

The (Economics and) Determinants of Free-Riding Potential and Utilization: Field Evidence from Triathlon

Felix Reichel*

January 23, 2026

Abstract

This paper studies the determinants (and economics) of free-riding in the swimming stage of mass-start triathlon. The theoretical framework builds upon non-discriminatory Tullock contests, modeling triathlon as a simplified two-stage game with quadratic effort costs, and extends to n -player and asymmetric settings in which drafting multiplicatively reduces heterogeneous effective marginal costs. Empirically, I construct an event-study measure of free-riding based on relative rank changes across the two stages to capture free-riding propagation, centered on the exceptional 2020 season, during which non-drafting rules were enforced due to the COVID-19 pandemic.

Causal identification combines pooled OLS (POLS) with athlete and event fixed effects (FE), as well as a regression discontinuity design (RDD) exploiting COVID-19 policy-induced rule changes—specifically, individually staggered starts—that mechanically reduced drafting potential and, consequently, its utilization. Additionally the economic analysis chapter uses elasticity, production frontier, and conditional quantile methods to primarily examine group-size effects.

The results show that: (i) free-riding was substantially positive prior to 2020, collapsed during the 2020 season, and only partially recovered thereafter; (ii) the drafting effects are stronger for weaker swimmers and athletes in deeper drafting positions, with modest gender differences; and (iii) age heterogeneity is pronounced, with middle-aged and older cohorts benefiting the most, suggesting that athlete experience plays a vital role in efficiently utilizing free-riding gains. However, this latter finding may also be significantly confounded by compositional changes in athletes' body characteristics over time.

Keywords: Free-riding; Contest theory; Tullock contests; Sports economics; Drafting externalities; Group-size effects; Event-study; Regression discontinuity design

JEL Codes: D72, L83.

*Department of Economics, Johannes Kepler University Linz

Contents

I	Introduction	9
II	Related Literature	11
III	Theoretical Framework	13
III.1	Players, Timing, and Information	13
III.2	Preferences and Outside Options	13
III.3	Stage 2: Tullock Contest with Drafting as a Cost Shifter	14
III.3.1	Best responses	14
III.3.2	Two-player closed form	14
III.3.3	Symmetric m -player benchmark	15
III.3.4	Asymmetric m -player equilibrium	15
III.4	Stage 1: Continuation Decisions and Subgame-Perfect Equilibrium	16
III.4.1	Continuation values	16
III.4.2	Cutoff rules	16
III.4.3	Stage 1 equilibrium and SPE existence	16
III.5	Drafting as a Cost-Shifting Technology	17
III.6	Welfare and Rent Dissipation	17
III.7	Testable Predictions	18
III.8	Figures	19
IV	Panel Data, Sample Restriction, and Variable Construction	20
IV.1	Core Variables	20
IV.2	Sample Construction	20
IV.3	Derived Measures	21
IV.4	Sample Restriction Procedure	22
IV.5	Implementation Notes	22
V	Descriptive Statistics	24
V.1	Overview and Unit of Analysis	24
V.2	Covariates: Original Dataset (Pre-Pipeline)	24
V.3	Covariates: Analysis Dataset (Post-Pipeline)	25
V.4	Outcomes: Free-Riding	25
V.5	Temporal Patterns by Heterogeneity Dimension	27
V.6	Heterogeneity Maps relative to 2019 (3D Contour Plots)	28
VI	Identification Strategy	29
VI.1	Evidence on Restricted Free-Riding Potential during COVID-19 (2020)	29
VI.2	Regression Discontinuity Design (RDD)	31
VI.3	Regression Discontinuity Design (RDD) on Rank-Improvement (2019 vs 2020)	32
VI.4	Identification via Rank-Improvement Probabilities	32
VII	Distributional Effects of the COVID Shock on Athletic Performance	37
VII.1	Conceptual Framework	37
VII.2	Empirical Design	37
VII.3	Fuzzy Quantile Regression Discontinuity	37
VII.4	Robustness and Validation	37
VIII	Empirical Framework	41
VIII.1	OLS with High-Dimensional Fixed Effects: Convexity, Grouped-Gram Estimation, and Fast Inference	41
VIII.2	Empirical Analysis using 2020 as reference year	42

VIII.3	Baseline Pooled Yearly Specification	43
VIII.4	Alternative Pooled Period Specification	44
IX	Empirical Results	46
IX.1	Effect Heterogeneity by Gender	46
IX.2	Effect Heterogeneity by Age	47
IX.3	Effect Heterogeneity by Drafter Position	48
IX.4	Effect Heterogeneity by Swimmer Ability	48
IX.5	Effect Heterogeneity by Group Gender Composition Mix	50
IX.6	Effect Heterogeneity by Group Size	50
X	Estimation Framework using 2019 as reference year	52
X.1	Heterogeneous Effects by Swim Ability on Free Riding	54
X.2	Heterogeneous Effects by Age on Free Riding	55
X.3	Age heterogeneity from the baseline FE model	55
X.4	Heterogeneous Effects by Drafter Position on Free Riding	57
XI	Robustness: Significant Interactions in Period-Pooled OLS Event Study	58
XII	On the Economics of Free Riding in Triathlon	62
XII.1	The Elasticity of Free Riding as a Public Good Problem	62
XII.2	The Production Frontier of Free Riding	64
XII.3	Engel Curve	66
XII.4	Beveridge Curve Analogy: Group Size and Slack	68
XII.5	Growth Incidence of Free Riding	71
XIII	Discussion	73
XIV	Conclusion	74
	Appendix	79
A.1	Data Dictionary	80
A.1.1	Original variables (from <i>data/preprocessed.csv</i>)	80
A.1.2	Computed variables (analysis pipeline)	82
A.2	Non-parametric Construction of Swimmer-Ability Labels (K-means Clustering)	83
A.3	Regression Results: Effect Heterogeneity	88
A.3.1	Heterogeneous Effects by Gender on Free Riding	88
A.3.2	Heterogeneous Effects by Swim Ability on Free Riding	89
A.3.3	Heterogeneous Effects by Age on Free Riding	90
A.3.4	Heterogeneous Effects by Drafter Position on Free Riding	91
A.4	Identification Robustness: CausalWizard.app Results For the Covid-19 Policy Shock As Treatment	92

List of Tables

1	Core variables in the event–athlete panel.	20
2	Descriptive statistics — original dataset sample	24
3	Summary statistics after pipeline ($n = 15,028$)	26
4	Categorical level counts after pipeline ($n = 15,028$)	26
5	Welch two-sample t-tests of mean rank improvement by period and swim skill	31
6	Sharp RDD estimates of rank improvement at the 2019/2020 cutoff	33
7	Baseline FE regression (outcome: free-riding)	43
8	Baseline FE regression (outcome: rank-improvement)	44
9	Pooled OLS with athlete and event-category FE (periods = pre/2020/post). Coefficients grouped by period and dimension, ordered by category values (e.g., age bins ascending; drafter positions 3–7, then other; pack sizes small→medium; team mix in the order shown). Estimate cells use a blue–white–red gradient with intensity proportional to $ \hat{\beta} $. Baselines are 2020 (period) and, within each dimension v , the reference category $\ell_0(v)$	58
12	Heterogeneous Effects by Gender on Free Riding	88
13	Heterogeneous Effects by Swim Ability on Free Riding	89
14	Year Fixed Effects on Free Riding by Age Bin	90
15	Year Fixed Effects on Free Riding by Drafter Position	91

List of Figures

1	Extensive Form of the Two-Stage Contest	19
2	Equilibrium Effort as a Function of Drafting	19
3	Flow of sample construction and derived measures in the event-athlete panel. The process includes data import, identification of $\mathcal{A}_{2020}$, application of filters (IV.1)–(IV.3), computation of $r_{a,e}^{\text{swim}}$ and $\Delta r_{a,e}$, periodization, definition of $FR_{a,e}$, and creation of k-means swim-ability clusters, resulting in the final dataset $\widehat{\mathcal{S}}$	23
4	Swim-out times by event category and period (original sample). Boxes show the interquartile range; whiskers extend to $1.5 \times \text{IQR}$	25
5	Temporal evolution of mean free-riding potential by subgroup and interaction dimensions. Each panel shows cell means without smoothing to highlight shifts around 2020.	27
6	Filled-contour descriptives: change in mean free-riding relative to 2019 across six dimensions.	28
7	Distributions of rank improvements by swim skill and period (pre, 2020, post). Conclusion: Pre $\succeq$ Post $\succeq$ 2020	30
8	Comparison of RDD fits using different polynomial specifications.	32
9	Share of TOP-5 athletes improving rank via drafting by period. Bars show proportions; whiskers denote 95% t-based confidence intervals.	34
10	Draft-induced rank improvements among TOP-10 athletes by two-year bins; the COVID season (2020) is shown separately.	34
11	Annual share of TOP-10 athletes improving rank via drafting, 2010–2024, with 95% confidence intervals.	35
12	Share of TOP-20 athletes improving rank via drafting by period.	35
13	Share of TOP-30 athletes improving rank via drafting by period.	36
14	Fuzzy Quantile Regression Discontinuity at the 10th Percentile	38
15	Fuzzy Quantile Regression Discontinuity at the Median	38
16	Fuzzy Quantile Regression Discontinuity at the 90th Percentile	39
17	Bandwidth Sensitivity of Median Fuzzy QTE	39
18	Placebo Cutoff at 2018	40
19	Placebo Cutoff at 2019	40
20	Year $\times$ gender coefficients for free-riding (points = estimates, bars = 95% CIs).	43
21	Year $\times$ gender coefficients for rank-improvement (95% CIs).	44
22	Year $\times$ gender event-study for free-riding. Candlestick body = 95% CI, tick = estimate $\hat{\beta}_{y,h}$. FE at athlete and event-category; SEs clustered by athlete and event; reference year 2020.	46
23	Year $\times$ age-bin event-study for free-riding. Bins ordered by lower bound; can- dlesticks show 95% CIs; ref = 2020.	47
24	Year $\times$ drafter position (bins) event-study for free-riding. 95% CIs; athlete and event-category FEs; clustered SEs; ref = 2020.	48
25	Year $\times$ swimmer ability event-study for free-riding (good/average/bad). 95% CIs; ref = 2020.	48
26	Year $\times$ team composition (gender mix) event-study for free-riding. Candlesticks = 95% CIs; ref = 2020.	50
27	Free-riding coefficients by group-size bins (95% CIs).	50
28	Rank-improvement coefficients by group-size bins (95% CIs).	51
29	Year FEs on Free Riding by Gender (Male vs. Female), baseline=2019	53
30	Year FEs by Swim Ability (Good/Bad/Average), baseline=2019.	54
31	Year FEs by Age bin (20–25, . . . , 60–65), baseline=2019.	55

32	Age heterogeneity (step function) including <i>all pack-size groups</i> . Lines show model-implied predictions by age bin using $\widehat{\Delta\text{RI}}(\alpha) = 0.0351(\alpha^2 - 25^2)$, with vertical shifts of +18.05 (medium) and +31.57 (large) relative to the small-group as baseline group.	56
33	Year FEs by Drafter Position (1–8), baseline=2019.	57
34	Pooled OLS interaction coefficients by period (pre/2020/post), sorted by $ \hat{\beta} $ within period $\times$ dimension. Points are colored by sign; vertical bars are 95% CIs. Baselines are 2020 (period) and, within each dimension v , the reference category $\ell_0(v)$	60
35	Elasticity sketch: $\log(1+\text{FR})$ vs. $\log(1+n)$. Estimated elasticity $\hat{\beta} \approx 0.427$ (95% CI: [0.393, 0.461]). For a finite change $n \rightarrow n'$, $\% \Delta(1+\text{FR}) = \left[\left(\frac{1+n'}{1+n} \right)^{\hat{\beta}} - 1 \right] \times 100\%$. In particular, doubling from n to $2n$ implies $\left[\left(\frac{1+2n}{1+n} \right)^{\hat{\beta}} - 1 \right] \times 100\%$, which approaches $(2^{\hat{\beta}} - 1) \times 100\% \approx 34.4\%$ for large n (95% CI $\approx [31.3\%, 37.6\%]$).	63
36	Estimated conditional quantile frontiers $Q_\tau(\text{FR} \mid r)$ at $\tau \in \{0.10, 0.50, 0.90\}$. The $\tau = 0.90$ curve serves as a proxy for the frontier $\text{FR}^*(r)$; higher τ values tighten the approximation, subject to sampling variability.	65
37	Engel-style curve: FR vs. $\log(1 + n)$, across periods (pre, covid, post).	66
38	Beveridge-style curve: group size n vs. slack S . Rising n initially expands matching opportunities ($\propto n^{1-\alpha}$) but eventually raises coordination costs $C(n)$; the interior maximum n^* solves $(1 - \alpha)\kappa_0 n^{-\alpha} = C'(n)$	70
39	Growth incidence curve: quantile-specific change in free riding (Post – Pre). Solid: $g(p)$; shading: simultaneous 95% band from the multiplier bootstrap.	71
40	Elbow plot: within-cluster dispersion W_k as a function of k	84
41	Average silhouette $\bar{s}(k)$	84
42	Gap statistic $\text{Gap}(k)$ with one-SE rule.	85
43	PCA projection of athletes with $k = 3$ clusters and 95% ellipses.	86
44	Cluster centroids across event categories (z-scored).	86
45	Swim times by cluster and event category (IQR boxes, $1.5 \times \text{IQR}$ whiskers). Median times follow <i>good < average < bad</i>	87
46	Cluster sizes (athlete counts).	87
47	Silhouette widths by athlete for $k = 3$	87
48	Directed acyclic graph (DAG) used for identification.	94

Declaration of Originality and Use of Artificial Intelligence

Erklärung zur Eigenständigkeit und Nutzung Künstlicher Intelligenz

English

This paper is the author's own original work. Assistive technologies, including Google Speech Dictation and the large language model *ChatGPT 5.2 Instant*, were used exclusively for transcription and minimal linguistic correction. All academic content, ideas, analyses, and conclusions were developed independently by the author.

In accordance with Directive (EU) 2019/790 and Regulation (EU) 2024/1689, the use of artificial intelligence tools does not affect the authorship or originality of this work. The paper has not been previously submitted or published.

Deutsch

Diese Arbeit wurde selbstständig und eigenverantwortlich verfasst. Zur Unterstützung wurden ausschließlich Hilfsmittel wie Google Speech Dictation sowie das große Sprachmodell *ChatGPT 5.2 Instant* zur Transkription und minimaler sprachlicher Korrektur verwendet. Sämtliche wissenschaftlichen Inhalte, Ideen, Analysen und Schlussfolgerungen stammen vom Autor.

In Übereinstimmung mit der Richtlinie (EU) 2019/790 sowie der Verordnung (EU) 2024/1689 bleibt die Urheberschaft und Originalität dieser Arbeit unberührt. Die Arbeit wurde weder ganz noch teilweise zuvor eingereicht oder veröffentlicht.

Glossary of Notation

This glossary consolidates the notation, symbols, and variables used throughout the paper. Indices: $\mathbf{a}$ or $\mathbf{i}$ denotes athletes/individual, $\mathbf{e}$ denotes events, and $\mathbf{t}$ denotes time/years/periods.

$\mathbf{a}$ or $\mathbf{i}$	Athlete identifier.
$\mathbf{e}$	Event identifier (unique race instance).
$\mathbf{d}(\mathbf{e})$	Event calendar date (ISO format).
$\mathbf{y}(\mathbf{e})$	Event year, derived from $\mathbf{d}(\mathbf{e})$.
$\mathbf{A}(2020)$	Set of athletes who competed at least once in 2020.
$\mathbf{m}(\mathbf{a}, \mathbf{e})$	Sex indicator: 1 = male, 0 = female.
$\mathbf{s}(\mathbf{a}, \mathbf{e})$	Swim split time (in seconds).
$\mathbf{r}^{\text{swim}}(\mathbf{a}, \mathbf{e})$	Within-event swim rank.
$\mathbf{r}^{\text{tot}}(\mathbf{a}, \mathbf{e})$	Final finishing rank.
$\mathbf{\ell}(\mathbf{a}, \mathbf{e})$	Swim leader indicator.
$\text{drafterPosition}(\mathbf{a}, \mathbf{e})$	Position within drafting swim group.
$\mathbf{g}(\mathbf{e})$	Swim-Group size in event $\mathbf{e}$.
$\Delta\mathbf{r}(\mathbf{a}, \mathbf{e})$	Simplified rank improvement: $\mathbf{r}^{\text{tot}}(\mathbf{a}, \mathbf{e}) - \mathbf{r}^{\text{swim}}(\mathbf{a}, \mathbf{e})$.
$\text{FR}(\mathbf{a}, \mathbf{e})$	Free-riding: $\Delta\mathbf{r}(\mathbf{a}, \mathbf{e}) - \Delta\mathbf{r}(2020)$ (centered at 2020).
period	Periods: <i>pre</i> (2015–2019), <i>COVID</i> (2020), <i>post</i> (2021–2024).
age, age^2	Athlete age in years; squared age term.
ageBin	Age categorized into 5-year bins.
groupSizeBin	Categorized swim group size: small, medium, large.
swimSkill	Ability clusters obtained via non-parametric k-means.
rankImprovement	Synonym for $\Delta\mathbf{r}(\mathbf{a}, \mathbf{e})$.
freeRiding	Rank improvement relative to 2020 baseline.
$\mathbf{Y}(\mathbf{i}, \mathbf{e}, \mathbf{t})$	Generic outcome (rank improvement RI or free-riding FR).
$\text{FR}(\mathbf{i}, \mathbf{e}, \mathbf{t})$	Free-riding outcome centered at 2020.
$\text{RI}(\mathbf{i}, \mathbf{e}, \mathbf{t})$	Raw rank improvement.
$\beta(\mathbf{y}, \mathbf{h})$	Coefficient on year $\mathbf{y}$ and heterogeneity group $\mathbf{h}$.
$\gamma^{(1)}, \gamma^{(2)}$	Coefficients on age and age^2 .
$\delta^{(1)}, \delta^{(2)}$	Coefficients on medium and large group dummies.
$\alpha(\mathbf{i})$	Athlete fixed effects.
$\lambda(\mathbf{c}(\mathbf{e}))$	Event-category fixed effects.
$\varepsilon(\mathbf{i}, \mathbf{e}, \mathbf{t})$	Regression error term.
$\mathbf{Y}(\mathbf{i}, 1), \mathbf{Y}(\mathbf{i}, 0)$	Rubin potential outcomes under treatment/control.
$\mathbf{T}(\mathbf{i})$	Treatment indicator (1 = treated, 0 = control).
ATE	Average Treatment Effect: $\mathbb{E}[\mathbf{Y}(1) - \mathbf{Y}(0)]$.
$\mathbf{x}(\mathbf{i})$	Running variable (race-day index in RDD).
$\mathbf{h}$	Bandwidth for local polynomial estimation.
$\mathbf{K}(\cdot)$	Triangular kernel weighting function.
$\alpha^{\text{L}}(\mathbf{x}), \alpha^{\text{R}}(\mathbf{x})$	Conditional mean functions left/right of cutoff.
τ	Sharp treatment effect at cutoff.

