. the final Tour rank, the cumulative final time and all the related time gaps, the cumulative bonus point obtained). Hence, it is reasonable to assume that assumption 4 holds as well.

4.3 Expected Results

The home advantage can be decomposed into several distinct components: familiarity with the region, support from the local crowd, and cultural recognition (Lucía et al. 2001; Jones 2013). In our study, we control for regional familiarity, allowing us to focus on the potential effects of crowd presence.

The literature on crowd effects remains inconclusive, particularly because different types of sports involve varying psychological and physiological demands. While one might expect a positive influence from crowd supportespecially for home athletes who may perceive local cheering as more personally motivatingthe evidence is mixed. Fischer and Haucap (2021) consider the crowd effect during the COVID-19 pandemic for German soccer team and found that, during the pandemic, the absence of crowd support reduced the home advantage for the first division players, but not for the second and the third. Psychological studies suggest that expectation pressure and collective fan behavior can, in some cases, impair athletic performance (for an overview, see Earnheardt et al. 2011).

For instance, performance may suffer if an athlete experiences psychological stress stemming from social and mass media exposure (MacPherson and Kerr 2021). Similarly, strong fan bases and heightened expectations from coaches or the public may lead athletes to overcommit, continuing to compete through minor injuries that, while not incapacitating, can still compromise performance (Endo et al. 2023).

These conflicting dynamics highlight that the net impact of home crowd support is likely heterogeneous and context-dependent.

5 Rider and Team Home Effect

We perform our estimation based on the sample with common support. A descriptive statistics is reported in Table 3. We will first present the results from MCF, and then append the binary choice estimation as a check of robustness complimentarily.

Table 3: **Mean Value by Full and Common Support Sample**

Variable	Full Sample		Common Support Sample	
	Untreated	Treated	Untreated	Treated
Area knowledge (%)	0.6	9.3	0.6	5.7
BMI (kg/m^2)	21.2	20.6	21.2	20.7
Captain (%)	6.8	8.8	6.9	4.8
Stage distance (km)	159.7	157.9	159.2	157.1
Tour completion (%)	50.7	53.2	50.5	52.2
Stage profile score	2.9	2.9	2.9	2.8
Stage rank	80.4	80.2	80.7	81.4
Rider's points previous year	479.9	392.3	481.0	438.6
Rider's specialty score	2.2	2.2	2.2	2.2
Rider's specialty TT (%)	21.5	10.6	20.3	12.6
Specialty match current stage (%)	35.7	34.3	35.7	35.6
Specialty match next stage (%)	33.6	33.4	33.6	32.0
Specialty match previous stage (%)	34.5	34.3	34.4	39.2
Stage type RR (%)	90.0	90.2	89.7	88.9
Team home country (%)	9.7	80.0	10.1	43.9
Team points previous year	3196.9	2646.2	3189.5	3218.2
Top 15 (%)	9.9	7.4	10.0	7.0
TT match current stage (%)	2.3	1.4	2.3	1.3

5.1 Average and Individualized Treatment Effects

We begin by reporting the average treatment effect (ATE) in the sample. The estimated ATE is -0.041 (-4.1%) with a p-value of 0.307%, suggesting a statistically significant negative effect of the treatment. In our context, this implies that the home effect on average hinders rider performance.

However, the treatment effect is likely to vary across individuals as shown in Figure 3. Figure 3a shows the estimated distribution of individualized average treatment effects (IATEs), sorted by quantile, and a comparison to ATE. We observe substantial heterogeneity: while the mean effect is negative, approximately 9.98% of IATEs are positive. The distribution is roughly symmetric around the mean. Moreover, 37.65% of riders exhibit IATEs that are significantly different from zero at the 5% significance level. Figure 3b shows the density of the distribution, where the red line represents ATE. We can see a relatively balanced distribution.

These findings suggest that the negative average effect of home effect is not uniform: for a subset of riders, the home environment is particularly detrimental. This result aligns with psychological research on performance pressure and crowd effects (Böheim et al. 2019; Harb-Wu and Krumer 2019; Scoppa 2021; Endo et al. 2023).

We next examine how IATEs vary with rider- and race-specific covariates to understand treatment effect heterogeneity. Figure 4 illustrates fitted trends of IATEs across selected variables, and we also draw a fitted line regressing the variable of interest on IATEs with an intercept.

Figure 4a shows a weak positive trend between BMI and IATE, though the estimated coefficient is statistically insignificant. This suggests that while higher BMI may be marginally associated with better performance under home conditions; the relationship is inconclusive. Figure 4b presents a more pronounced and significant negative relationship between IATEs and a riders previous-year PCS points. This finding supports the hypothesis that riders with higher prior performance face greater public expectations and psychological pres-

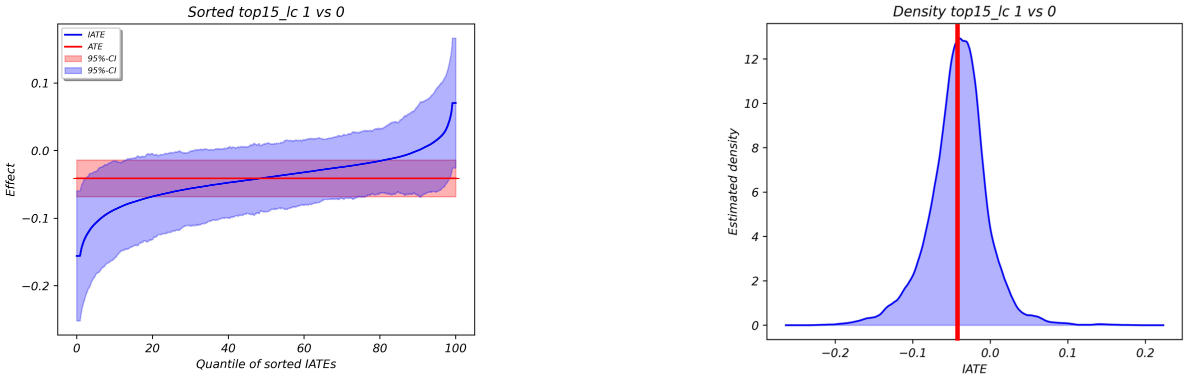

(a) Sorted Distribution

(b) Density Function

Figure 3: **Average Treatment Effect with Individualized Effect Distribution**

NOTE: This figure shows the sorted distribution of IATEs. Red elements indicate the average treatment effect (ATE, -0.041). Blue elements show the IATEs, both with 95% confidence intervals. The left panel is the sorted distribution and the right panel is the density.

sure, leading to more negative home effects. Conversely, riders with fewer past points may feel less external pressure, rendering them less susceptible to home-related performance deterioration. We also found similar trend for the team-level IATEs, although not statistical significant: one explanation is that team-level points average over heterogeneous individual experiences, thus ruling out any direct relation. Figure 4c illustrates a slight difference across different age groups, such that mid-young riders (24 - 30) have higher variances for effects, but getting converged for older riders, and very young riders (20 - 24) have also small variances. Our interpretation here is that the 24 - 30 age group is the major participant of the tour, which gives higher variances because many of them take part in this for the first time. But a relatively older rider who still takes part in the tour means that this person had good performances and better training in face of crowds, which explains the converge of the outliers. And a very young rider can take part in the tour because they are *particularly* good to be selected into a world tour team and has the candidacy. Those riders may also have very good psychological support as well. Figure 4d displays the relationships of IATEs with stage distance. No strong patterns are observed. Regarding stage distance, the flat trend is consistent with the races pre-determined structure, which is not likely to interact with home effects. Notably, Figure 4d reveals two clusters: one corresponding to time trial (TT) stages, typically held in cities or in mountain with more crowd presence, and another corresponding to longer stages that possibly cross also more remote areas. The fitted line for the TT cluster indicates a lower average IATE, suggesting that higher crowd density may be associated with poorer performancefurther reinforcing the psychological crowd-pressure hypothesis.