I. Introduction

Drafting or also called *Slipstreaming*—minimizing aerodynamic or hydrodynamic resistance in water, which is about 800 to 830 times denser than air, deliberately through motion—is a defining feature of peloton sports (Matthes and Piazzolo 2024) and open-water swimming (Chatard et al. 1990; Millet, Chatard, and Chollet 1999; Reichel 2025d). In triathlon, swim-leg drafting and swim group formation shape the strategic environment for the subsequent bike and run segments/stages: athletes who conserve energy or attach to faster swim groups exit the water with effectively lower marginal costs (Reichel 2025b; Olbrecht 2011). These cost advantages then propagate to the later stages; consequently, relative rank improvements (or declines in sometimes—but usually not in majority) can be observed and are therefore used in this paper as a free-riding reverse proxy variable.

Additionally, the real and/or real-time formation of swim groups is always unobserved and may change multiple times during the race. However, my main argument for using the rank-improvement proxy is that overtaking during the first/swimming stage is relatively rare—and even more so in shorter race formats, for which most cross sections are observed. Thus, the swim exit time for group formation remains a valid approximation for the purposes of this paper.

This paper therefore aims to study the economics and determinants of free-riding in mass-start triathlon. The theoretical framework builds upon non-discriminatory Tullock contests (Tullock 1980), modeling triathlon as a two-stage game with quadratic first-stage costs, and extends to n -player and asymmetric cases in which drafting reduces heterogeneous marginal effort costs. Empirically, I construct an event-study measure of free-riding based on rank improvements, centered around the exceptional 2020 season, when COVID-19 response policies (individually staggered starts, distancing, and reduced congestion) sharply reduced drafting opportunities (Reichel 2025d). Identification combines pooled OLS with athlete and event fixed effects and a regression discontinuity design (RDD) that exploits this quasi-experimental policy setting, which induced local rule changes through individually staggered swim starts. The economic analysis then uses elasticity, frontier, and conditional quantile methods to examine primarily group-size effects.

The mechanism underlying the theoretical framework and model is well captured by contest theory: when drafting lowers effective costs, equilibrium efforts and winning probabilities shift monotonically with the cost shifter (Tullock 1980; Skaperdas 1996; Konrad 2009). Hydrodynamics provide the physical channel, as trailing swimmers face substantially less drag and oxygen uptake than leaders (Chatard et al. 1990). Prior empirical studies (Reichel 2025d) confirm that modified race formats mechanically reduced swim group formation and drafting opportunities, validating the regression discontinuity design (RDD) used in the present study. Complementary empirical studies extending this paper model triathlon as a finite multi-stage game with endogenous effort and drafting position, investigating nonlinear and diminishing returns to drafting depth (Reichel 2025b).

The remainder of the paper is organized as follows. Section II reviews the related literature, connecting contest theory, soft-matter physics, applied physiology and medical literature (Ohmichi, Takamoto, and Miyashita 1983), quasi-experimental policy variation, and treatment effects under heterogeneity analysis. Section III presents the core theoretical setup, beginning with the two-player case and extending to general n -player and asymmetric contests. Section III also formalizes the contest environment, showing how drafting operates as a cost shifter within Tullock-style games. Section IV presents the data, variable construction, and implementation procedures, while Section V provides descriptive statistics for free-riding outcomes. Section VI outlines the identification strategy, visualizing evidence from the 2020 disruption and the regression discontinuity design (RDD). Section VIII details the empirical framework, including the baseline pooled OLS (POLS) model with fixed effects (FE) and pooled-period specifications. Section XII develops an economic interpretation of free-riding,

drawing on elasticity, free-riding production frontiers, Engel and Beveridge camel-hump curve analogies from microeconomics and labor economics, and growth-incidence curves (GIC) as used in macroeconomics and the economics of inequality (e.g., pro-poor growth).

Section IX turns to heterogeneity across gender, age, drafter position, swimmer ability, gender group composition, and group-size effects. Section XIII discusses broader implications, and Section XIV concludes.

II. Related Literature

This paper connects to four different fields of literature—(1) contest theory and rent seeking; (2) drafting physics and swim group dynamics; (3) COVID–19 policy as quasi–experimental variation; and (4) team composition, heterogeneity, and distributional methods—and maps them implicitly to the theoretical and empirical interpretations documented across Sections III–IX. The discussion then highlights findings on the 2020 collapse and partial rebound of free–riding, group–size elasticities, distributional production frontiers, and heterogeneous effects by ability, drafter position, gender, and age.

Contests and Rent Seeking

We simply model the bike–run stage as a nondiscriminatory Tullock–style contest with drafting as a multiplicative cost shifter. Foundations on contest success functions (CSF), existence, uniqueness, and comparative statics generally follow (Tullock 1980; Skaperdas 1996; Szidarovszky and Okuguchi 1997), with comprehensive treatments and surveys in (Konrad 2009; Corchón Diaz and Serena 2016). Classical rent–seeking and all–pay results provide the welfare and design background for effort elicitation and prize allocation (Krueger 1974; Hillman and Riley 1989; Baye, Kovenock, and Vries 1996; Moldovanu and Sela 2001; Che and Gale 2003).

Recent work on multi–contest environments and risk–taking, and sex–based differences in risk–taking—if drafting is interpreted as decreased risk exposure—clarifies behavioral responses under intensified competition (Matthes and Piazzolo 2024; Gürtler, Struth, and Thon 2023; Song and Houser 2021; Lackner and Sonnabend 2025), while multiplicity in contests informs the interpretation of nonmonotonic patterns (Modak Chowdhury and Sheremeta 2011). In our framework, drafting–driven cost relief shifts equilibrium efforts and success probabilities monotonically, consistent with the observed collapse of relative multi–stage rank improvements—used here as a free–riding proxy—in 2020 under format changes, and the partial post–2020 rebound (see Section VI and Section IX).

Drafting Physics and Swim Group Dynamics

Open–water studies document sizable drag and metabolic reductions for trailing swimmers at short distances (Chatard et al. 1990; Chatard and Wilson 2003), with technique and kicking interacting with drafting among elite triathletes (Millet, Chatard, and Chollet 1999). Equipment and environmental conditions further influence resistance (Toussaint, Bruinink, Coster, et al. 1989; Ohmichi, Takamoto, and Miyashita 1983), and scale–model experiments document substantial drag shielding in formations (Bolon et al. 2020). Coaching syntheses and sport–science reviews emphasize how early positioning propagates to later race segments (Olbrecht 2011) and how social and pacing environments shape overall race performance (Sakalidis et al. 2022).

In our data, these mechanisms manifest as stronger free–riding gains for deeper drafter positions and weaker swimmers, as well as increasing returns with swim group size—patterns we quantify using year $\times$ group event studies and group–size gradients (see Section VIII and Section IX).

COVID–19 Policy as Quasi–Experimental Variation

Pandemic response policies (rolling or wave starts, enforced spacing, and reduced congestion) plausibly curtailed drafting opportunities (Leitner 2021; Österreichischer Triathlonverband (ÖTRV) 2020). We leverage this exogenous shift through an event study centered on the

exceptional year 2020 and a sharp RDD at the 2019/2020 cutoff, consistent with the causal inference framework of (Angrist and Pischke 2009; Imbens and Rubin 2015; Pearl 2009; Rubin 1974) and panel fixed-effects practice in (Wooldridge 2010).

The results show a significant downturn in free-riding in 2020 with partial recovery thereafter, and a significant discontinuity at the cutoff in local polynomial specifications—consistent with mechanically constrained swim group formation (Reichel 2025d). Semi-log specifications and transformation issues are addressed following (Huntington-Klein 2023; Santos Silva and Tenreiro 2006; Bellemare and Wichman 2020). Detailed evidence is provided in Section VI.

Team Composition, Heterogeneity, and Strategic Complementarities

The peer-effects and participation literatures motivate our design choices and identification strategies for social interactions (Manski 1993; Bramoullé, Djebbari, and Fortin 2009; Angrist 2014; Guryan, Kroft, and Notowidigdo 2009; De Giorgi, Frederiksen, and Pistaferri 2020; Nishihata 2022). In strategic environments with complementarities, composition and dispersion shape equilibrium responses (Milgrom and Shannon 1994; Vives 2005; Tadelis 2013). Empirically, we find stronger effects for weaker swimmers, increasing benefits with deeper drafter positions, modest gender gaps, and pronounced age heterogeneity peaking among middle-aged cohorts—each consistent with drafting shifting heterogeneous marginal costs and continuation values. These patterns are reported in year $\times$ dimension event studies and fixed-effects regressions with athlete and event-category controls (see Section IX).

Measurement, Distributional, and Macro-Micro Analogies

The economic analysis and interpretation employ elasticities, frontiers, and distributional methods. The estimated elasticity of free-riding with respect to group size (log-log slope $\hat{\beta} \approx 0.427$) maps percentage changes in swim group size to free-riding and rank-improvement responses, following guidance on log and inverse hyperbolic sine interpretations (Bellemare and Wichman 2020; Santos Silva and Tenreiro 2006; Huntington-Klein 2023). We approximate production frontiers via conditional quantile regressions (Koenker 2005; S. Firpo, Fortin, and Lemieux 2009; S. P. Firpo, Fortin, and Lemieux 2018; Chernozhukov, Fernández-Val, and Melly 2013), and investigate distributional dominance (Davidson and Duclos 2000).

To further investigate outcome dispersion, we draw on information-theoretic measures (Theil 1967; Cowell 2005). Engel- and Beveridge-curve analogies link our group-size and slack relationships to classic search and matching frameworks (Pissarides 2000; Petrongolo and Pissarides 2001; Blanchard and Diamond 1989; Lubik 2013; Barnichon and Figura 2010; Acemoglu 2020) and to consumption responses and aggregation (Deaton 1992). These literatures support the empirical chapters that estimate group-size elasticities, quantile frontiers, Beveridge-style interior optima, and growth-incidence curves (GIC; see Section XII and Section IX).

Quick Links to Results. Consistent with the above literatures, this paper documents: (a) a substantial pre-2020 level of free-riding, a collapse in 2020, and partial rebound thereafter (Section VI); (b) significant discontinuities at the 2019/2020 cutoff (RDD; Section VI); (c) positive and economically meaningful group-size elasticities (Section XII); (d) larger gains for weaker swimmers and deeper drafter positions; (e) modest gender differences; and (f) pronounced age heterogeneity (Section IX).

III. Theoretical Framework

This section develops a contest-theoretic framework as already published (Reichel 2025c) in which swim drafting operates as a multiplicative cost shifter in subsequent competitive stages. The framework is tailored to mass-start triathlon but applies more broadly to other multi-stage competitions in which early positioning or result realizations affect later effort costs and participation incentives.

The model builds on the canonical Tullock contest (Tullock 1980; Skaperdas 1996; Konrad 2009) and embeds it in a sequential finite-horizon environment with endogenous continuation decisions. Athletes first observe post-swim outcomes, including drafting exposure, and decide whether to continue racing. Those who continue enter a bike-run contest in which equilibrium effort, winning probabilities, and welfare depend on drafting-induced cost heterogeneity.

The key mechanism is that drafting in the swim reduces effective marginal effort costs later on. This cost relief affects both (i) optimal effort choices in the contest stage and (ii) the incentive to continue competing after the swim. The framework therefore links swim group formation directly to free-riding behavior, continuation, and aggregate effort.

III.1 Players, Timing, and Information

There are $n \geq 2$ athletes indexed by $i \in N = \{1, \dots, n\}$. Competition unfolds over two stages.

Stage 1: Post-swim continuation. At the end of the swim segment, each athlete i observes a publicly known state vector

$$s_i = (t_i^S, r_i^S, D_i, G),$$

where t_i^S is swim time, r_i^S is swim rank, $D_i \in [0, 1]$ is the share of swim time spent drafting, and G denotes the directed drafting network formed during the swim. These objects summarize early performance and the extent of hydrodynamic cost reduction.

Each athlete simultaneously chooses whether to *continue* (C) or *terminate* (T) the race. Let $S \subseteq N$ denote the set of athletes who continue, with group size $m = |S| \geq 1$.

Stage 2: Bike-run contest. If $m \geq 1$, athletes in S compete in a Tullock contest by choosing nonnegative efforts $e_i \geq 0$ simultaneously. Effort determines winning probabilities and is costly. Drafting realized in Stage 1 affects these costs multiplicatively.

All information is complete at both stages, and the solution concept is subgame-perfect equilibrium (SPE) which is generally solvable via backward induction.

III.2 Preferences and Outside Options

If athlete i terminates in Stage 1, her payoff equals

$$u_i^T = -\alpha t_i^S - \beta r_i^S + \vartheta_i, \quad (\text{III.1})$$

where $\alpha, \beta > 0$ measure disutility from swim time and rank, and ϑ_i captures idiosyncratic benefits of terminating (e.g., training value, injury avoidance).

If athlete i continues to Stage 2, expected utility is determined by contest outcomes and effort costs. These are described next.

III.3 Stage 2: Tullock Contest with Drafting as a Cost Shifter

To solve for the SPE via backward induction we start formally with Stage 2. Conditional on a nonempty continuation set $S \subseteq N$, athletes in S compete in a bike-run contest. Effort provision is simultaneous and costly, and winning probabilities follow a Tullock contest success function (CSF).

Contest technology. Let $w_i > 0$ denote a contest weight capturing effectiveness unrelated to effort costs (e.g., technical ability or discipline specialization). The probability that athlete $i \in S$ wins is simply give as

$$p_i(e) = \frac{w_i e_i}{\sum_{j \in S} w_j e_j}, \quad e_i \geq 0. \quad (\text{III.2})$$

Effort costs and drafting. Effort costs are quadratic with effective cost coefficient

$$C_i(e_i) = \frac{1}{2} k_i e_i^2, \quad k_i := \frac{c_i}{\psi_i}, \quad (\text{III.3})$$

where $c_i > 0$ is a baseline cost parameter and $\psi_i > 0$ is a swim drafting multiplier. Higher drafting exposure in Stage 1 increases ψ_i and lowers marginal effort costs.

Prizes and expected utility. Let $\Delta_i > 0$ denote athlete i 's prize differential (winner's prize minus loser's prize). Dropping constants independent of effort, expected utility in Stage 2 is

$$U_i(e) = p_i(e) \Delta_i - \frac{1}{2} k_i e_i^2. \quad (\text{III.4})$$

The equilibrium analysis proceeds by characterizing best responses, existence, and comparative statics.

III.3.1 Best responses

Let $E = \sum_{j \in S} e_j$ denote total effort. Since $p_i = e_i/E$, differentiation yields

$$\frac{\partial p_i}{\partial e_i} = \frac{p_i(1-p_i)}{e_i}, \quad \frac{\partial p_i}{\partial e_j} = -\frac{p_i p_j}{e_j}, \quad j \neq i. \quad (\text{III.5})$$

Athlete i solves

$$\max_{e_i \geq 0} \Delta_i p_i(e) - \frac{1}{2} k_i e_i^2.$$

For any interior solution $e_i > 0$, the first-order condition is

$$\Delta_i \frac{p_i(1-p_i)}{e_i} - k_i e_i = 0 \iff e_i^2 = \frac{\Delta_i}{k_i} p_i(1-p_i). \quad (\text{III.6})$$

Equation (III.6) links equilibrium effort to contest probabilities and drafting-adjusted costs.

III.3.2 Two-player closed form

Suppose $S = \{i, j\}$. Define the effective advantage ratio

$$R := \frac{\Delta_i \psi_i / c_i}{\Delta_j \psi_j / c_j} = \frac{\Delta_i / k_i}{\Delta_j / k_j}. \quad (\text{III.7})$$

Combining (III.6) with $p_i/(1-p_i) = (w_i e_i)/(w_j e_j)$ yields

$$\frac{p_i}{1-p_i} = R^{1/2} \frac{w_i}{w_j}. \quad (\text{III.8})$$

Thus $p_i^* = \rho/(1+\rho)$, where $\rho = R^{1/2}(w_i/w_j)$. Substituting into (III.6) gives equilibrium efforts:

$$\begin{aligned} e_i^* &= \sqrt{\frac{\Delta_i \psi_i}{c_i}} \frac{\rho^{1/2}}{1+\rho}, \\ e_j^* &= \sqrt{\frac{\Delta_j \psi_j}{c_j}} \frac{\rho^{1/2}}{1+\rho}. \end{aligned} \quad (\text{III.9})$$

Effort is increasing in the prize Δ_i and drafting multiplier ψ_i , and decreasing in baseline cost c_i .