5.2 Group Heterogeneity

The heterogeneity in the IATEs suggests the necessity of considering group-based variations in treatment effects. We examine several categorical variables — rider specialty, team home effect, captain status, and stage profilethat may explain effect heterogeneity at the aggregate level. Figure 5 presents the estimated results.

Figures 5a and 5b illustrate the group-aggregated effects by captain status. The home effect is –0.034 for non-captain riders (p-value = 1.19%) and –0.052 for captain riders (p-value = 1.94%). However, the difference between these ATEs is not statistically significant. Hence we cannot say we identify heterogeneity for the group of captains.

Figures 5c and 5d show the heterogeneity by rider specialty. The specialty variable ranges from 1 to 3, representing increasing affinity for hilly to mountainous terrains, based on past performance. Statistically significant negative effects are observed for riders specialized in hilly (–0.035, $p = 0.32\%$) and mountainous regions (–0.062, $p = 0.02\%$). Sprinters (specialty = 1) exhibit a slightly positive effect (0.03, $p = 0.66\%$) in terms of its heterogeneity to the ATE, while mountainous riders show further negative deviations (–0.02, $p = 0.97\%$). This is also interpretable because the sprinters specialty is *endogenously* determined by their previous race results. This implicitly suggests that those sprinters are better prepared for those races are classified