III.3.3 Symmetric m-player benchmark

Assume $w_i \equiv 1$, $c_i \equiv c$, $\psi_i \equiv \psi$, and $\Delta_i \equiv \Delta$ for all $i \in S$, with $m = |S|$. Symmetry implies $p_i^* = 1/m$. Equation (III.6) yields

$$(e^*)^2 = \frac{\Delta \psi}{c} \frac{m-1}{m^2}, \quad e^* = \sqrt{\frac{\Delta \psi}{c}} \sqrt{\frac{m-1}{m^2}}. \quad (\text{III.10})$$

Equilibrium effort decreases in group size m due to diluted winning probabilities, but increases in drafting intensity ψ .

III.3.4 Asymmetric m-player equilibrium

For general heterogeneity, summing (III.6) across $i \in S$ implies

$$e_i = \frac{\Delta_i E}{c_i E^2 / \psi_i + \Delta_i}, \quad E = \sum_{j \in S} e_j. \quad (\text{III.11})$$

Total effort E satisfies the fixed-point equation

$$g(E) := \sum_{i \in S} \frac{\Delta_i}{c_i E^2 / \psi_i + \Delta_i} - 1 = 0. \quad (\text{III.12})$$

Proposition III.1 (Existence and uniqueness). *Assume $c_i > 0$, $\psi_i \in [\underline{\psi}, \bar{\psi}] \subset (0, \infty)$, and $\Delta_i > 0$. Then $g(E)$ is continuous and strictly decreasing on $E \geq 0$, with $g(0) = m-1 > 0$ and $\lim_{E \rightarrow \infty} g(E) = -1$. Hence there exists a unique $E^* > 0$ solving (III.12). The implied effort profile (III.11) is the unique interior Nash equilibrium of Stage 2.*

Proof. Each term $\Delta_i/(c_i E^2 / \psi_i + \Delta_i)$ is strictly decreasing in E , equal to one at $E = 0$, and converges to zero as $E \rightarrow \infty$. Strict concavity of U_i in e_i ensures uniqueness. $\square$

Corollary III.1 (Equilibrium probabilities and efforts). *At the unique equilibrium,*

$$p_i^* = \frac{\Delta_i}{c_i (E^*)^2 / \psi_i + \Delta_i}, \quad e_i^* = p_i^* E^*.$$

Proposition III.2 (Comparative statics). *Total effort E^* and individual efforts e_i^* increase in ψ_i and Δ_i , and decrease in c_i . For any $j \neq i$, p_j^* weakly decreases when ψ_i or Δ_i increases.*

III.4 Stage 1: Continuation Decisions and Subgame-Perfect Equilibrium

After observing post-swim outcomes, athletes decide whether to continue competing or to terminate the race. This section characterizes continuation values, cutoff rules, and the existence of subgame-perfect equilibria.

III.4.1 Continuation values

Fix a post-swim state vector (t_i^S, r_i^S, D_i, G) for each athlete i , and suppose a continuation set $S \subseteq N$ with $i \in S$ proceeds to Stage 2. Let $E^*(S)$ and $\{p_i^*(S), e_i^*(S)\}_{i \in S}$ denote the unique equilibrium characterized in Proposition III.1.

Athlete i 's continuation value is

$$\begin{aligned} W_i(S) &= p_i^*(S) \Delta_i - \frac{1}{2} k_i (e_i^*(S))^2 \\ &= \frac{\Delta_i^2}{c_i (E^*(S))^2 / \psi_i + \Delta_i} - \frac{1}{2} \frac{c_i}{\psi_i} (p_i^*(S) E^*(S))^2. \end{aligned} \quad (\text{III.13})$$

The net gain from continuing relative to termination is given as

$$V_i(S) = W_i(S) - U_i^T, \quad (\text{III.14})$$

where U_i^T is defined in (III.1).

III.4.2 Cutoff rules

Continuation incentives depend monotonically on drafting exposure through ψ_i .

Proposition III.3 (Cutoff in drafting-induced cost reduction). *Fix (t_i^S, r_i^S) , (c_i, Δ_i) , and a continuation set S with $i \in S$. The net continuation value $V_i(S)$ is continuous and strictly increasing in the swim drafting multiplier ψ_i . Hence, there exists a unique cutoff $\psi_i^*(S)$ such that*

$$V_i(S) \geq 0 \iff \psi_i \geq \psi_i^*(S).$$

Proof. By Proposition III.2, both $E^*(S)$ and $p_i^*(S)$ increase in ψ_i , while the effective cost parameter $k_i = c_i / \psi_i$ decreases. Therefore $W_i(S)$ in (III.13) is strictly increasing in ψ_i . Since U_i^T does not depend on ψ_i , the result follows by continuity and the intermediate value theorem (IVT). $\square$

III.4.3 Stage 1 equilibrium and SPE existence

At Stage 1, athletes simultaneously choose $a_i \in \{C, T\}$. Let $\mathcal{S}$ denote the family of nonempty subsets of N .

Definition III.1 (Equilibrium continuation set). A set $S^* \in \mathcal{S}$ is a Stage 1 equilibrium continuation set if, for all $i \in N$,

$$i \in S^* \Rightarrow V_i(S^*) \geq 0, \quad i \notin S^* \Rightarrow V_i(S^* \cup \{i\}) \leq 0. \quad (\text{III.15})$$

Theorem III.1 (Existence of a subgame-perfect equilibrium (SPE)). *Under the maintained assumptions, there exists at least one subgame-perfect equilibrium. Any ψ -monotone selection rule that iteratively admits athletes with $V_i(S) \geq 0$ and removes those with $V_i(S) < 0$ converges to an equilibrium continuation set S^* .*

Proof. Define the operator $\Phi(S) = \{i \in N : V_i(S) \geq 0\}$. Since N is finite, iterative application of Φ converges in finitely many steps to a fixed point S^* satisfying $\Phi(S^*) = S^*$, which fulfills (III.15). Backward induction then delivers a subgame-perfect equilibrium (SPE). $\square$

Remark III.1 (Comparative statics of continuation). The cutoff $\psi_i^*(S)$ decreases in the prize differential Δ_i and increases in the baseline cost c_i . In symmetric environments, $\psi_i^*(S)$ weakly increases with group size $|S| > 1$, since larger contests dilute winning probabilities.

III.5 Drafting as a Cost-Shifting Technology

I now formalize how swim drafting propagates into a cost shifter in the contest stage.

Let $D_i \in [0, 1]$ denote athlete i 's draft exposure share, defined as the fraction of swim time spent in a hydrodynamic wake. Draft exposure depends on swim rank, group size, and network position.

I simply model drafting as a multiplicative cost shifter,

$$\psi_i = \Psi(D_i, m, r_i^S, G) \in [\underline{\psi}, \bar{\psi}] \subset (0, \infty), \quad (\text{III.16})$$

with Ψ weakly increasing in D_i and Lipschitz continuous in (D_i, m) .

Reduced-drag specification. A "удобный" fact is

$$\psi_i = \frac{1}{1 - \eta D_i}, \quad \eta \in (0, 1), \quad (\text{III.17})$$

which implies

$$k_i = \frac{c_i}{\psi_i} = c_i(1 - \eta D_i), \quad \frac{\partial k_i}{\partial D_i} = -\eta c_i < 0.$$

Thus drafting linearly reduces marginal effort costs in the bike-run stage.

III.6 Welfare and Rent Dissipation

Aggregate welfare in Stage 2 is simply the sum of equilibrium utilities:

$$\sum_{i \in S} u_i(e^*) = \sum_{i \in S} \left[\frac{\Delta_i^2}{c_i(E^*)^2 / \psi_i + \Delta_i} - \frac{1}{2} \frac{c_i}{\psi_i} (p_i^* E^*)^2 \right]. \quad (\text{III.18})$$

Symmetric benchmark. Under symmetry (Δ, c, ψ) with group size m ,

$$\frac{\text{total effort cost}}{\text{expected prize intake}} = \frac{1}{2} \cdot \frac{m-1}{m}.$$

Rent dissipation increases with group size and converges to $1/2$ as $m \rightarrow \infty$.

III.7 Testable Predictions


The model yields the following predictions:

- (P1) Higher draft exposure D_i increases ψ_i , raising equilibrium effort e_i^* and winning probability p_i^* .
- (P2) Larger group size m lowers individual effort through diluted winning odds but may increase aggregate effort if drafting access rises with m .
- (P3) Continuation is more likely for athletes with higher draft exposure and lower baseline costs—thus leaders might drop out earlier.
- (P4) Free-riding gains are strongest for weaker swimmers and deeper drafting positions.

III.8 Figures


Extensive Form

Figure 1: Extensive Form of the Two-Stage Contest



Comparative Statics

Figure 2: Equilibrium Effort as a Function of Drafting



IV. Panel Data, Sample Restriction, and Variable Construction

The empirical analysis relies on an assembled *event–athlete panel* that links individual performances to race characteristics over a ten-year horizon. The raw source is the preprocessed file, from which we construct a structured panel indexed by athlete $\mathbf{a}$ and event $\mathbf{e}$. Each record contains identifiers, performance outcomes, and demographics, forming the basis for all measures of drafting and free-riding.

Each observation $(\mathbf{a}, \mathbf{e})$ provides: a unique athlete identifier $\mathbf{a}$, a unique event identifier $\mathbf{e}$, the athlete’s overall finishing rank $r_{\mathbf{a},\mathbf{e}}^{\text{tot}}$, the swim split $s_{\mathbf{a},\mathbf{e}}$ (seconds), a sex indicator $m_{\mathbf{a},\mathbf{e}} \in \{0, 1\}$ and the event date $d_{\mathbf{e}}$, shared by all participants in event $\mathbf{e}$.

From the event date we extract the calendar year,

$$y_{\mathbf{e}} = \text{year}(d_{\mathbf{e}}), \quad \mathcal{Y} = \{2015, \dots, 2024\},$$

and define the subset of athletes who competed at least once during the pandemic year:

$$\mathcal{A}_{2020} = \{ \mathbf{a} : \exists \mathbf{e} \text{ with } y_{\mathbf{e}} = 2020 \text{ and observation } (\mathbf{a}, \mathbf{e}) \}.$$

Restricting the panel to $\mathcal{A}_{2020}$ centers the analysis on the COVID–19 season, which provides the quasi-experimental variation exploited in the causal identification strategy chapter.

IV.1 Core Variables

Table 1 summarizes the main variables. In addition to raw identifiers and split times, the panel includes within-event ranks, drafting positions, and group size measures that connect contest theory to observed race behavior.

Name	Type	Unit	Definition / Notes
$\mathbf{a}$	ID	–	Athlete identifier (stable across events).
$\mathbf{e}$	ID	–	Event identifier (unique race instance).
$d_{\mathbf{e}}$	date	ISO	Event date.
$y_{\mathbf{e}}$	integer	year	Calendar year extracted from $d_{\mathbf{e}}$.
$m_{\mathbf{a},\mathbf{e}}$	binary	$\{0, 1\}$	Sex indicator (1 = male, 0 = female).
$s_{\mathbf{a},\mathbf{e}}$	time	s	Swim split time (seconds).
$r_{\mathbf{a},\mathbf{e}}^{\text{swim}}$	rank	integer	Within-event swim rank (ascending; average ties).
$r_{\mathbf{a},\mathbf{e}}^{\text{tot}}$	rank	integer	Overall finishing rank.
$\text{leader}_{\mathbf{a},\mathbf{e}}$	binary	$\{0, 1\}$	Indicator for swim leader (1 if led the swim pack).
$\text{drafter_position}_{\mathbf{a},\mathbf{e}}$	integer	places	Within-pack drafting position (1 = pack leader, which is also dropped).
$g_{\mathbf{e}}$	count	athletes	Group size (number of unique athletes in event $\mathbf{e}$).

Table 1: Core variables in the event–athlete panel.

IV.2 Sample Construction

The initial sample is

$$\mathcal{S}_{\text{raw}} = \{(\mathbf{a}, \mathbf{e}) : y_{\mathbf{e}} \in \mathcal{Y}\}.$$

I restrict causal attention to athletes who at least competed once in 2020:

$$\mathcal{S} = \{(\mathbf{a}, \mathbf{e}) \in \mathcal{S}_{\text{raw}} : \mathbf{a} \in \mathcal{A}_{2020}\}.$$

Three additional data quality filters refine the analytic sample:

$$(i) \ r_{a,e}^{\text{tot}} \leq 500, \quad (\text{IV.1})$$

$$(ii) \ \text{leader}_{a,e} = 0, \quad (\text{IV.2})$$

$$(iii) \ 1 < \text{drafter_position}_{a,e} \leq 20. \quad (\text{IV.3})$$

These conditions limit the sample to competitive finishers, exclude swim leaders (for whom drafting is irrelevant), and focus on realistic drafting positions. The final analytic panel is

$$\tilde{\mathcal{S}} = \{(a, e) \in \mathcal{S} : (\text{IV.1}) \wedge (\text{IV.2}) \wedge (\text{IV.3})\}.$$

I define three time periods:

$$\text{Pre-COVID: } y_e \in \{2015, \dots, 2019\},$$

$$\text{COVID: } y_e = 2020,$$

$$\text{Post-COVID: } y_e \in \{2021, \dots, 2024\}.$$

IV.3 Derived Measures

Several derived measures operationalize drafting and free-riding behavior.

Swim rank and rank improvement. Within each event e ,

$$r_{a,e}^{\text{swim}} = \text{rank}_e(s_{a,e}),$$

where ties are averaged downwards (to reduce this ranking bias). Rank improvement is defined as

$$\Delta r_{a,e} = r_{a,e}^{\text{tot}} - r_{a,e}^{\text{swim}},$$

so $\Delta r_{a,e} < 0$ indicates gains relative to the swim leg/stage.

Free-riding outcome. To normalize outcomes relative to the COVID baseline, define

$$\text{FR}_{a,e} = \Delta r_{a,e} - \overline{\Delta r}_{2020}, \quad \overline{\Delta r}_{2020} = \frac{1}{|\{(a, e) : y_e = 2020\}|} \sum_{\{(a, e) : y_e = 2020\}} \Delta r_{a,e}.$$

Hence, $\text{FR}_{a,e}$ measures athlete a 's deviation from the 2020 mean rank improvement.

Binning and ability classes. Group size g_e is categorized as small (< 5), medium (5–10), or large (> 10). Age is grouped into five-year bands. Swim ability is proxied by k-means clustering ($k = 3$) of standardized best swim times by category. Clusters are ordered by mean swim speed and labeled *good*, *average*, and *poor*.

IV.4 Sample Restriction Procedure

Figure 3 summarizes the procedure:

- (a) Load and retain all observations from 2015–2024.
- (b) Identify $\mathcal{A}_{2020}$ and construct a balanced panel across years.
- (c) Apply the quality filters (IV.1)–(IV.3).
- (d) Compute within-event swim ranks and rank improvements.
- (e) Create period indicators, group-size bins, and age classes.
- (f) Compute free-riding outcomes $\text{FR}_{a,e}$.
- (g) Derive swim-ability clusters via k-means with 3 classes.

IV.5 Implementation Notes

Practical implementation follows these conventions:

- (i) *Ranking ties*: average ranks are used; dense or minimum ranks yield nearly identical results.
- (ii) *Leader exclusion*: athletes identified as swim leaders are removed from drafting analyses.
- (iii) *Drafting positions*: only positions 2 through 20 are retained.
- (iv) *Outlier handling*: extreme swim or total times are winsorized at the 0.5 and 99.5 percentiles.
- (v) *Reproducibility*: all clustering employs standardized variables, a fixed random seed, and k-means++ initialization.




Figure 3: Flow of sample construction and derived measures in the event-athlete panel. The process includes data import, identification of $\mathcal{A}_{2020}$, application of filters (IV.1)–(IV.3), computation of $r_{a,e}^{\text{swim}}$ and $\Delta r_{a,e}$, periodization, definition of $\text{FR}_{a,e}$, and creation of k -means swim-ability clusters, resulting in the final dataset $\tilde{\mathcal{S}}$.

V. Descriptive Statistics

This section documents the distribution of key covariates and outcomes *before* and *after* the data processing pipeline described in Section IV. We begin with the original dataset sample, conduct basic validity checks, and then present the post-pipeline analysis sample that implements the “same-athletes” restriction and data quality filters (IV.1)–(IV.3). We conclude with descriptive evidence on the evolution of the outcome over time and across subgroups.

V.1 Overview and Unit of Analysis

The unit of observation is athlete–event ($\mathbf{a}, \mathbf{e}$). For each record we observe identifiers, timing splits, ranks, demographic indicators, and variables that proxy drafting opportunity (e.g., drafter position, swim group/cluster, field size). Periodization follows Section IV: *pre* (2015–2019), *COVID* (2020), and *post* (2021–2024).

V.2 Covariates: Original Dataset (Pre-Pipeline)

Table 2 summarizes the main variables in the original dataset prior to any pipeline restrictions. These moments characterize the raw distribution of times and ranks, the mix of male/female athletes, and the breadth of observed group structures (leaders, drafters, positions). Note that this pre-pipeline sample still includes swim leaders and very deep drafter positions.

Table 2: Descriptive statistics — original dataset sample

Variable	Min	1st Quartile	Median	Mean	3rd Quartile	Max
Total	165	4965	8232	11789	16809	61282
Swim-Out times (sec)	165	4825	1431	1625	2097	29700
Race rank	0	286	91	156.9	151	2623
Male	0.0000	1.0000	1.0000	0.7898	1.0000	1.0000
Year	1922	1969	1977	1977	1985	2009
COVID	0.0000	0.0000	0.0000	0.1097	0.0000	1.0000
Post COVID	0.0000	0.0000	0.0000	0.08563	0.0000	1.0000
Week running	0.0	159.0	279.0	317.64	475.0	751.0
Age (years)	8.00	31.00	38.00	38.56	46.00	99.00
Age ²	64	961	1444	1600	2116	9801
Event year	2010	2013	2015	2016	2019	2024
Cluster (swim group)	1.00	4.00	10.00	16.95	23.00	224.00
Leader	0.0000	0.0000	0.0000	0.3134	1.0000	1.0000
Drafter	0.0000	0.0000	1.0000	0.6866	1.0000	1.0000
Drafter position	1.00	1.00	3.00	17.15	11.00	937.00
First drafter	0.0000	0.0000	0.0000	0.2961	1.0000	1.0000
Second drafter	0.0000	0.0000	0.0000	0.13420	0.0000	1.0000
Third drafter	0.0000	0.0000	0.0000	0.08329	0.0000	1.0000
Fourth drafter	0.0000	0.0000	0.0000	0.05846	0.0000	1.0000
Fifth drafter	0.0000	0.0000	0.0000	0.0440	0.0000	1.0000
Last drafter	0.0000	0.0000	0.0000	0.2961	1.0000	1.0000

Validity checks. We confirm that variables lie in plausible ranges and that coverage is balanced across periods and categories. Figure 4 compares swim-out times by event category and period; spreads are broadly comparable, with the anticipated shifts around 2020. Extreme

swim times are rare and retained in the raw view; robustness checks (elsewhere) trim at $1.5 \times \text{IQR}$ as needed.




Figure 4: Swim-out times by event category and period (original sample). Boxes show the interquartile range; whiskers extend to $1.5 \times \text{IQR}$.

V.3 Covariates: Analysis Dataset (Post-Pipeline)

We now turn to the analysis dataset produced by the pipeline in Section IV. The “same-athletes” restriction centers the panel on 2020 (the quasi-experimental pivot), and quality filters (IV.1)–(IV.3) remove leaders and implausible drafting positions. As a result, all retained observations are drafters in positions 2–20, and leaders are absent by construction.

To complement the continuous summaries, Table 4 reports level counts for categorical covariates in the post-pipeline sample. These distributions show the modal event categories, the composition by gender and swim skill, and the mass of observations across drafter positions and group-size bins.

V.4 Outcomes: Free-Riding

Definition and Construction

We measure *free-riding potential* by centering athlete-event rank improvements relative to the 2020 benchmark year. For athlete i in event t :

$$\text{FR}_{it} = \text{RankImprovement}_{it} - \overline{\text{RankImprovement}}_{\{\text{year}=2020\}},$$

where the bar denotes the sample mean among all 2020 observations. This normalization places 2020 at zero by construction and highlights deviations before and after the COVID-19 disruption. The data sample for all figures below mirrors the pipeline in Section IV (same-athletes, filters (IV.1)–(IV.3), consistent outlier handling).

Table 3: Summary statistics after pipeline (n = 15,028)

Variable	Min	1st Quartile	Median	Mean	3rd Quartile	Max
X (row index)	71530	108936	132896	126833	145434	168391
Total (sec)	201	4674	7636	9716	10614	51755
Swimming (sec)	116.0	815.8	1272.5	1434.1	1942.0	6131.0
Rank	0.00	20.00	46.00	73.83	98.00	500.00
Male	0.0000	1.0000	1.0000	0.7678	1.0000	1.0000
Year	2015	2017	2019	2019	2021	2024
COVID	0.0000	0.0000	0.0000	0.3374	1.0000	1.0000
Post COVID	0.0000	0.0000	0.0000	0.1414	0.0000	1.0000
Week running	262.0	377.0	488.0	499.5	594.0	751.0
Age (years)	13.0	32.0	40.0	40.2	48.0	79.0
Age ²	169	1024	1600	1744	2304	6241
Event year	2015	2017	2019	2019	2021	2024
Cluster (swim group)	1.00	4.00	10.00	14.71	20.00	175.00
Leader	0	0	0	0	0	0
Drafter	1	1	1	1	1	1
Drafter position	2.000	3.000	4.000	5.953	8.000	20.000
First drafter	0	0	0	0	0	0
Second drafter	0.0000	0.0000	0.0000	0.2429	0.0000	1.0000
Third drafter	0.0000	0.0000	0.0000	0.1680	0.0000	1.0000
Fourth drafter	0.0000	0.0000	0.0000	0.1153	0.0000	1.0000
Fifth drafter (fift drafter)	0.00000	0.00000	0.00000	0.08491	0.00000	1.00000
Last drafter	0.0000	0.0000	0.0000	0.2484	0.0000	1.0000
Swim rank	1.00	7.00	17.00	24.67	33.00	216.00
Total rank	0.00	20.00	46.00	73.83	98.00	500.00
Rank improvement	-150.00	5.00	25.00	49.16	67.00	484.00
Free riding	-168.375	-13.375	6.625	30.783	48.625	465.625
Field size	1.00	24.00	39.00	48.45	64.00	216.00
log(field size)	0.000	3.178	3.664	3.588	4.159	5.375
Male share	0.0000	0.7222	0.7857	0.7678	0.8333	1.0000
Share good (swim skill)	0.0000	0.2500	0.3906	0.3956	0.5306	1.0000

Table 4: Categorical level counts after pipeline (n = 15,028)

Variable	Levels (count)
Event category	Kurz: 4554; Lang: 426; Mittel: 2858; Sprint: 7190
Gender (factor)	Female: 3490; Male: 11538
Swim skill	good: 5945; average: 6285; bad: 2798
Group size bin	small: 270; medium: 854; large: 13904
Drafter position (factor)	2: 3650; 3: 2524; 4: 1733; 5: 1276; 6: 960; 7: 782; other: 4103
Gender mix category	pure male: 569; male 80–20: 5608; balanced 50–50: 134; pure female: 142; other: 8575
Field tercile	T1 (small): 5020; T2 (medium): 5106; T3 (large): 4902

V.5 Temporal Patterns by Heterogeneity Dimension

We first describe mean free-riding potential over time by single dimensions—swim skill, gender, field size, and age. The figures report cell means without smoothing to make shifts around 2020.




Figure 5: Temporal evolution of mean free-riding potential by subgroup and interaction dimensions. Each panel shows cell means without smoothing to highlight shifts around 2020.

Field-size terciles summarize congestion as a primitive determinant of drafting scope, while

age bins proxy physiological and experience gradients that may interact with drafting benefits. Cross-dimensional panels test whether improvements post-2020 concentrated among weaker swimmers in larger fields or among specific age-skill combinations.

V.6 Heterogeneity Maps relative to 2019 (3D Contour Plots)

Finally, we present filled-contour summaries that highlight how changes in mean free-riding (relative to 2019) vary across six dimensions: gender, age, drafter position, swim ability, gender mix, and group size. Read these as smoothed heatmaps of cell means; darker regions indicate larger deviations.




Figure 6: Filled-contour descriptives: change in mean free-riding relative to 2019 across six dimensions.

Notes. All series are cell means without smoothing unless explicitly smoothed for the contour visualizations. Vertical reference lines at 2020 (where included) mark the centering year used to construct FR_{it} . The underlying selection, cleaning, and outlier handling replicate the main data pipeline in Section IV, ensuring comparability across panels and subgroups.

VI. Identification Strategy

The empirical strategy combines two complementary approaches. First, we provide descriptive evidence that free-riding potential was restricted during the 2020 COVID season using simple t -tests. Second, we formalize this intuition within a regression discontinuity design (RDD), exploiting the abrupt change in race formats induced by pandemic response policies.

VI.1 Evidence on Restricted Free-Riding Potential during COVID-19 (2020)

We begin with a direct comparison of mean rank improvements across periods. To assess whether the COVID-19 season (calendar year 2020) altered within-race dynamics, we employ Welch's two-sample t -test, which does not assume equal variances across groups. For each skill class, we test the null hypothesis

$$H_0 : \mu_A = \mu_B \quad \text{against} \quad H_1 : \mu_A \neq \mu_B,$$

where μ_A and μ_B denote the average rank improvement in two comparison periods.

The test statistic is

$$t = \frac{\bar{X}_A - \bar{X}_B}{\sqrt{\frac{s_A^2}{n_A} + \frac{s_B^2}{n_B}}},$$

with degrees of freedom approximated using the Satterthwaite formula:

$$v = \frac{\left(\frac{s_A^2}{n_A} + \frac{s_B^2}{n_B}\right)^2}{\frac{(s_A^2/n_A)^2}{n_A-1} + \frac{(s_B^2/n_B)^2}{n_B-1}}.$$

Figure 7 shows the empirical distributions of rank improvements by swim skill and period, while Table 5 reports test statistics, degrees of freedom, and p -values.

Three regularities emerge clearly:

- (i) *Pre-COVID vs. 2020*. Rank improvements are significantly lower in 2020 compared to the pre-COVID baseline, across all skill groups. This holds for good swimmers ($t = -2.20$, $p = 0.028$), average swimmers ($t = -5.25$, $p < 10^{-6}$), and bad swimmers ($t = -6.60$, $p < 10^{-10}$).
- (ii) *Post-COVID vs. 2020*. The post-COVID period exhibits even lower improvements relative to 2020, with $t = 10.08$ ($p < 10^{-22}$) for good swimmers, $t = 12.03$ ($p < 10^{-31}$) for average swimmers, and $t = 14.70$ ($p < 10^{-45}$) for bad swimmers.
- (iii) *Pre-COVID vs. Post-COVID*. Direct comparisons confirm a monotonic decrease in improvements over time, with strongly negative t -statistics: -10.81 ($p < 10^{-25}$) for good swimmers, -14.52 ($p < 10^{-45}$) for average swimmers, and -20.61 ($p < 10^{-90}$) for bad swimmers.

The magnitude of the t -statistics is largest for weaker swimmers. This is consistent with the interpretation that athletes starting further back benefit disproportionately when pack dynamics and course management reduce opportunities for others to free-ride. In short, rather than suppressing advancement, the constraints of 2020 coincide with—and are followed by—systematically larger rank improvements. **Conclusion: Pre $\succeq$ Post $\succeq$ 2020**




Figure 7: Distributions of rank improvements by swim skill and period (pre, 2020, post). **Conclusion:** $\text{Pre} \succeq \text{Post} \succeq 2020$

Table 5: Welch two-sample t-tests of mean rank improvement by period and swim skill

Skill	t-statistic	p-value		N _A	N _B
<i>Panel A. Pre-COVID vs 2020</i>					
good	−2.20	0.03 **		2638.00	430.00
average	−5.25	0.00 ***		3212.00	796.00
bad	−6.60	0.00 ***		1662.00	531.00
<i>Panel B. Post-COVID vs 2020</i>					
good	10.08	0.00 ***		2007.00	430.00
average	12.03	0.00 ***		2292.00	796.00
bad	14.70	0.00 ***		1139.00	531.00
<i>Panel C. Pre-COVID vs Post-COVID</i>					
good	−10.81	0.00 ***		2638.00	2007.00
average	−14.52	0.00 ***		3212.00	2292.00
bad	−20.61	0.00 ***		1662.00	1139.00

Notes: Welch unequal-variance two-sample t-tests; two-sided p-values. N_A and N_B denote the number of observations in the first and second period listed in each panel, respectively. A negative t implies that the second period has a higher mean rank improvement. Significance levels: * p<0.10, ** p<0.05, *** p<0.01. Extremely small p-values are rounded to four decimals. Degrees of freedom are calculated using the Satterthwaite approximation.

VI.2 Regression Discontinuity Design (RDD)

While the t-tests provide descriptive evidence, they do not by themselves establish causality. To address this, we turn to a regression discontinuity design (RDD) that leverages the discontinuous change in start formats introduced in 2020. The RDD framework sharpens identification by comparing athletes just on either side of the policy break, holding constant unobserved heterogeneity and long-run trends.

VI.3 Regression Discontinuity Design (RDD) on Rank-Improvement (2019 vs 2020)

Design. The outcome variable is rank improvement, defined as

$$Y_i = \text{TotalRank}_i - \text{SwimRank}_i,$$

for athlete i . The running variable is the available-day index x_i , constructed as the sequential race-day index after removing empty days. Observations from 2019 have $x_i < 0$ and observations from 2020 have $x_i > 0$.

We estimate the discontinuity at the cutoff $x = 0$ using a sharp regression discontinuity design (RDD). Let $\alpha_L(x)$ and $\alpha_R(x)$ denote the conditional mean functions on the left ($x < 0$) and right ($x > 0$) of the cutoff. The sharp treatment effect is

$$\hat{\tau} = \hat{\alpha}_R(0) - \hat{\alpha}_L(0).$$


Estimation. Following standard RDD practice, we use local polynomial regressions weighted by a simple default triangular kernel:

$$\hat{\alpha}_s(0) = \arg \min_{\alpha, \beta, \gamma} \sum_{i: s_i = s} K\left(\frac{x_i}{h}\right) [Y_i - \alpha - \beta x_i - \gamma x_i^2 - \dots]^2,$$

where $s \in \{L, R\}$ indexes the side of the cutoff, h is a symmetric bandwidth, and $K(u) = (1 - |u|)\mathbf{1}\{|u| < 1\}$ is the default triangular kernel.

We report both local linear (degree 1, meaning a slope) and local cubic (degree 3) specifications. Inference relies on Huber-White heteroskedasticity-robust (HC0) standard errors. Outliers are removed within each category-year cell using Tukey's rule $Q_1 \pm 3 \text{ IQR}$.

Bandwidth. The bandwidth h is selected using the Imbens–Kalyanaraman (2012) plug-in method, which minimizes the asymptotic mean squared error (MSE) of the sharp RDD estimator. This choice balances bias and variance in local polynomial fitting. Robustness checks with alternative bandwidths (e.g., Calonico–Cattaneo–Titiunik bias-corrected intervals) yield consistent results.



(a) Local cubic polynomial fit ($p = 3$); the x -axis is truncated at the chosen bandwidth h .

(b) Local linear fit ($p = 1$) with triangular kernel weights. Blue points: 2019 ($x < 0$), red points: 2020 ($x > 0$).

Figure 8: Comparison of RDD fits using different polynomial specifications.

VI.4 Identification via Rank-Improvement Probabilities

We provide a complementary, model-free identification argument based on the *proportion of athletes who improve their rank* via drafting.

Table 6: Sharp RDD estimates of rank improvement at the 2019/2020 cutoff

Model	h	n _L	n _R	$\hat{\tau}$	t / p
Local linear (p = 1)	13	718	1756	-30.654* (SE = 17.280)	-1.77 / 0.0761
Local polynomial (p = 3)	13	718	1756	-91.941*** (SE = 30.771)	-2.99 / 0.0028

Notes. Sharp RDD estimates computed with triangular kernel weights and bandwidth h selected via the Imbens–Kalyanaraman (2012) method. Heteroskedasticity-robust (HC0) standard errors in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Counts n_L , n_R denote the number of observations within the bandwidth after outlier removal by Tukey’s rule ($Q_1 \pm 3 \text{ IQR}$).

Setup (Rubin causal framework). Let

$$Y_{it} = \mathbf{1}\{\text{athlete } i \text{ improves rank in season } t\},$$

and define treatment

$$D_t = \mathbf{1}\{\text{season } t \text{ is non-COVID}\}.$$

Each athlete has potential outcomes $Y_{it}(1)$ and $Y_{it}(0)$ corresponding to non-COVID and COVID environments.

The estimand of interest is

$$\tau = \mathbb{E}[Y_{it}(1) - Y_{it}(0)],$$

interpreted as the causal change in the probability of rank improvement.

Proxy estimator. We proxy $\mathbb{E}[Y_{it}(d)]$ by the sample proportion

$$\hat{p}_d = \mathbb{E}[Y_{it} \mid D_t = d], \quad d \in \{0, 1\},$$

and estimate τ by the difference in proportions

$$\hat{\tau} = \hat{p}_1 - \hat{p}_0,$$

with standard t-based standard errors for binary outcomes.

Identification for compliers. Assume: (i) *Consistency*: $Y_{it} = Y_{it}(D_t)$; (ii) *Monotonicity*: $Y_{it}(1) \geq Y_{it}(0)$ for all i ; (iii) *Independence within rank group*: $\{Y_{it}(1), Y_{it}(0)\} \perp D_t$ conditional on belonging to a fixed top-K rank group.

Proposition VI.1 (Identification via proportions). *Under (i)–(iii),*

$$\hat{\tau} = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid \text{compliers}] \cdot \Pr(\text{complier}),$$

so $\hat{\tau} > 0$ implies a positive causal effect of non-COVID conditions on rank improvement for compliers.

Proof. By monotonicity, only compliers contribute to $Y(1) - Y(0)$. Independence implies $\mathbb{E}[Y(d)] = \mathbb{E}[Y \mid D = d]$ for $d \in \{0, 1\}$, so the difference in observed proportions identifies the stated quantity. $\square$

Empirical evidence. Figures 9–13 show that $\hat{p}_1 > \hat{p}_0$ across all top-K thresholds, indicating that rank improvement via drafting is systematically less likely during the COVID season.




Figure 9: Share of TOP-5 athletes improving rank via drafting by period. Bars show proportions; whiskers denote 95% t-based confidence intervals.




Figure 10: Draft-induced rank improvements among TOP-10 athletes by two-year bins; the COVID season (2020) is shown separately.




Figure 11: Annual share of TOP-10 athletes improving rank via drafting, 2010–2024, with 95% confidence intervals.




Figure 12: Share of TOP-20 athletes improving rank via drafting by period.




Figure 13: Share of TOP-30 athletes improving rank via drafting by period.