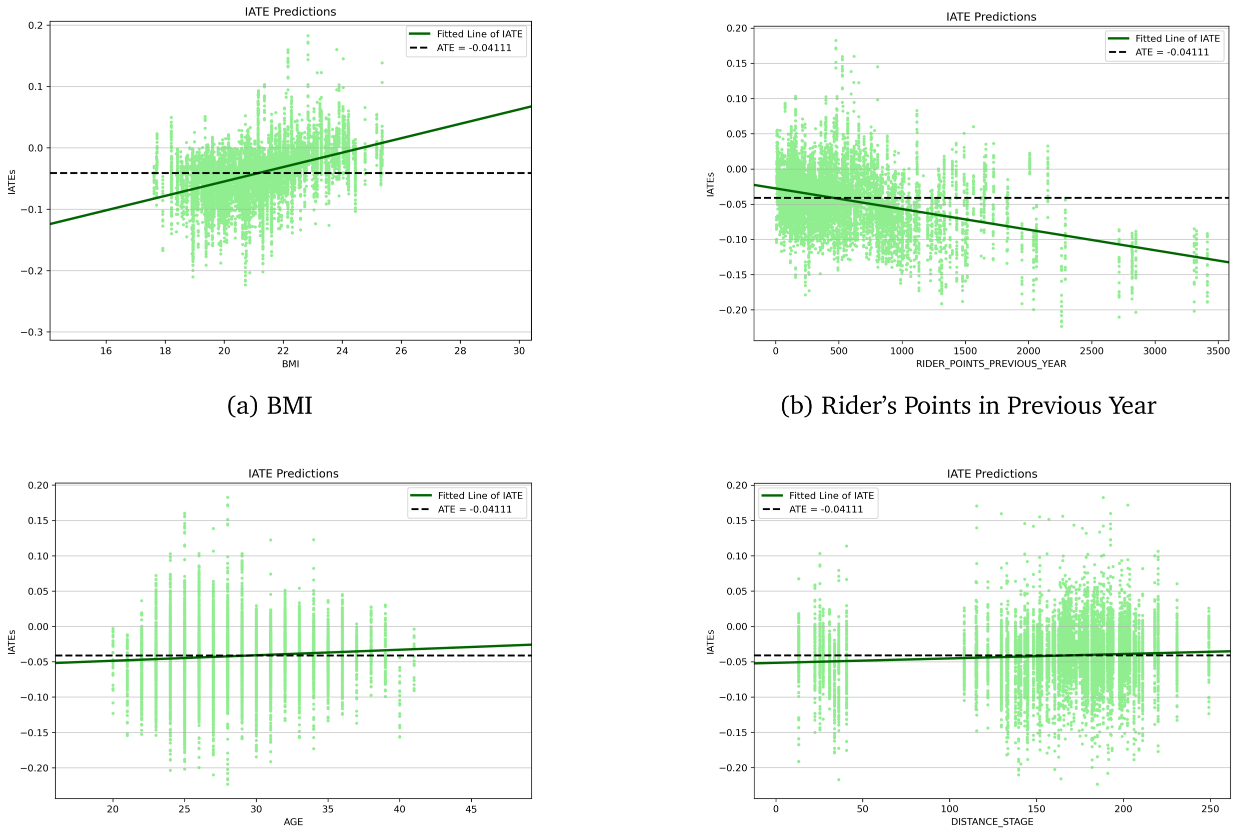

(a) BMI

(b) Rider's Points in Previous Year

(c) Age

(d) Stage Distance

Figure 4: **Heterogeneity of IATEs with Respect to Key Covariates**

NOTE: Each panel shows the relationship between IATEs and a covariate. Light green dots represent individual IATE values; the dashed black line indicates the ATE; the solid green line shows the fitted OLS trend. Only selected interpretable covariates are shown.

as sprinting ("flat") stages, which have higher audiences and crowd along the race because they are always going across city centers and major roads. Those people, according to our theory of psychological pressure, by definition are better prepared than the others. And this effect is indeed eased with the increasing of the steepness, i.e., less access for the audience.

In contrast, Figures 5e and 5f, which group by stage profile (parcours), show that while all GATEs are significantly different from zero at the 99% level, the differences across groups are not statistically distinguishable from zero. This suggests that the home effect is largely invariant to external stage conditions and is more influenced by rider-specific attributes. This conclusion aligns with our earlier finding that the rider's area knowledge variable was not statistically significant.

Finally, Figures 5g and 5h explore the role of team home status. Riders not cycling for a home team experience home effect (-0.041,p-value 0.35%) if they are home riders, but riders cycling for a home team do not experience it. Meanwhile, the differences of such effect to the ATE are not statistically distinguishable.

5.3 Home Effect is Individually Psychological

The results we had presented above indicate one thing that we identified: home effect is heterogeneous across people. In general, it is negative, but is eased on those who are younger and who have higher chances to be exposed in the audience. And it is especially severe for those who have high records from the previous years, because they potentially experience higher public expectations and overcommitment. This corresponds to other scholars' research results, as argued before.

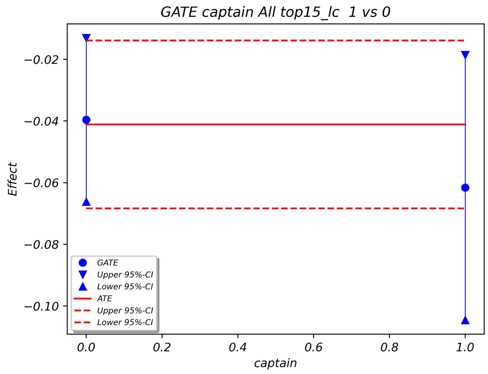

(a) Captain (True Effect)

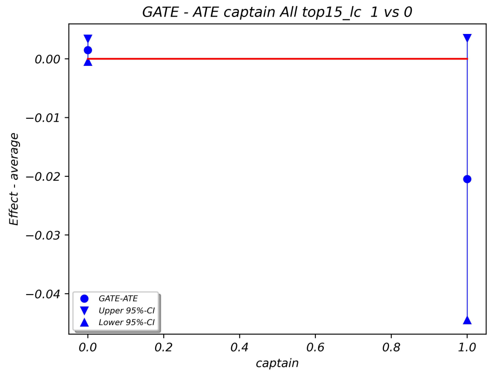

(b) Captain (Difference)

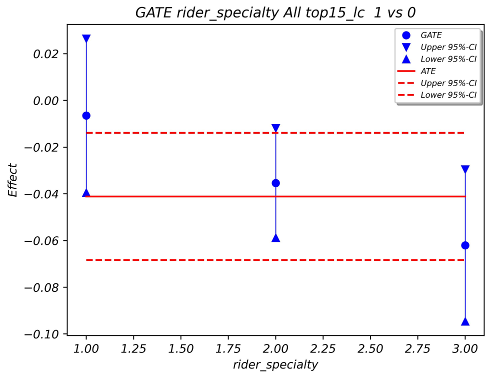

(c) Rider's Specialty (True Effect)