Corollary VI.1 (Local average treatment effect interpretation). *Under Assumptions (i)–(iii) of Proposition VI.1, the difference in rank-improvement probabilities between non-COVID and COVID seasons identifies a local average treatment effect (LATE) for athletes whose rank outcomes are affected by the competitive environment:*

$$\text{LATE} = \mathbb{E}[Y_{it}(1) - Y_{it}(0) \mid \text{compliers}] = \frac{\hat{p}_1 - \hat{p}_0}{\Pr(\text{complier})}.$$

Connection to the event-study estimates. The positive differences in rank-improvement probabilities documented above correspond directly to positive average shifts in the conditional mean of rank improvement, which is precisely the object captured by the non-COVID event-study coefficients in equations (VIII.4)–(XI.1).

VII. Distributional Effects of the COVID Shock on Athletic Performance

VII.1 Conceptual Framework

The COVID-19 pandemic represents a sharp and plausibly exogenous shock to competitive environments. Public health restrictions altered training opportunities, race schedules, and participation margins in ways that are unlikely to affect all athletes symmetrically. Standard mean-based treatment effects may therefore mask substantial heterogeneity in responses across the performance distribution.

We conceptualize performance as the outcome of both individual ability and strategic adaptation. Athletes with higher baseline ability may benefit disproportionately from reduced competition density, altered selection into events, or improved relative preparation during periods of disruption. Conversely, lower-ranked athletes may experience muted or even negative effects if access to training and competition deteriorates.

These considerations motivate a distributional approach that explicitly allows treatment effects to vary across quantiles of the outcome distribution.

VII.2 Empirical Design

We study the effect of the COVID-era regime using a regression discontinuity design (RDD) in event time. The running variable is event year normalized relative to 2020, with a cutoff at zero. The treatment of interest is exposure to the post-COVID regime, which changes discontinuously—but not deterministically—at the cutoff due to staggered restrictions and partial compliance. This motivates a fuzzy RDD.

The outcome variable is percentile rank improvement, computed within event-by-year cells. This transformation ensures comparability across heterogeneous races and captures relative performance gains rather than absolute finishing times.

VII.3 Fuzzy Quantile Regression Discontinuity

To estimate heterogeneous effects, we implement fuzzy quantile regression discontinuity designs at the 10th, 50th, and 90th percentiles of the outcome distribution. Identification follows a local control-function approach. In the first stage, treatment status is regressed on the running variable separately on either side of the cutoff. In the second stage, quantile regressions include the treatment indicator, the running variable, and the first-stage residual.

Estimation uses a triangular kernel with a baseline bandwidth of ± 3 years. All specifications are local linear.

The figures reveal a clear pattern of increasing treatment effects across the performance distribution. While median effects are positive and statistically meaningful, effects at the upper tail are substantially larger, indicating that the COVID-era regime disproportionately benefited top-performing athletes.

VII.4 Robustness and Validation

We conduct several robustness checks to validate the design.

First, Figure 17 reports bandwidth sensitivity of the median fuzzy quantile treatment effect. Estimates are stable across a wide range of bandwidth choices.




Figure 14: Fuzzy Quantile Regression Discontinuity at the 10th Percentile




Figure 15: Fuzzy Quantile Regression Discontinuity at the Median




Figure 16: Fuzzy Quantile Regression Discontinuity at the 90th Percentile




Figure 17: Bandwidth Sensitivity of Median Fuzzy QTE

Second, we implement placebo tests using artificial cutoffs in 2018 and 2019. Figures 18 and 19 show no evidence of discontinuities at these placebo thresholds, supporting the identifying assumption that the observed effects are unique to the COVID shock.




Figure 18: Placebo Cutoff at 2018




Figure 19: Placebo Cutoff at 2019

Taken together, these results provide strong evidence that the COVID-era regime induced substantial and heterogeneous changes in relative athletic performance, with the largest gains accruing to athletes at the top of the performance distribution.

VIII. Empirical Framework

VIII.1 OLS with High-Dimensional Fixed Effects: Convexity, Grouped-Gram Estimation, and Fast Inference

This subsection formalizes the fixed-effects OLS estimator underlying (VIII.4)–(VIII.5), establishes its equivalence to a grouped-Gram representation, and shows that computational cost of such models is asymptotically dominated by Gram matrix construction rather than by fixed-effect absorption itself.

Convex OLS problem with fixed effects. Let $\mathbf{y} \in \mathbb{R}^n$ denote the stacked outcome, $\mathbf{X} \in \mathbb{R}^{n \times k}$ the matrix of structural regressors, $\mathbf{D}_1 \in \{0, 1\}^{n \times N}$ athlete fixed effects, and $\mathbf{D}_2 \in \{0, 1\}^{n \times C}$ event-category fixed effects. The estimator solves the convex quadratic program

$$\min_{\beta, \alpha, \lambda} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta - \mathbf{D}_1\alpha - \mathbf{D}_2\lambda\|_2^2. \quad (\text{VIII.1})$$

After standard normalizations on (α, λ) , the objective is strictly convex and admits a unique minimizer.

Within representation. Let $\mathbf{D} = [\mathbf{D}_1 \ \mathbf{D}_2]$ and define the annihilator matrix $\mathbf{M}_\mathbf{D} = \mathbf{I} - \mathbf{D}(\mathbf{D}'\mathbf{D})^\dagger \mathbf{D}'$. The first-order conditions imply that $\hat{\beta}$ satisfies

$$\mathbf{X}'\mathbf{M}_\mathbf{D}(\mathbf{y} - \mathbf{X}\beta) = 0, \quad (\text{VIII.2})$$

corresponding to OLS on variables residualized with respect to all fixed effects.

Theorem VIII.1 (Frisch–Waugh–Lovell with multiple fixed effects). *The coefficient vector $\hat{\beta}$ solving (VIII.1) is identical to the OLS estimator from the projected regression $\mathbf{M}_\mathbf{D}\mathbf{y} = \mathbf{M}_\mathbf{D}\mathbf{X}\beta + \mathbf{M}_\mathbf{D}\varepsilon$.*

Proof. Substituting the optimal fixed effects $\hat{\delta}(\beta) = (\mathbf{D}'\mathbf{D})^\dagger \mathbf{D}'(\mathbf{y} - \mathbf{X}\beta)$ into the normal equations yields (VIII.2). Because the matrix $\mathbf{M}_\mathbf{D}$ is symmetric and idempotent, the projected regression defines a strictly convex quadratic in β with the same minimizer as in the original problem. $\square$

Grouped-Gram formulation. Classical within estimation explicitly constructs demeaned arrays, which is memory-intensive for large n and k . An alternative approach computes the required cross-products using grouped sufficient statistics. For each fixed-effect unit i , define

$$s_{x,i} = \sum_t x_{it}, \quad s_{y,i} = \sum_t y_{it}, \quad G_i = \sum_t x_{it}x'_{it}.$$

Aggregating $\{s_{x,i}, s_{y,i}, G_i\}$ across units reproduces the cross-products $\mathbf{X}'\mathbf{M}_\mathbf{D}\mathbf{X}$ and $\mathbf{X}'\mathbf{M}_\mathbf{D}\mathbf{y}$ exactly, without forming demeaned observations.

Theorem VIII.2 (Grouped-Gram equivalence). *The grouped-Gram estimator constructed from unit-level sufficient statistics yields exactly the same $\hat{\beta}$ as the classical within estimator solving (VIII.1).*

Proof. Both estimators solve the same normal equations $\mathbf{X}'\mathbf{M}_\mathbf{D}\mathbf{X}\beta = \mathbf{X}'\mathbf{M}_\mathbf{D}\mathbf{y}$. The grouped-Gram construction reproduces these matrices algebraically using only linearity of summation. Hence the resulting $\hat{\beta}$ is identical. $\square$

Lemma VIII.1 (Fixed-effects estimation reduces to Gram formation). *For linear models with additive fixed effects, the computational cost of estimating $\hat{\beta}$ is asymptotically dominated by forming the Gram matrix $X'M_D X$ and vector $X'M_D y$. Absorption of fixed effects requires no operations beyond those needed to compute these cross-products.*

Proof. By Theorem VIII.1, estimation depends only on the projected normal equations. The grouped-Gram formulation constructs $X'M_D X$ and $X'M_D y$ directly from sufficient statistics without explicit demeaning or inversion of $D'D$. Solving the resulting $k \times k$ linear system is $\mathcal{O}(k^3)$ and negligible relative to forming the Gram matrices when $n \gg k$. Thus overall cost is driven by Gram formation. $\square$

Corollary VIII.1 (Computational efficiency). *Grouped-Gram estimation reduces memory movement from $\mathcal{O}(nk)$ to $\mathcal{O}(Nk^2)$ per fixed-effect dimension and avoids explicit within transformations, yielding substantial runtime gains in large panels.*

Some formal analysis and benchmarks of this estimator are provided by (Reichel 2026).

Fast variance–covariance construction. Furthermore, statistical inference on $\hat{\beta}$ requires covariance matrices of regressors and residuals. (Reichel 2025a) shows that the standard covariance matrix admits the uncentered Gram representation

$$\text{Cov}(X) = \frac{1}{n-1} \left(X^\top X - \frac{1}{n} s s^\top \right), \quad s = X^\top \mathbf{1}_n,$$

which is algebraically identical to the centered formulation but avoids explicit demeaning, reducing computation to a single outer product and subtraction.

Because both estimation and inference reduce to Gram computations, optimized Gram kernels directly translate into end-to-end (e2e) speedups. In particular, (Rybin, Zhang, and Luo 2025) show that specialized routines for computing Gram products and $XX^\top$ which can be substantially faster than generic or outdated/non-updated (Open-)BLAS implementations by exploiting data locality and hardware-level parallelism.¹²

Implication for the empirical analysis. All specifications estimated in Sections VIII–VIII.4 are exact solutions to the convex problem (VIII.1). Modern grouped-Gram and uncentered-Gram techniques affect only computational efficiency, not identification, estimation, or inference.

VIII.2 Empirical Analysis using 2020 as reference year

We study the outcome FR_{iet} (“free-riding”) defined by centering raw rank-improvement at year 2020:

$$FR_{iet} \equiv RI_{iet} - \overline{RI}_{y=2020}, \quad \Rightarrow \quad \mathbb{E}[FR_{iet} \mid y=2020] = 0. \quad (\text{VIII.3})$$

For each heterogeneity dimension H_i (gender, age bin, drafter position, swimmer ability, team composition) we estimate a pooled event–study with athlete and event–category fixed effects:

$$FR_{iet} = \sum_{y \in \mathcal{T}} \sum_{h \in \mathcal{H}} \beta_{y,h} \mathbf{1}\{year_t = y\} \mathbf{1}\{H_i = h\} + \alpha_i + \lambda_{c(e)} + \varepsilon_{iet}, \quad (\text{VIII.4})$$

1. For example, the OpenBLAS project (<https://github.com/OpenMathLib/OpenBLAS>) provides an open-source optimized implementation of the BLAS and LAPACK APIs for modern CPUs.

2. For instance, the OpenBLAS project (<https://github.com/OpenMathLib/OpenBLAS/commit/71261a7b3f77363ccf15a5f893d3687c026c04e1>) implements optimized summation kernels by derivation from existing SASUM/DASUM vectorized kernels in basic BLAS procedures.

where α_i are athlete FEs and $\lambda_{c(e)}$ are event–category FEs. Standard errors are two–way clustered by *athlete* and *event*. We plot year–by–group coefficients $\hat{\beta}_{y,h}$ as candlesticks: the *rectangle* is the 95% CI, the *horizontal tick* is the point estimate; no connecting lines are drawn. All years (2015–2024) appear on the horizontal axis, and groups are shown only if they have at least one observation in 2020 (so the reference $y=2020$ is meaningful). No additional covariates are included beyond the fixed effects.

VIII.3 Baseline Pooled Yearly Specification

We begin with a flexible year–by–gender specification that also controls for athlete age and event field size. At the athlete–event level, we estimate two outcomes: free–riding FR_{iet} (rank improvement centered at 2020) and raw rank–improvement RI_{iet} . For a generic outcome $Y_{iet} \in \{FR_{iet}, RI_{iet}\}$, the regression is

$$\begin{aligned}
 Y_{iet} = & \sum_{y \in \mathcal{Y}} \sum_{g \in \{F,M\}} \beta_{y,g} \mathbf{1}\{\text{year}_t = y\} \mathbf{1}\{\text{gender}_{it} = g\} \\
 & + \gamma_1 \text{age}_{it} + \gamma_2 \text{age}_{it}^2 \\
 & + \delta_1 \mathbf{1}\{\text{medium}_e\} + \delta_2 \mathbf{1}\{\text{large}_e\} + \alpha_i + \lambda_{c(e)} + \varepsilon_{iet}.
 \end{aligned}
 \tag{VIII.5}$$

Here, α_i are athlete fixed effects, $\lambda_{c(e)}$ are event–category fixed effects, and standard errors are two–way clustered by athlete and event.

Outcome: Free–riding.

$$FR_{iet} = \text{RHS of (VIII.5)}. \tag{VIII.6}$$




Figure 20: Year×gender coefficients for free–riding (points = estimates, bars = 95% CIs).

Table 7: Baseline FE regression (outcome: free–riding)

Regressor	Estimate	Std. Error	t	p
Age ²	0.0351	0.0121	2.895	0.00393
1{group size = medium}	18.049	6.756	2.671	0.00775
1{group size = large}	31.565	5.952	5.303	1.57×10^{-7}
RMSE				49.1
Adj. R ²				0.509
Within R ²				0.070
Observations	15,028; athletes = 1,965; categories = 4			
SEs	clustered by athlete and event			

Outcome: Rank-improvement.

$$RI_{i\text{et}} = \text{RHS of (VIII.5).} \tag{VIII.7}$$




Figure 21: Year×gender coefficients for rank-improvement (95% CIs).

Table 8: Baseline FE regression (outcome: rank-improvement)

Regressor	Estimate	Std. Error	t	p
Age ²	0.0351	0.0121	2.904	0.00381
1{group size = medium}	18.049	6.759	2.670	0.00777
1{group size = large}	31.565	5.956	5.300	1.60 × 10 ^{−7}
RMSE				49.1
Adj. R ²				0.509
Within R ²				0.070
Observations	15,028; athletes = 1,965; categories = 4			
SEs	clustered by athlete and event			

Remarks. Across both outcomes, year×gender coefficients are statistically indistinguishable from zero, while larger event fields are associated with significantly higher free-riding and rank-improvement. The convex age profile ($\gamma_2 > 0$) is small but precisely estimated.

As our parametrically linear regression models should absorb the same (source of) variance both outcome variables lead to same magnitude and SE in coefficients.

VIII.4 Alternative Pooled Period Specification

While the baseline is flexible, it can be noisy when estimating year-specific effects. We therefore also employ a pooled event-study design that collapses years into three periods: pre-2020, 2020, and post-2020. This provides a more parsimonious estimate of systematic shifts:

$$\begin{aligned}
FR_{iet} = & \beta_{pre} \mathbf{1}\{pre\} + \beta_{post} \mathbf{1}\{post\} \\
& + \sum_v \sum_{\ell \neq \ell_0(v)} \left[\beta_{pre, v=\ell} \mathbf{1}\{pre\} \mathbf{1}\{v = \ell\} + \beta_{post, v=\ell} \mathbf{1}\{post\} \mathbf{1}\{v = \ell\} \right] \\
& + \alpha_i + \lambda_{c(e)} + \varepsilon_{iet}.
\end{aligned} \tag{VIII.8}$$

Athlete FEs (α_i) absorb permanent ability, and event-category FEs ($\lambda_{c(e)}$) capture persistent race differences. Standard errors are two-way clustered by athlete and event. The omitted period is 2020, so all coefficients are relative to the COVID season.

IX. Empirical Results

IX.1 Effect Heterogeneity by Gender

Specification. In (VIII.4) set $H_i = \text{gender} \in \{\text{Female}, \text{Male}\}$ (Unknown/Missing are omitted from the plot when absent or unsupported in 2020).




Figure 22: Year×gender event-study for free-riding. Candlestick body = 95% CI, tick = estimate $\hat{\beta}_{y,h}$. FE at athlete and event-category; SEs clustered by athlete and event; reference year 2020.

Interpretation. Female and Male series broadly follow the aggregate profile: positive pre-2020 effects, normalization at 2020, and a partial post-2020 rebound. Across most years the gender gap is modest and often within the 95% CI, indicating limited systematic gender heterogeneity.

IX.2 Effect Heterogeneity by Age

Specification. Let $H_i = \text{age_bin}$ (5-year bins over the observed range; the code retains a “Missing” bin when present).




Figure 23: Year×age-bin event-study for free-riding. Bins ordered by lower bound; candlesticks show 95% CIs; ref = 2020.

Interpretation. All age bins show positive pre-2020 effects; post-2020 estimates remain above zero but below pre-period magnitudes. Differences across age bins are small with overlapping CIs in most years.

IX.3 Effect Heterogeneity by Drafter Position

Specification. Let $H_i = \text{drafter_position}$ binned into $\{2, 3, 4, 5, 6, 7, \text{other}\}$.




Figure 24: Year×drafter position (bins) event–study for free–riding. 95% CIs; athlete and event–category FEs; clustered SEs; ref = 2020.

Interpretation. Pre-2020 coefficients are larger for later (higher) drafter positions, consistent with more scope to free-ride from the back of the pack. Post-2020 values increase again across bins, but many pairwise differences are not statistically distinct at the 5% level.

IX.4 Effect Heterogeneity by Swimmer Ability

Specification. Let $H_i = \text{swim_skill} \in \{\text{good}, \text{average}, \text{bad}\}$, computed via k-means on athletes' best swim times by category (fastest = *good*).




Figure 25: Year×swimmer ability event–study for free–riding (good/average/bad). 95% CIs; ref = 2020.

Interpretation. Ability sorts the magnitudes: the *bad* group shows the largest pre-2020 effects and a clearer rebound after 2020; *good* swimmers display the smallest swings. CIs often separate *bad* from *good*, indicating meaningful heterogeneity by ability.

IX.5 Effect Heterogeneity by Group Gender Composition Mix

Specification. Let $H_i = \text{gender_mix_category} \in \{\text{pure male, male 80–20, balanced 50–50, pure female, other}\}$. Small cells are included only when they have at least one 2020 observation.




Figure 26: Year×team composition (gender mix) event–study for free–riding. Candlesticks = 95% CIs; ref = 2020.

Interpretation. Larger mixes (e.g., *pure male, male 80–20*) display precise positive effects pre-2020 and a partial recovery post-2020; rarer mixes are more variable with wider CIs. Where sample sizes are adequate, cross-mix differences are moderate.

Estimation notes. All results are from pooled OLS with athlete and event–category fixed effects. Standard errors are two–way clustered by athlete and event. Groups are plotted only if they have at least one observation in 2020; all years 2015–2024 appear on the x–axis. No additional covariates beyond fixed effects are included.

IX.6 Effect Heterogeneity by Group Size




Figure 27: Free–riding coefficients by group–size bins (95% CIs).




Figure 28: Rank-improvement coefficients by group-size bins (95% CIs).

X. Estimation Framework using 2019 as reference year

Model. Let y_{it} denote athlete i 's Free Riding in year t . We estimate a fully interacted (pooled) two-way OLS with high-dimensional fixed effects:

$$y_{it} = \alpha_i + \mathbf{F}_{it}'\gamma + \sum_{s \in \mathcal{S}} \mathbf{1}\{S_i = s\} \left(\sum_{y \in \mathcal{Y} \setminus \{2019\}} \beta_y^{(s)} \mathbf{1}\{\text{year} = y\} \right) + \varepsilon_{it}, \quad (\text{X.1})$$

where α_i are athlete fixed effects and $\mathbf{F}_{it}$ collects the fixed effects used in each table (group size, group gender composition mix, drafter position, event category, event name, event location, and—when applicable—athlete swim skill). Year 2019 is the omitted (baseline) category, so $\beta_y^{(s)}$ measures the deviation in Free Riding in year y relative to 2019 for subgroup s . Standard errors are clustered at the athlete level (Tables A.8.1, A.8.2, A.8.4) and at the event level (Table A.8.3), matching the tables.

Subgroup contrasts. For any two subgroups s and s' , the pooled model implies the cross-group year contrast can be simply calculated as

$$\Delta_y^{(s-s')} \equiv \beta_y^{(s)} - \beta_y^{(s')},$$

which is the difference in year- y deviations relative to 2019 between s and s' .

A.8.1 Heterogeneous Effects by Gender on Free Riding

Gender. Free riding (FR) moves markedly over time for both sexes, with a sharper 2020 contraction for *women*. In the gender splits, FR falls in 2020 by -21.851 for men versus -35.052 for women, from generally positive pre-2020 levels (e.g., 2015: 26.205 for men, 13.076 for women). By 2024, women show a clearer rebound (9.997) while the men’s coefficient is smaller (2.535). These patterns suggest the 2020 format changes suppressed FR broadly, with a larger immediate impact among women and only a partial, female-led recovery thereafter.




Figure 29: Year FEs on Free Riding by Gender (Male vs. Female), baseline=2019

X.1 Heterogeneous Effects by Swim Ability on Free Riding

Swim ability. Heterogeneity by swim skill is pronounced. “Bad” and “Average” swimmers exhibit large positive FR pre-2020 (e.g., 2015: 45.087 for Bad, 25.862 for Average) and the steepest 2020 drops (−32.809 and −30.837, respectively). By contrast, “Good” swimmers’ FR is only moderately negative in 2020 (−7.957). Post-2020 recovery is tepid: Bad shows a marginal uptick in 2024 (10.474), while Good and Average are near zero and imprecise. This pattern is consistent with drafting benefits being especially consequential for weaker swimmers; when drafting was curtailed, their relative FR contracted the most.




Figure 30: Year FEs by Swim Ability (Good/Bad/Average), baseline=2019.

X.2 Heterogeneous Effects by Age on Free Riding

Age. Year-by-age estimates show positive pre-2020 FR in many bins (e.g., 2017: 55.286 for ages 50–55), followed by broad contractions in 2020 across all age groups (e.g., 40–45: –28.724; 45–50: –32.547; 60–65: –34.659). Post-2020 coefficients are mixed and mostly imprecise at annual frequency, with scattered negatives (e.g., 2024: –19.976 for ages 35–40) and few clear positives. Taken together, the sharp 2020 declines are ubiquitous by age, while post-2020 rebounds appear selective rather than uniform.




Figure 31: Year FEs by Age bin (20–25, . . . , 60–65), baseline=2019.

X.3 Age heterogeneity from the baseline FE model

Interpretation. The baseline fixed-effects regression of rank-improvement includes a quadratic age term and two pack-size dummies:

$$\widehat{\text{RI}}_{\text{iet}} = \beta_{A^2} \text{Age}_{\text{it}}^2 + \beta_M \mathbf{1}\{\text{group size} = \text{medium}\}_{\text{et}} + \beta_L \mathbf{1}\{\text{group size} = \text{large}\}_{\text{et}} + \alpha_i + \lambda_e + \varepsilon_{\text{iet}}.$$

Estimated coefficients are $\hat{\beta}_{A^2} = 0.0351$ (SE 0.0121, $t = 2.904$, $p = 0.0038$); $\hat{\beta}_M = 18.049$ (SE 6.759, $p = 0.0078$); $\hat{\beta}_L = 31.565$ (SE 5.956, $p = 1.6 \times 10^{-7}$). The positive $\hat{\beta}_{A^2}$ implies a *convex* age profile: within athlete, rank-improvement rises with age at an accelerating rate. Medium and large packs shift the profile upward by about +18 and +31.6 points, respectively. Fit diagnostics (RMSE = 49.1, Adj. $R^2 = 0.509$, Within $R^2 = 0.070$) indicate that while fixed effects capture substantial cross-sectional structure, the within-variation explained by these age/pack terms alone is modest, as expected.

Predictive mapping (relative to a reference age). Because levels are absorbed by fixed effects, we plot *relative* predictions normalized at a reference age of $a_0 = 25$:

$$\widehat{\Delta \text{RI}}(a; a_0, g) = \hat{\beta}_{A^2} (a^2 - a_0^2) + \mathbf{1}\{g = \text{medium}\} \hat{\beta}_M + \mathbf{1}\{g = \text{large}\} \hat{\beta}_L,$$

with $g \in \{\text{small (baseline), medium, large}\}$. For example, moving from $a_0 = 25$ to $a = 40$ (holding group size fixed) changes predicted rank-improvement by $0.0351 \times (40^2 - 25^2) =$




Figure 32: Age heterogeneity (step function) including *all pack-size groups*. Lines show model-implied predictions by age bin using $\widehat{\Delta RI}(a) = 0.0351(a^2 - 25^2)$, with vertical shifts of +18.05 (medium) and +31.57 (large) relative to the small-group as baseline group.

$$0.0351 \times 975 \approx 34.2 \text{ points.}$$

Reading the figure. At younger ages near $a_0 = 25$, predicted differences are small by construction; as age increases, the convex term $\hat{\beta}_{A^2} > 0$ steepens the slope, implying larger gains in rank-improvement at older ages (within athlete). Independent of age, medium and large packs add about +18 and +31.6 points, respectively. Magnitudes are relative and should be read as fixed-effects contrasts; without a linear age term, curvature is the primary within-athlete age signal in this specification.

X.4 Heterogeneous Effects by Drafter Position on Free Riding

Drafter position. FR varies by pack position. Before 2020, several positions show sizable positives (e.g., 2015: 43.445 at position 3). In 2020, FR contracts at multiple positions, with especially large or significant drops at positions 1 (−31.538), 2 (−17.886), 4 (−53.006), and 6 (−23.429). Later, some positions show strong positives in 2024 (e.g., position 5: 72.760; position 8: 37.775), pointing to a return of drafting-compatible race conditions that amplify FR deeper in the pack.




Figure 33: Year FEs by Drafter Position (1–8), baseline=2019.

Synpaper with period-pooled interactions. The pooled event study confirms higher FR pre-2020 (main effect +13.43, 95% CI [0.18, 26.67]) and clarifies where post-2020 heterogeneity concentrates: *mid-age bins* (roughly 30–55) show the clearest post increases; *male-heavy team mixes* exhibit *lower* FR after 2020 (pure male: −41.18; male 80–20: −32.20); *smaller/medium packs* are associated with *lower* FR (post, medium: −15.67); and *back-of-pack positions* carry *higher* FR post-2020 (“other”: 15.67; position 7: 10.60; position 6: 9.39). Gender differences are modest overall, with a small male post-2020 increase (10.89). Altogether, the split-sample and pooled results paint a coherent story: drafting constraints in 2020 suppressed FR broadly but *heterogeneously*—most for weaker swimmers and many age groups—and the subsequent partial normalization raised FR especially in mid-age cohorts and deeper pack positions, while male-heavy fields and thinner packs tempered FR.

XI. Robustness: Significant Interactions in Period-Pooled OLS Event Study

Specification

We estimate a pooled OLS with athlete fixed effects and event-category fixed effects, interacting three period indicators (pre, 2020, post) with observed heterogeneity along several dimensions:

$$v \in \{\text{gender, age bin, drafter position, team gender mix, pack size}\}.$$

Let t index periods and $c(e)$ denote the event-category of event e . Using 2020 as the baseline period and, within each heterogeneity dimension v , a reference category $\ell_0(v)$, the estimating equation is

$$\begin{aligned} FR_{iet} = & \beta_{pre} \mathbf{1}\{t = pre\} + \beta_{post} \mathbf{1}\{t = post\} \\ & + \sum_v \sum_{\ell \neq \ell_0(v)} \left[\beta_{pre, v=\ell} \mathbf{1}\{t = pre\} \mathbf{1}\{v_{ie} = \ell\} + \beta_{post, v=\ell} \mathbf{1}\{t = post\} \mathbf{1}\{v_{ie} = \ell\} \right] \\ & + \alpha_i + \lambda_{c(e)} + \varepsilon_{iet}, \end{aligned} \tag{XI.1}$$

where α_i are athlete fixed effects and $\lambda_{c(e)}$ are event-category fixed effects. Standard errors are two-way clustered by athlete and event.

Table 9: Pooled OLS with athlete and event-category FE (periods = pre/2020/post). Coefficients grouped by period and dimension, ordered by category values (e.g., age bins ascending; drafter positions 3–7, then other; pack sizes small→medium; team mix in the order shown). Estimate cells use a blue–white–red gradient with intensity proportional to $|\hat{\beta}|$. Baselines are 2020 (period) and, within each dimension v , the reference category $\ell_0(v)$.

Term	Estimate	Lower 95%	Upper 95%	Sig.
Period: Pre				
<i>Main effect</i>				
Pre (vs. 2020)	13.43	0.18	26.67	✓
<i>Age bin (ascending)</i>				
Pre × Age bin: [10, 15]	−6.31	−25.57	12.95	
Pre × Age bin: (20, 25]	23.42	15.60	31.25	✓
Pre × Age bin: (25, 30]	32.98	24.46	41.50	✓
Pre × Age bin: (30, 35]	27.31	18.76	35.86	✓
Pre × Age bin: (35, 40]	35.09	26.43	43.74	✓
Pre × Age bin: (40, 45]	34.39	25.84	42.94	✓
Pre × Age bin: (50, 55]	19.21	9.28	29.14	✓
Pre × Age bin: (55, 60]	13.34	0.91	25.78	✓
Pre × Age bin: (60, 65]	9.81	−6.07	25.68	
Pre × Age bin: (65, 70]	13.86	−11.42	39.14	
Pre × Age bin: (70, 75]	67.15	−57.42	191.72	
<i>Pack size (small → medium)</i>				
Pre × Pack size: small	−30.28	−40.75	−19.82	✓
Pre × Pack size: medium	−11.52	−17.51	−5.52	✓
<i>Team mix (pure female, male 80–20, pure male, other)</i>				
Pre × Team mix: pure female	25.30	11.27	39.33	✓
Pre × Team mix: male 80–20	−10.88	−21.87	0.11	
Pre × Team mix: pure male	1.88	−10.27	14.03	
Pre × Team mix: other	−7.21	−18.06	3.64	
<i>Drafter position (3,4,5,6,7, other)</i>				

cont'd on next page

<i>(cont'd)</i> Pooled OLS with athlete and event-category FE (periods = pre/2020/post).				
Term	Estimate	Lower 95%	Upper 95%	Sig.
Pre × Drafter position: 3	2.43	−2.02	6.89	
Pre × Drafter position: 4	7.50	2.57	12.42	✓
Pre × Drafter position: 5	2.71	−2.26	7.67	
Pre × Drafter position: 6	3.22	−2.30	8.74	
Pre × Drafter position: 7	3.77	−2.29	9.84	
Pre × Drafter position: other	6.39	2.71	10.07	✓
<i>Gender</i>				
Pre × Gender: male	3.66	−2.59	9.90	
Period: Post				
<i>Main effect</i>				
Post (vs. 2020)	5.72	−12.24	23.68	
<i>Age bin (ascending)</i>				
Post × Age bin: [10, 15]	−9.88	−56.61	36.86	
Post × Age bin: [20, 25]	5.21	−6.99	17.40	
Post × Age bin: (25, 30]	19.85	6.11	33.59	✓
Post × Age bin: (30, 35]	33.52	20.20	46.83	✓
Post × Age bin: (35, 40]	33.09	19.65	46.52	✓
Post × Age bin: (40, 45]	34.91	21.60	48.22	✓
Post × Age bin: (45, 50]	39.30	25.84	52.75	✓
Post × Age bin: (50, 55]	35.48	21.78	49.19	✓
Post × Age bin: (55, 60]	21.83	7.75	35.92	✓
Post × Age bin: (60, 65]	26.02	8.65	43.39	✓
Post × Age bin: (65, 70]	22.25	3.16	41.34	✓
Post × Age bin: (70, 75]	13.47	−18.96	45.91	
Post × Age bin: (75, 80]	−11.65	−74.83	51.52	
<i>Pack size (small → medium)</i>				
Post × Pack size: small	−16.71	−36.05	2.64	
Post × Pack size: medium	−15.67	−23.79	−7.55	✓
<i>Team mix (pure male, male 80–20, other, pure female)</i>				
Post × Team mix: pure male	−41.18	−59.53	−22.82	✓
Post × Team mix: male 80–20	−32.20	−45.32	−19.09	✓
Post × Team mix: other	−15.70	−28.60	−2.79	✓
Post × Team mix: pure female	4.06	−24.88	32.99	
<i>Drafter position (3,4,5,6,7, other)</i>				
Post × Drafter position: 3	4.41	−0.21	9.03	
Post × Drafter position: 4	4.02	−1.01	9.06	
Post × Drafter position: 5	5.48	−0.89	11.86	
Post × Drafter position: 6	9.39	2.58	16.19	✓
Post × Drafter position: 7	10.60	2.25	18.96	✓
Post × Drafter position: other	15.67	11.20	20.15	✓
<i>Gender</i>				
Post × Gender: male	10.89	4.79	16.99	✓
<i>Notes:</i> Pooled OLS with athlete FE (α_i) and event-category FE ($\lambda_{c(e)}$). Standard errors are two-way clustered by athlete and event. Period dummies use 2020 as the reference. Heterogeneity groups (Age bin, Pack size, Team mix, Drafter position, Gender) enter as interactions with Pre and Post dummies, relative to each group's reference category. Rows within each subgroup are ordered by category values. A checkmark indicates the 95% confidence interval excludes zero.				

Interpretation.

- **Pre vs. 2020 (main effect).** Free-riding (FR) is higher pre-2020 relative to 2020:




Figure 34: Pooled OLS interaction coefficients by period (pre/2020/post), sorted by $|\hat{\beta}|$ within period $\times$ dimension. Points are colored by sign; vertical bars are 95% CIs. Baselines are 2020 (period) and, within each dimension v , the reference category $\ell_0(v)$.

$\hat{\beta}_{\text{pre}} = 13.43$ (95% CI [0.18, 26.67]).

- **Age heterogeneity (post vs. 2020).** Post-2020 increases in FR are concentrated in mid-age bins: (45, 50] (39.30 [25.84, 52.75]), (50, 55] (35.48 [21.78, 49.19]), (40, 45] (34.91 [21.60, 48.22]), (30, 35] (33.52 [20.20, 46.83]). Very young/old bins are imprecisely estimated.
- **Team mix (post vs. 2020).** Male-heavy fields are associated with *lower* FR: pure male (−41.18 [−59.53, −22.82]); male 80–20 (−32.20 [−45.32, −19.09]). “Other” is modestly negative (−15.70 [−28.60, −2.79]).
- **Pack size.** Smaller/medium packs correlate with lower FR (pre, small: −30.28 [−40.75, −19.82]; post, medium: −15.67 [−23.79, −7.55]).
- **Drafter position (post vs. 2020).** Back-of-pack positions show higher FR (“other”: 15.67 [11.20, 20.15]; 7: 10.60 [2.25, 18.96]; 6: 9.39 [2.58, 16.19]).
- **Gender.** Modest differences: males show a small post-2020 increase (10.89 [4.79, 16.99]); the pre-period male effect is not significant.

Overall, FR is higher before 2020, with selective post-2020 rebounds concentrated in mid-age groups, attenuated in male-heavy and thinner fields, and amplified for back-of-pack positions—patterns consistent with drafting-based mechanisms rather than broad gender-level shifts.

XII. On the Economics of Free Riding in Triathlon

The phenomenon of *free riding* or *slipstreaming* in endurance sports such as triathlon provides a natural field laboratory for the application of microeconomic concepts. Free riding occurs when athletes strategically exploit drafting effects in swimming or cycling, thereby conserving energy while benefiting from the effort of others. We analyze free riding using classical tools of economic theory, including elasticity, production frontiers, Engel curves, Beveridge curves, inequality measurement, and growth incidence curve (GIC) analysis.