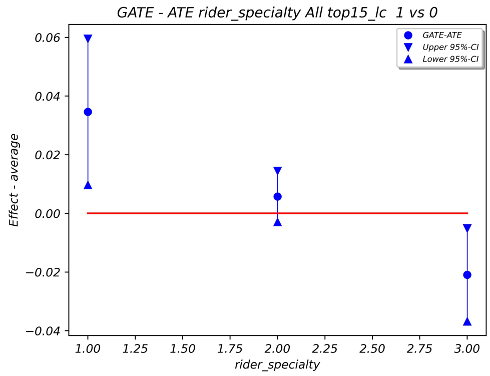

(d) Rider's Specialty (Difference)

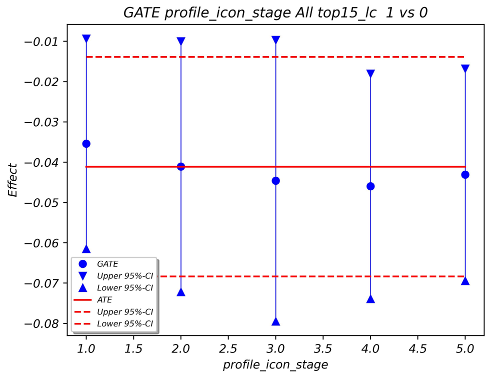

(e) Parcours Profile (True Effect)

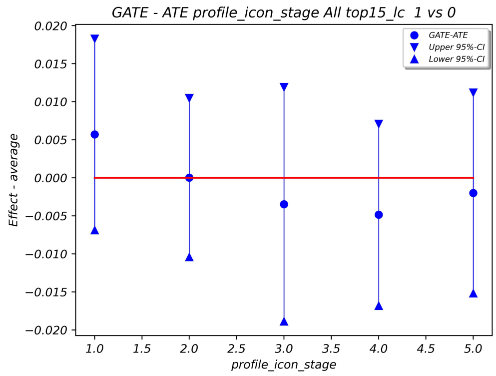

(f) Parcours Profile (Difference)

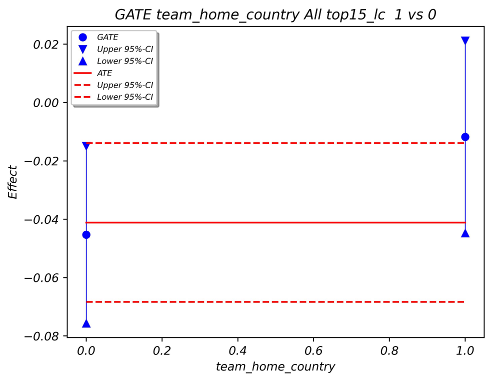

(g) Team Home Effect (True Effect)

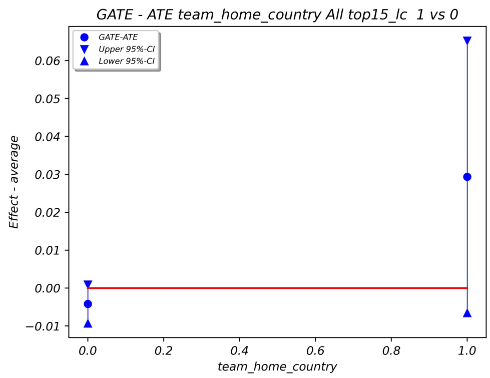

(h) Team Home Effect (Difference)

Figure 5: **True Effects and Differences in GATEs**

Note: This figure displays the group average treatment effects (GATE) for different variables. The left panel presents the GATE while the right panel shows its difference from the ATE (both with 95% confidence intervals).

5.4 Results of Binary Choice: A Robustness Check

As a robustness check, we estimate a binary choice model using logistic regression. This classical model is valued for its interpretability and its ability to assess the association between covariates and the binary outcome.