XII.1 The Elasticity of Free Riding as a Public Good Problem

Simply consider a swim group of size n , where each member contributes effort toward a collective outcome (e.g., maintaining pace, sharing training resources). Let g denote the aggregate group good produced by total contributions. Each individual i chooses effort $e_i \geq 0$, which entails a private cost $c(e_i)$ with $c' > 0$, $c'' > 0$. The group good is given by

$$g = \sum_{i=1}^n e_i,$$

and generates drafting benefits $B(g)$, with $B' > 0$, $B'' < 0$, that are non-excludable within the group.

In this setting, free riding arises because the marginal private benefit of an additional unit of effort is diluted by group size:

$$\frac{\partial U_i}{\partial e_i} = \frac{1}{n} B'(g) - c'(e_i),$$

where each swimmer captures only $\frac{1}{n}$ of the marginal group benefit but bears the full marginal cost of effort. As n increases, the incentive to reduce own contribution—and thus to free ride—also rises.

Social Optimum vs. Nash Equilibrium. A social planner maximizing group welfare would choose $\{e_i\}$ such that

$$B'(g) = c'(e_i), \quad \forall i,$$

ensuring that the marginal social benefit of effort equals its marginal cost. In contrast, in the non-cooperative Nash equilibrium, each individual contributes only until

$$\frac{1}{n} B'(g) = c'(e_i).$$

Because $\frac{1}{n} B'(g) < B'(g)$ for any $n > 1$, equilibrium effort is inefficiently low relative to the social optimum. The gap widens with group size, reflecting the classical free-rider problem as it appears in public good provision.

Empirical Estimation. To quantify this relationship empirically, we model observed free riding (FR) as a function of group size n using a log-log specification:

$$\log(1 + FR_i) = \alpha + \beta \log(1 + n_i) + \varepsilon_i.$$

Here, β represents the elasticity of free riding with respect to group size:

$$\varepsilon_{FR,n} = \frac{\partial FR}{\partial n} \cdot \frac{n}{FR}.$$




Figure 35: Elasticity sketch: $\log(1+FR)$ vs. $\log(1+n)$. Estimated elasticity $\hat{\beta} \approx 0.427$ (95% CI: $[0.393, 0.461]$). For a finite change $n \rightarrow n'$, $\% \Delta(1+FR) = \left[\left(\frac{1+n'}{1+n} \right)^{\hat{\beta}} - 1 \right] \times 100\%$. In particular, doubling from n to $2n$ implies $\left[\left(\frac{1+2n}{1+n} \right)^{\hat{\beta}} - 1 \right] \times 100\%$, which approaches $(2^{\hat{\beta}} - 1) \times 100\% \approx 34.4\%$ for large n (95% CI $\approx [31.3\%, 37.6\%]$).

Interpretation. In the model $\log(1+FR_i) = \alpha + \beta \log(1+n_i) + \varepsilon_i$, β is the elasticity of $(1+FR)$ with respect to $(1+n)$: a 1% increase in $(1+n)$ is associated with a 0.427% increase in $(1+FR)$. For finite changes, $(1+FR)$ is multiplied by $\left(\frac{1+n'}{1+n} \right)^{\beta}$; hence the effect of “doubling the group size” depends on the baseline n via $\frac{1+2n}{1+n}$ and converges to about 34.4% for large n . This pattern aligns with the theory: larger groups raise scope for opportunism, but congestion yields diminishing returns to so-called “shirking”.

XII.2 The Production Frontier of Free Riding

Let $\text{Rank} \in \mathbb{N}$ denote swimmer position (lower = stronger). Define the attainable set:

$$\mathcal{F}(\mathbf{r}) = \{\text{FR} \geq 0 : (\text{FR}, \mathbf{r}) \text{ feasible}\}.$$

The *free riding frontier* is the supremum (least upper bound) of the swim rank $\mathbf{r}$.

$$\text{FR}^*(\mathbf{r}) = \sup(\mathcal{F}(\mathbf{r})),$$

analogous to a production frontier. Assume $\text{FR}^*(\mathbf{r})$ is increasing and concave:

$$\frac{\partial \text{FR}^*}{\partial \mathbf{r}} > 0, \quad \frac{\partial^2 \text{FR}^*}{\partial \mathbf{r}^2} < 0.$$

Quantile Approximation. Let $\text{FR} \in [0, \infty)$ denote free riding and $\mathbf{r} \in \mathbb{R}$ the (scalar) rank covariate. Define the conditional cdf $F_{\text{FR}|\mathbf{r}}(\mathbf{y} | \mathbf{r})$ and its (left-continuous) conditional quantile function $Q_\tau(\text{FR} | \mathbf{r}) := \inf\{\mathbf{y} \in \mathbb{R} : F_{\text{FR}|\mathbf{r}}(\mathbf{y} | \mathbf{r}) \geq \tau\}$ for $\tau \in (0, 1)$. We formalize the (upper) *frontier* as the conditional essential supremum

$$\text{FR}^*(\mathbf{r}) := \text{ess sup}\{\mathbf{y} : \mathbb{P}(\text{FR} \leq \mathbf{y} | \mathbf{r}) < 1\} \in [0, \infty],$$

i.e., the smallest $\mathbf{y}$ such that $\mathbb{P}(\text{FR} \leq \mathbf{y} | \mathbf{r}) = 1$.³

Identification via extreme quantiles. If the conditional upper support is bounded (finite $\text{FR}^*(\mathbf{r})$) and $F_{\text{FR}|\mathbf{r}}(\cdot | \mathbf{r})$ is continuous at $\text{FR}^*(\mathbf{r})$ from the left, then

$$\lim_{\tau \uparrow 1} Q_\tau(\text{FR} | \mathbf{r}) = \text{FR}^*(\mathbf{r}).$$

If the support is unbounded but the conditional tail obeys a regular variation (or other EVT) condition, then $Q_\tau(\text{FR} | \mathbf{r})$ is a consistent *frontier proxy* as $\tau \uparrow 1$, with a bias that vanishes under extremal-quantile rates (e.g., $\tau_n = 1 - k_n/n$, $k_n \rightarrow \infty$, $k_n/n \rightarrow 0$).

Approximation. I approximate high conditional quantiles by a linear-in- $\mathbf{r}$ specification

$$Q_\tau(\text{FR} | \mathbf{r}) = \alpha(\tau) + \beta(\tau) \mathbf{r}, \quad \tau \in (0, 1),$$

and take $Q_\tau(\text{FR} | \mathbf{r})$ with *large* τ (e.g., $\tau \in \{0.90, 0.95\}$) as an estimate of $\text{FR}^*(\mathbf{r})$. Estimation uses quantile regression:

$$(\hat{\alpha}(\tau), \hat{\beta}(\tau)) \in \arg \min_{\alpha, \beta} \sum_{i=1}^n \rho_\tau(\text{FR}_i - \alpha - \beta \mathbf{r}_i), \quad \rho_\tau(u) = u(\tau - \mathbf{1}\{u < 0\}).$$

Under standard conditions (i.i.d., conditional quantile linearity, mild regularity), $(\hat{\alpha}(\tau), \hat{\beta}(\tau))$ is consistent and asymptotically normal for fixed $\tau \in (0, 1)$; for extremal $\tau = \tau_n \uparrow 1$, inference requires extremal-quantile theory or subsampling.

Shape and smoothing. Because $\text{FR}^*(\mathbf{r})$ need not be globally linear in $\mathbf{r}$, we also consider a flexible frontier $Q_\tau(\text{FR} | \mathbf{r}) = f_\tau(\mathbf{r})$ with f_τ approximated by B-splines or local polynomials. When theory implies concavity of the frontier, we can impose $f''_\tau(\mathbf{r}) \leq 0$ via shape-constrained (isotonic/concave) quantile regression, ensuring a concave envelope for high- τ quantiles.

3. If the conditional support is unbounded, $\text{FR}^*(\mathbf{r}) = \infty$ and the frontier is interpreted asymptotically; see discussion below.




Figure 36: Estimated conditional quantile frontiers $Q_{\tau}(\text{FR} \mid r)$ at $\tau \in \{0.10, 0.50, 0.90\}$. The $\tau = 0.90$ curve serves as a proxy for the frontier $\text{FR}^*(r)$; higher τ values tighten the approximation, subject to sampling variability.

Results. Figure 36 indicates $Q_{0.90}(\text{FR} \mid r) \approx 0$ for small r (leaders) and increases with r (followers), suggesting $\text{FR}^*(r)$ is (i) near zero at low ranks and (ii) increasing in r on average. Imposing concavity on f_{τ} yields a frontier with $f'_{\tau}(r) \geq 0$ and $f''_{\tau}(r) \leq 0$, i.e., *diminishing marginal frontier effects* of rank on free riding:

$$\frac{\partial}{\partial r} \text{FR}^*(r) \downarrow \quad \text{as } r \text{ rises,} \quad \text{equivalently} \quad \frac{\partial^2}{\partial r^2} \text{FR}^*(r) \leq 0.$$

Economically, higher r expands the scope for opportunism (a larger feasible set), but congestion/monitoring frictions reduce the incremental gain from further rank increases, producing a concave frontier. For policy/counterfactual reading, extremal quantile standard errors (or subsampling) around the $\tau = 0.90$ curve should be reported to quantify the sampling uncertainty of the frontier proxy.

XII.3 Engel Curve

In classical consumer theory, an Engel curve traces Marshallian demand $x(p, m)$ for a good as a function of income m , holding prices p fixed (Deaton 1992). By analogy, let free riding (FR) be the consumption good and group size (n) the analogue of income.

Feasible set. Define the indirect opportunity set for free riding:

$$\mathcal{B}(n) = \{FR \in \mathbb{R}_+ : FR \text{ is feasible when group size is } n\}.$$

We assume $\mathcal{B}(n)$ is nondecreasing in n : if $n' > n$ then $\mathcal{B}(n) \subseteq \mathcal{B}(n')$, reflecting the idea that larger groups allow (weakly) more scope for shirking.

Engel curve analogue. The *Engel curve of free riding* is the typical choice conditional on n , which we represent as the conditional mean function

$$E[FR | n] = f(\log(1 + n)),$$

where the $\log(1 + n)$ transformation ensures diminishing marginal sensitivity to group size. We estimate f nonparametrically, imposing weak regularity:

$$f'(\cdot) > 0 \quad (\text{monotonicity}), \quad f''(\cdot) < 0 \quad (\text{concavity}).$$




Figure 37: Engel-style curve: FR vs. $\log(1 + n)$, across periods (pre, covid, post).

Comparative Statics. Figure 37 suggests the following properties:

(i) *Monotonicity and concavity.* Differentiating the conditional expectation, we obtain

$$\begin{aligned} \frac{\partial}{\partial n} E[FR | n] &= f'(\log(1 + n)) \cdot \frac{1}{1 + n} > 0, \\ \frac{\partial^2}{\partial n^2} E[FR | n] &= f''(\log(1 + n)) \cdot \frac{1}{(1 + n)^2} - f'(\log(1 + n)) \cdot \frac{1}{(1 + n)^2} < 0. \end{aligned}$$

Therefore, FR rises with group size but at a decreasing rate.

(ii) *COVID-19 restrictions as inward shifts.* Let $f_t(\cdot)$ denote the Engel curve in period $t \in \{\text{pre}, \text{covid}, \text{post}\}$. During COVID restrictions, groups are smaller and the conditional mean frontier contracts, so that

$$f_{\text{covid}}(z) < f_{\text{pre}}(z), \quad \forall z.$$

(iii) *Post-period outward shift.* In the recovery period, we find

$$f_{\text{post}}(z) > f_{\text{pre}}(z), \quad \forall z,$$

analogous to an outward shift in a classical Engel curve following income growth.

Interpretation. Group size n acts as a “budget” that constrains the opportunity set for free riding. The Engel curve $f(\log(1 + n))$ captures how the intensity of free riding adjusts with available opportunities, in direct analogy to how household consumption responds to income. Concavity reflects diminishing marginal “returns” to group size in generating free riding.

XII.4 Beveridge Curve Analogy: Group Size and Slack

In micro-/labor economics, the Beveridge style curve relation arises from a matching function $M(u, v)$ that links unemployed workers u and vacancies v ; the Cobb–Douglas form is a standard benchmark (Pissarides 2000; Petrongolo and Pissarides 2001). We adapt this structure to swim groups by defining *slack* S as effective rank improvement (drafting gains) created by pairwise/trio “matches” among swimmers.

Matching Function Analogy. Let n be the analogue of vacancies and the individual rank disadvantage $r \geq 0$ the analogue of unemployment. We posit a Cobb–Douglas matching technology

$$M(r, n) = A r^\alpha n^{1-\alpha}, \quad \alpha \in (0, 1), A > 0,$$

where A summarizes matching efficiency (environment, coordination protocols). Slack for a swimmer with disadvantage r is proportional to matching output:

$$S(r, n) = \phi M(r, n) = \phi A r^\alpha n^{1-\alpha}, \quad \phi > 0.$$

Aggregation. Let r have density $f(r)$ on $[0, \bar{r}]$ with finite $\mathbb{E}[r^\alpha]$. The mean slack is

$$\bar{S}(n) = \phi A n^{1-\alpha} \int_0^{\bar{r}} r^\alpha f(r) dr = \underbrace{(\phi A \mathbb{E}[r^\alpha])}_{:= \kappa_0} n^{1-\alpha},$$

so, holding the r -distribution fixed, $\bar{S}(n) \propto n^{1-\alpha}$, a standard implication of Cobb–Douglas matching under constant returns in (r, n) (Petrongolo and Pissarides 2001; Acemoglu 2020).

Congestion/Coordination Costs. To capture congestion at high n , subtract a cost $C(n)$ that is increasing and convex:

$$C'(n) > 0, \quad C''(n) \geq 0.$$

Define effective (average) slack:

$$S(n) = \bar{S}(n) - C(n) = \kappa_0 n^{1-\alpha} - C(n).$$

Curvature, Elasticities, and the Camel-Hump. First derivative and elasticity with respect to n :

$$S'(n) = (1 - \alpha)\kappa_0 n^{-\alpha} - C'(n), \quad \eta_{S,n}(n) \equiv \frac{\partial \log S(n)}{\partial \log n} = \frac{nS'(n)}{S(n)}.$$

Second derivative:

$$S''(n) = -\alpha(1 - \alpha)\kappa_0 n^{-\alpha-1} - C''(n) < 0 \quad \text{if } C''(n) > 0,$$

so congestion steepens concavity relative to the frictionless $n^{1-\alpha}$ frontier.

Inverted-U (Beveridge-shape). If there exists n^* such that $S'(n^*) = 0$ and $S''(n^*) < 0$, then S is increasing for $n < n^*$ and decreasing for $n > n^*$:

$$\underbrace{(1 - \alpha)\kappa_0 (n^*)^{-\alpha}}_{\text{marginal match benefit}} = \underbrace{C'(n^*)}_{\text{marginal congestion cost}}, \quad S''(n^*) < 0.$$

For a power congestion cost $C(n) = \kappa_1 n^\gamma$ with $\kappa_1 > 0$ and $\gamma > 1$,

$$n^* = \left(\frac{(1-\alpha)\kappa_0}{\kappa_1\gamma} \right)^{\frac{1}{\gamma-1}}, \quad S(n^*) = \kappa_0(n^*)^{1-\alpha} - \kappa_1(n^*)^\gamma.$$

The hump exists whenever $\gamma > 1$ (strictly convex C) and parameters yield $S(n^*) > 0$.

Comparative Statics on the Peak. From the closed form above:

$$\frac{\partial n^*}{\partial \kappa_0} > 0, \quad \frac{\partial n^*}{\partial \kappa_1} < 0, \quad \frac{\partial n^*}{\partial \alpha} < 0.$$

Thus, higher matching efficiency (A or ϕ) or more dispersion in ability (higher $\mathbb{E}[r^\alpha]$) raise the optimal group size; stronger congestion (κ_1, γ) or a higher weight on r (larger α) reduce it. Changes in A (matching efficiency) shift the *entire* S curve, echoing “efficiency” shifts of the Beveridge relation in labor markets (Barnichon and Figura 2010; Blanchard and Diamond 1989).

Alternative Functional Forms and Robustness. A flexible congestion term such as $C(n) = \kappa_1 n \log n$ (weakly convex) yields

$$S'(n) = (1-\alpha)\kappa_0 n^{-\alpha} - \kappa_1(1 + \log n),$$

again generating a unique interior n^* when parameters imply an intersection. More generally, any C with $C'(n)$ eventually dominating $n^{-\alpha}$ produces an inverted-U. Endogenizing A (e.g., coordination protocols or signaling) produces outward/inward shifts akin to Beveridge-curve “efficiency” shifts (Petrongolo and Pissarides 2001; Lubik 2013).

Estimation and Tests. A parsimonious parametric specification suitable for estimation is

$$S_i = \theta_0 n_i^{1-\alpha} - \kappa_1 n_i^\gamma + \varepsilon_i,$$

with (α, γ) identified from curvature and the interior maximum.⁴ Alternatively, a semi-nonparametric sieve (e.g., B-splines in $\log n$) with the shape restrictions $S'(n) \gtrless 0$ before/after n^* and $S''(n) \leq 0$ can recover the hump while remaining agnostic about C .

Micro–Macro Link. At the micro level, mean slack decomposes as $S(n) = \phi A \mathbb{E}[r^\alpha] n^{1-\alpha} - C(n)$, so changes in the *distribution* of ranks (e.g., training that compresses r) shift $\mathbb{E}[r^\alpha]$ and therefore the entire curve. This mirrors the way composition and efficiency shift empirical Beveridge curves in labor markets (Petrongolo and Pissarides 2001; Barnichon and Figura 2010).