The specification mirrors that of the MCF model, with log-transformations applied to stage distance, and rider and team points from the previous year to address scale differences. Due to space limitations and the limited interpretability of log-odds, we report only marginal effects (Table 4).

Table 4: Marginal Effects from Logistic Regression (in Percentage Points)

Variable	**Marginal Effect (%)**	***p*-value**
year_tour_2021	3.09^{***} (0.66)	0.000
year_tour_2022	1.88^{**} (0.60)	0.002
year_tour_2023	0.57 (0.60)	0.346
year_tour_2024	0.62 (0.62)	0.316
rider_points_previous_year	4.82^{***} (0.23)	0.000
bmi	0.17 (0.12)	0.164
captain	3.74^{***} (0.58)	0.000
age	-0.34^{***} (0.05)	0.000
team_points_previous_year	–0.33 (0.42)	0.438
distance_stage	0.21 (0.33)	0.520
profile_icon_stage	-0.44^{***} (0.13)	0.001
perc_tour_completed	0.02^{***} (0.01)	0.001
rider_home_country	–1.19 (0.84)	0.154
team_home_country	-1.60^{*} (0.65)	0.014
tt_match	8.60^{***} (1.14)	0.000
specialty_match	7.99^{***} (0.39)	0.000
specialty_match_prev	-1.08^{**} (0.39)	0.006
specialty_match_next	–0.69 (0.39)	0.079
area_knowledge	–2.29 (2.31)	0.321

Note: We report marginal effects in percentage points, with standard errors in parentheses. Stars represent the significance level: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

The treatment variable *rider_home_country* shows a negative coefficient, consistent with the MCF estimation. However, the effect is not statistically significant, with a p-value of 17.6%. This may reflect model limitations, suggesting that logistic regression is not ideally suited for capturing the home effect in our context.

We identify a statistically significant negative association for *team_home_country* and a significant positive effect for *captain*, aligning with expectations regarding team strategy. Captains, often the strongest or most prominent riders, benefit more from team support, which may translate into higher likelihoods of finishing in the top 15.

The variable *area_knowledge* is not statistically associated with the outcome. This reinforces our core argument: the home effect is driven primarily by psychological pressure rather than geographic familiarity. Furthermore, given that many Tour de France stages (especially in mountainous regions) are repeated annually and announced in advance, prior local knowledge likely has limited marginal value.

Finally, we examine stage-rider specialty matching. We find positive and statistically significant associations for both *specialty_match* and *tt_match*, indicating that alignment between a rider's specialty and the stage profile increases the probability of a top finish. Interestingly, *specialty_match_prev* shows a small but significant negative effect ($p = 2.3\%$), possibly reflecting rider fatigue from prior exertion in favorable stages, which may impair performance in subsequent ones.

We acknowledge that logistic regression may not be the optimal specification in our context. This limitation arises from the binary transformation: riders ranked 16th and 150th are both coded as zero, despite representing different levels of performance. Although we considered alternatives such as ordered logit models, they require setting thresholds to define performance categories – a process that introduces arbitrariness again. These concerns again highlights the advantage of MCF approach, which does not impose a parametric assumption on $\mathbb{E}(Y \mid D)$ and therefore allows for a more flexible estimation of treatment effects.

6 Conclusions

There has been long discussion of the home effects to team competition effects, with limited attention in individual sports (Jones 2013). Our research extends the analysis of teamed-up individual competitions based on professional cyclings. We exploit a 5-year dataset of Tour de France to study the stage-level impact of home race. Our major result is that the home effect, in fact, decreases home riders' probability of getting into the top 15 of a given stage, which is useful not only for the general Tour classification but also grants them the UCI professional points and the visibility needed to get a renewal or a new contract.

We analyze the cause of this and find that this could be related to their social and public pressure, especially for those who have higher achievements (and, potentially, higher public exposure) in the past year. We identify an negative average effect of -4.1%, which suggests that the home effect slightly hinders professional cyclists' changes to rank in the top 15 of a given stage. We also found out that such effect is eased when riders have specialties of sprinting, as they tend to be more exposed to the public audience.

Our research still has limitations in terms of the data size and availability. Even though we consider psychological factors to be the main driving issue, we are not able to acquire their pressure level per se, or find variables to proxy such factor. Further, since we considered only the *Tour de France*, our treatment variable (and so, the home effect) is strongly concentrated on French riders, which makes it difficult to generalize the results for the other non-French cyclists riding a stage in their home country. Similarly, we cannot take any conclusion about home effect in cycling in general, i.e. regardless of the *Grand Tour* considered and the nationality of the rider. The next step would be to investigate how different factors influence the pressures heterogeneously and generalize our results by including other *Grand Tours* (e.g. *Giro d'Italia*, *La Vuelta*).

References

Böheim, R., D. Grübl, and M. Lackner (2019). "Choking under pressure Evidence of the causal effect of audience size on performance". In: *Journal of Economic Behavior & Organization* 168, pp. 76–93.

Brandenburger, A. M. and B. J. Nalebuff (2011). *Co-Opetition*. Crown. 305 pp.

Breteau, P. P. (2024). "Tour de France de 1903 à 2025 : quels sont les départements gagnants et perdants du tracé ?" In.

Correia-Oliveira, C. R. and V. A. Anrade-Souza (2022). "Home advantage in soccer after the break due to COVID-19 pandemic: does crowd support matter?" In: *International Journal of Sport and Exercise Psychology* 20.4, pp. 1245–1256.

Courneya, K. S. and A. V. Carron (1992). "The Home Advantage In Sport Competitions: A Literature Review". In.

Earnheardt, A., P. Haridakis, and B. Hugenberg (2011). *Sports Fans, Identity, and Socialization: Exploring the Fandemonium*. Lexington Books. 313 pp.

Endo, T., H. Sekiya, and C. Raima (2023). "Psychological pressure on athletes during matches and practices". In: *Asian Journal of Sport and Exercise Psychology*. Relationships between Physical Activity, Psychological and Cognitive Functioning: Evidence and Health Implications 3.3, pp. 161–170.

Fischer, K. and J. Haucap (2021). "Does Crowd Support Drive the Home Advantage in Professional Football? Evidence from German Ghost Games during the COVID-19 Pandemic". In: *Journal of Sports Economics* 22.8, pp. 982–1008.

Harb-Wu, K. and A. Krumer (2019). "Choking under pressure in front of a supportive audience: Evidence from professional biathlon". In: *Journal of Economic Behavior & Organization* 166, pp. 246–262.

Jones, M. B. (2013). "The home advantage in individual sports: An augmented review". In: *Psychology of Sport and Exercise* 14.3, pp. 397–404.

Koning, R. H. (2011). "Home advantage in professional tennis". In: *Journal of Sports Sciences* 29.1, pp. 19–27.

Larson, D. and J. Maxcy (2013). *Uncertainty of Outcome and Radio Policy in Professional Road Cycling*. Rochester, NY.

Larson, D. J. and J. G. Maxcy (2016). “Human Capital Development in Professional Cycling”. In: *The Economics of Professional Road Cycling*. Ed. by D. Van Reeth and D. J. Larson. Cham: Springer International Publishing, pp. 129–145.
Lechner, M. (2001). “Identification and estimation of causal effects of multiple treatments under the conditional independence assumption”. In: *Econometric Evaluation of Labour Market Policies*. Ed. by M. Lechner and F. Pfeiffer. Heidelberg: Physica-Verlag HD, pp. 43–58.
Lechner, M. (2019). *Modified Causal Forests for Estimating Heterogeneous Causal Effects*. arXiv: 1812.09487 [econ].
Lechner, M. and J. Mareckova (2022). *Modified Causal Forest*. arXiv: 2209.03744 [econ].
Lucía, A., J. Hoyos, and J. L. Chicharro (2001). “Physiology of Professional Road Cycling”. In: *Sports Medicine* 31.5, pp. 325–337.
MacPherson, E. and G. Kerr (2021). “Sport fans responses on social media to professional athletes norm violations”. In: *International Journal of Sport and Exercise Psychology* 19.1, pp. 102–119.
Matthes, J. and D. Piazolo (2024). “Dont put all your legs in one basket: Theory and evidence on coopetition in road cycling”. In: *European Economic Review* 170, p. 104892.
McCullagh, P. (1980). “Regression Models for Ordinal Data”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 42.2, pp. 109–142.
Menaspa, P., C. Abbiss, and D. Martin (2013). “Performance Analysis of a World-Class Sprinter During Cycling Grand Tours”. In: *International Journal of Sports Physiology and Performance* 8.3, pp. 336–340.
Nevill, A. M., R. L. Holder, A. Bardsley, H. Calvert, and S. Jones (1997). “Identifying home advantage in international tennis and golf tournaments”. In: *Journal of Sports Sciences*.
Phillips, K. E. and W. G. Hopkins (2020). “Determinants of Cycling Performance: a Review of the Dimensions and Features Regulating Performance in Elite Cycling Competitions”. In: *Sports Medicine - Open* 6.1, p. 23.
Pollard, R. and G. Pollard (2005). “Long-term trends in home advantage in professional team sports in North America and England (1876 2003)”. In: *Journal of Sports Sciences* 23.4, pp. 337–350.
Ramchandani, G. and D. Wilson (2020). “Does a home advantage effect exist in international track and field athletics?” In: *Managing Sport and Leisure* 25.5, pp. 376–389.
Rebeggiani, L. (2016). “The Organizational Structure of Professional Road Cycling”. In: *The Economics of Professional Road Cycling*. Ed. by D. Van Reeth and D. J. Larson. Cham: Springer International Publishing, pp. 33–54.
Rubin, D. (1972). “Estimating Causal Effects of Treatments in Experimental and Observational Studies”. In: *ETS Research Bulletin Series* 1972.2, pp. i–31.
Scoppa, V. (2021). “Social pressure in the stadiums: Do agents change behavior without crowd support?” In: *Journal of Economic Psychology* 82, p. 102344.
Torgler, B. (2007). “La Grande Boucle: Determinants of Success at the Tour de France”. In: *Journal of Sports Economics* 8.3, pp. 317–331.
Van Reeth, D. (2022a). “Globalization in Professional Road Cycling”. In: *The Economics of Professional Road Cycling*. Ed. by D. Van Reeth. Cham: Springer International Publishing, pp. 337–367.
Van Reeth, D. (2022b). “The Finances of Professional Cycling Teams”. In: *The Economics of Professional Road Cycling*. Ed. by D. Van Reeth. Cham: Springer International Publishing, pp. 35–67.
Van Reeth, D. and D. J. Larson, eds. (2016). *The Economics of Professional Road Cycling*. Vol. 11. Sports Economics, Management and Policy. Cham: Springer International Publishing.
Vandeweghe, H. (2022). “Doping in Cycling”. In: *The Economics of Professional Road Cycling*. Ed. by D. Van Reeth. Cham: Springer International Publishing, pp. 283–308.
Wikipedia (2025). *List of doping cases in cycling*. In: *Wikipedia*.

Author Contributions Declaration

The authors' contributions to this project are organized as follows:

- **Conceptualization of the research topic**: Eren Celik, Xiaoyi Guo, Angelo Mimmo
- **Data scraping from API and geocoding**: Xiaoyi Guo
- **Data curation and feature engineering**: Angelo Mimmo
- **Empirical strategy, identification, and assumption**: Xiaoyi Guo, Angelo Mimmo
- **Estimation, statistical description, and visualization**: Eren Celik
- **Documentation organization, writing, and review**: Eren Celik, Xiaoyi Guo, Angelo Mimmo

We hereby confirm that the above allocation of contributions accurately reflects the actual division of work. The chronological sequence in which the work was carried out corresponds to the order in which the contributions are listed.

Date and Place: 20.05.2025, St. Gallen

Signatures of the authors:

Celik, Eren	Guo, Xiaoyi	Mimmo, Angelo