4. Identification improves by pooling across periods with common (α, γ) and period-specific efficiencies $\theta_{0,t}$, capturing Beveridge-style *shifts* (Barnichon and Figura 2010).




Figure 38: Beveridge-style curve: group size n vs. slack S . Rising n initially expands matching opportunities ($\propto n^{1-\alpha}$) but eventually raises coordination costs $C(n)$; the interior maximum n^* solves $(1 - \alpha)\kappa_0 n^{-\alpha} = C'(n)$.

Interpretation. As in search-and-matching models, both scarcity (small n) and congestion (large n) reduce effective matches. Slack $S(n)$ is maximized at the interior n^* where marginal benefits equal marginal congestion costs. Exogenous shifts in matching efficiency (A) or composition ($\mathbb{E}[r^\alpha]$) move the entire curve (Beveridge-style shifts), while changes in C rotate/twist the curve by altering curvature (Blanchard and Diamond 1989; Barnichon and Figura 2010).

XII.5 Growth Incidence of Free Riding

Following Ravallion and Chen (2003) and Ravallion (2004), the growth incidence curve (GIC) describes quantile-specific distributional change between two periods $t \in \{\text{Pre}, \text{Post}\}$. Let F_t be the cdf of FR in period t , and define the (left-continuous) unconditional quantile function

$$Q_t(p) := \inf\{y \in \mathbb{R} : F_t(y) \geq p\}, \quad p \in (0, 1).$$

The *absolute* and *relative* GIC are, respectively,

$$\Delta Q(p) = Q_{\text{Post}}(p) - Q_{\text{Pre}}(p), \quad g(p) = \frac{Q_{\text{Post}}(p) - Q_{\text{Pre}}(p)}{Q_{\text{Pre}}(p)} = \frac{Q_{\text{Post}}(p)}{Q_{\text{Pre}}(p)} - 1.$$

When zeros are prevalent (as with FR), a *log-GIC* avoids division by zero:

$$G_{\log}(p; c) = \log(Q_{\text{Post}}(p) + c) - \log(Q_{\text{Pre}}(p) + c),$$

with $c \geq 0$ (e.g., $c = 1$ to mirror the $\log(1 + \text{FR})$ transform). For small changes, $G_{\log}(p; c) \approx g(p)$.

Quantile-density calculus. Let f_t be the density of FR in period t and $q_t(p) := \partial_p Q_t(p) = 1/f_t(Q_t(p))$ the *quantile density*. If f_t is positive and continuous at $Q_t(p)$, then

$$\Delta Q'(p) = q_{\text{Post}}(p) - q_{\text{Pre}}(p), \quad g'(p) = \frac{q_{\text{Post}}(p)Q_{\text{Pre}}(p) - q_{\text{Pre}}(p)Q_{\text{Post}}(p)}{(Q_{\text{Pre}}(p))^2}.$$

Hence, “top-heavy” growth on $(p_0, 1)$ obtains if $g'(p) > 0$ there; sufficient conditions include $q_{\text{Post}}(p)/q_{\text{Pre}}(p) > Q_{\text{Post}}(p)/Q_{\text{Pre}}(p)$ for $p \in (p_0, 1)$.




Figure 39: Growth incidence curve: quantile-specific change in free riding (Post – Pre). Solid: $g(p)$; shading: simultaneous 95% band from the multiplier bootstrap.

Estimation and inference. Let $\hat{Q}_t$ be sample quantiles (with monotone rearrangement to avoid crossings if needed). Pointwise, $\sqrt{n_t}(\hat{Q}_t(p) - Q_t(p)) \Rightarrow \mathcal{N}(0, p(1-p)/f_t(Q_t(p))^2)$.

For the entire curve $p \in \mathcal{P} \subset (0, 1)$, simultaneous bands for Q_t , ΔQ , g , or $G_{\log}$ follow from functional CLTs; valid uniform inference can be obtained by the (weighted) bootstrap or multiplier bootstrap over p .⁵ Composition-adjusted GICs (“counterfactual GICs”)—holding covariates fixed across periods—can be constructed either by reweighting or via RIF-regressions for unconditional quantiles.⁶ Top-heaviness can be tested using a supremum-type test of $H_0 : g'(p) \leq 0$ on $p \in [p_0, 1)$ or by testing integral contrasts $\int_{p_0}^1 g(p) dp > 0$ with uniform bands.⁷

Result. Figure 39 shows $g(p) \approx 0$ over most quantiles but $g(p) \gg 0$ in the upper tail ($p > 0.90$). Formally, we find

$$\frac{\partial g(p)}{\partial p} > 0 \quad \text{for } p \in (0.9, 1),$$

with the null of flat upper-tail growth rejected by the uniform band. Thus, distributional change is *top-heavy*: post-period gains in free riding accrue disproportionately to those already at high levels.

Interpretation and Summary

The formal analysis yields several key results:

1. **Elasticity:** Free riding responds positively but inelastically to group size (log-log estimate $\hat{\varepsilon} \approx 0.43$).
2. **Frontier:** The free-riding production frontier $FR^*(r)$ is increasing and concave in rank r .
3. **Engel Analogy:** FR follows a concave Engel-style curve in group size, saturating as n grows, with downward (COVID) and upward (post) shifts.
4. **Beveridge Analogy:** Slack $S(n)$ is hump-shaped:

$$\exists n^* : \frac{\partial S}{\partial n} > 0 \ (n < n^*), \quad \frac{\partial S}{\partial n} < 0 \ (n > n^*).$$

Medium groups maximize slack; very small or very large groups underperform.

5. **Inequality:** Overall dispersion (Theil-T) fell during COVID due to restricted opportunities, then rose above pre-period levels, consistent with the decomposability and sensitivity of entropy indices to upper-tail changes (Theil 1967; Cowell 2005).
6. **Growth Incidence:** Distributional change is top-heavy:

$$g(p) \approx 0 \ (p < 0.9), \quad g(p) \gg 0 \ (p > 0.9),$$

with upper-tail increases statistically significant under uniform inference.

5. See Koenker (2005) for quantile estimation; Chernozhukov, Fernández-Val, and Melly (2013) for joint inference on distribution/quantile functionals and counterfactuals.

6. See S. Firpo, Fortin, and Lemieux (2009) and S. P. Firpo, Fortin, and Lemieux (2018).

7. On functional hypotheses over distributions (e.g., stochastic dominance), see Davidson and Duclos (2000).

XIII. Discussion

This paper studies free riding in triathlon using complementary empirical strategies that exploit sharp institutional rule-based variation, panel data structure, and distributional decomposition tools. Across all designs, the evidence points to a common pattern: drafting-related rank gains contract sharply in 2020 and recover only partially afterward, with substantial heterogeneity across athletes and race environments.

The regression discontinuity (RDD) at the COVID (2020) cutoff provides a design-based estimate that is difficult to reconcile with smooth trends alone. Local linear estimates indicate a drop in rank improvement of about -31 points (SE 17.3), while local cubic specifications yield a much larger effect of roughly -92 points (SE 30.8). The magnitude difference across polynomial orders reflects curvature near the cutoff, yet both specifications point in the same direction. These results match the institutional reality of stricter non-drafting enforcement and altered race formats in 2020.

The identification strategy based on rank-improvement probabilities reinforces this conclusion without relying on linear functional forms. The proportion of athletes improving their rank is consistently lower in 2020 than in pre- and post-periods across all top-K thresholds (namely: TOP-5, TOP-10, TOP-20, TOP-30). Under monotonicity and within-rank-group independence, the difference in proportions identifies the average causal effect for compliers. The fact that $\hat{p}_1 > \hat{p}_0$ throughout implies that non-COVID conditions raise the probability of rank improvement for athletes whose outcomes respond to the competitive environment.

Panel regressions with athlete and event-category fixed effects connect these probability shifts to changes in conditional means (*c. p.* analysis). Event-study coefficients show positive deviations before 2020, normalization by construction in 2020, and a selective rebound afterward. The pooled period specification estimates a pre-2020 main effect of $+13.43$ (95% CI $[0.18, 26.67]$), while the post-period main effect is smaller and imprecise. These average patterns mask clear heterogeneity. Medium and large packs increase rank improvement by about $+18.0$ and $+31.6$ points, respectively, while the quadratic age term ($\hat{\gamma}_2 \approx 0.035$) implies a convex age profile within athlete.

Disaggregated results sharpen the economic interpretation. Weaker swimmers, later drafter positions, and mid-age groups experience the largest contractions in 2020 and the clearest rebounds afterward. Male-heavy fields and thinner packs dampen post-2020 recovery, while back-of-pack positions show higher free riding once restrictions ease. Distributional tools reach the same conclusion from another angle: upper quantiles of free riding shift inward during COVID and expand afterward, producing top-heavy growth in the recovery period. The growth incidence curve shows near-zero change below the 90th percentile and pronounced gains above it.

Theoretical exercises in the final sections help organize these findings. The public-good model links larger packs to stronger incentives to shirk. The production-frontier and quantile analysis show that feasible free riding rises with rank disadvantage but at a decreasing rate. Engel-curve and Beveridge-style analogies describe how group size creates opportunities up to a point, after which congestion and coordination costs dominate. Across these frameworks, the COVID shock acts as a contraction of the feasible set, with recovery restoring opportunities unevenly rather than uniformly.

Methodologically, the paper shows that high-dimensional fixed effects and large panels need not limit the empirical scope. Grouped-Gram estimation and uncentered Gram variance formulas allow exact OLS estimation and inference while shifting computational burden to cross-product formation. This matters for designs with many fixed effects, rich interactions, and long panels such as those used here, but also for much larger panels with the use of modern notebooks/without the use of supercomputing.

XIV. Conclusion

This study provides causal and descriptive evidence on free riding in triathlon, focusing on how drafting-related rank gains respond to institutional constraints. A sharp break at the 2019/2020 boundary reveals a large and statistically meaningful decline in rank improvement during COVID. Differences in rank-improvement probabilities confirm the same pattern under weak assumptions and identify positive effects for compliers in non-COVID periods. Fixed-effects event studies trace the time path of these effects and document substantial heterogeneity across age, ability, pack size, team composition, and drafter position.

The results point to a clear economic mechanism. Drafting opportunities expand with group size and positional disadvantage, while restrictions compress the opportunity set and disproportionately affect athletes who benefit most from pack dynamics. Post-2020 rebounds recovers these opportunities only partially and unevenly, leading to top-heavy gains rather than broad-based normalization.

Beyond the application, the paper demonstrates a unified empirical workflow that combines regression discontinuity, panel methods, probability-based identification, and distributional analysis, supported by computational techniques that scale to high-dimensional fixed effects. The approach can be applied to other settings where individual outcomes depend on group structure, enforcement regimes, or congestion.

Several limitations remain. The analysis focuses on observed race outcomes rather than direct measures of effort or drafting behavior. External validity outside triathlon depends on how closely other endurance sports share the same institutional and strategic features. Future work could combine tracking data, physiological measures, or experimental enforcement changes to refine the mechanism further.

Overall, the evidence shows that free riding in competitive endurance settings responds sharply to institutional constraints, with persistent heterogeneity across athletes and race environments. Understanding these responses helps clarify how rules, group structure, and incentives (e.g. such as time penalties) shape performance in settings where individual outcomes depend on collective group behavior.

References

- Acemoglu, Daron. 2020. *Search, matching, unemployment: Lecture notes*. Lecture slides. Massachusetts Institute of Technology.
- Angrist, Joshua D. 2014. “The perils of peer effects.” *Labour Economics* 30:98–108. <https://doi.org/10.1016/j.labeco.2014.05.008>.
- Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics*. Princeton University Press.
- Barnichon, Régis, and Andrew Figura. 2010. *What drives movements in the unemployment rate? A decomposition of the Beveridge curve*. Finance and Economics Discussion Series 2010-48. Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/pubs/feds/2010/201048/201048pap.pdf>.
- Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1996. “The all-pay auction with complete information.” *Economic Theory* 8 (2): 291–305.
- Bellemare, Marc F., and Casey J. Wichman. 2020. “Elasticities and the inverse hyperbolic sine transformation.” *Oxford Bulletin of Economics and Statistics* 82 (1): 50–61. <https://doi.org/10.1111/obes.12325>.
- Blanchard, Olivier Jean, and Peter Diamond. 1989. “The Beveridge curve.” *Brookings Papers on Economic Activity* 1989 (1): 1–76. <https://doi.org/10.2307/2534495>. <https://www.brookings.edu/articles/the-beveridge-curve/>.
- Bolon, Benoit, Cyril Pretot, Christophe Clanet, Fabien Larrarte, and Remy Carmigniani. 2020. “Drafting of two passive swimmer scale models for open-water races.” *Soft Matter* 16 (31): 7270–7273.
- Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. “Identification of peer effects through social networks.” *Journal of Econometrics* 150 (1): 41–55.
- Chatard, Jean-Claude, Jean-Marc Lavoie, Bernard Bourgoin, and Jean-Robert Lacour. 1990. “The contribution of passive drag as a determinant of swimming performance.” *International Journal of Sports Medicine* 11 (4): 367–372.
- Chatard, Jean-Claude, and Barry Wilson. 2003. “Drafting distance in swimming.” *Medicine & Science in Sports & Exercise* 35 (7): 1176–1180. <https://doi.org/10.1249/01.MSS.0000074564.06106.1F>.
- Che, Yeon-Koo, and Ian L. Gale. 2003. “Optimal design of research contests.” *American Economic Review* 93 (3): 646–671.
- Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly. 2013. “Inference on counterfactual distributions.” *Econometrica* 81 (6): 2205–2268. <https://doi.org/10.3982/ECTA10582>.
- Corchón Díaz, Luis Carlos, and Marco Serena. 2016. “Contest theory: A survey.” Forthcoming, *Handbook of Game Theory and Industrial Organization*, <https://doi.org/10.2139/ssrn.2811686>. <https://ssrn.com/abstract=2811686>.
- Cowell, Frank A. 2005. *Theil, inequality indices and decomposition*. Working Paper 2005-01. ECINEQ. <https://www.ecineq.org/milano/WP/ECINEQ2005-01.pdf>.
- Davidson, Russell, and Jean-Yves Duclos. 2000. “Statistical inference for stochastic dominance and for the measurement of poverty and inequality.” *Econometrica* 68 (6): 1435–1464. <https://doi.org/10.1111/1468-0262.00167>.

- De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri. 2020. "Consumption network effects." *Review of Economic Studies* 87 (1): 130–163. <https://doi.org/10.1093/restud/rdy039>.
- Deaton, Angus. 1992. *Understanding consumption*. Oxford: Oxford University Press. ISBN: 9780198288244. <https://ideas.repec.org/b/oxp/obooks/9780198288244.html>.
- Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional quantile regressions." *Econometrica* 77 (3): 953–973. <https://doi.org/10.3982/ECTA6822>.
- Firpo, Sergio P., Nicole M. Fortin, and Thomas Lemieux. 2018. "Decomposing wage distributions using recentered influence function regressions." *Econometrics* 6 (2): 28. <https://doi.org/10.3390/econometrics6020028>.
- Gürtler, Oliver, Lennart Struth, and Max Thon. 2023. "Competition and risk-taking." *European Economic Review* 160:104592. <https://doi.org/10.1016/j.euroecorev.2023.104592>.
- Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo. 2009. "Peer effects in the workplace: Evidence from random groupings in professional golf tournaments." *American Economic Journal: Applied Economics* 1 (4): 34–68. <https://doi.org/10.1257/app.1.4.34>.
- Hillman, Arye L., and John G. Riley. 1989. "Politically Contestable Rents and Transfers." *Economics and Politics* 1, no. 1 (March): 17–39.
- Huntington-Klein, Nick. 2023. "Linear rescaling to accurately interpret logarithms." *Journal of Econometric Methods* 12 (1): 139–147. <https://doi.org/10.1515/jem-2021-0029>.
- Imbens, Guido W., and Donald B. Rubin. 2015. *Causal inference for statistics, social, and biomedical sciences: An introduction*. New York: Cambridge University Press. ISBN: 9781107084523. <https://doi.org/10.1017/CBO9781139025751>.
- Koenker, Roger. 2005. *Quantile regression*. Cambridge: Cambridge University Press. ISBN: 9780521845731. <https://www.cambridge.org/core/books/quantile-regression/C18AE7BCF3EC43C16937390D44A328B1>.
- Konrad, Kai A. 2009. *Strategy and Dynamics in Contests*. Oxford: Oxford University Press. ISBN: 9780199549009.
- Krueger, Anne O. 1974. "The political economy of the rent-seeking society." *American Economic Review* 64 (3): 291–303.
- Lackner, Mario, and Hendrik Sonnabend. 2025. *Positive Feedback Shapes Gender Gaps in Adolescent Risk-Taking – Causal Evidence from Real-Risk Competitions*. Economics Working Papers 2025–10. Department of Economics, Johannes Kepler University Linz.
- Leitner, Stefan. 2021. *Massenstart oder Einzelstart?* Accessed February 12, 2025. <https://www.trinews.at/massenstart-oder-einzelstart/>.
- Lubik, Thomas A. 2013. *The shifting and twisting Beveridge curve: An aggregate perspective*. Working Paper 13-16. Federal Reserve Bank of Richmond. https://www.richmondfed.org/publications/research/working_papers/2013/wp_13-16.
- Manski, Charles F. 1993. "Identification of endogenous social effects: The reflection problem." *Review of Economic Studies* 60 (3): 531–542. <https://doi.org/10.2307/2298123>.
- Matthes, Julian, and David Piazzolo. 2024. "Don't put all your legs in one basket: Theory and evidence on coopetition in road cycling." *European Economic Review* 170:104892. <https://doi.org/10.1016/j.euroecorev.2024.104892>.
- Milgrom, Paul, and Chris Shannon. 1994. "Monotone comparative statics." *Econometrica* 62 (1): 157–180. <https://doi.org/10.2307/2951479>.

- Millet, Grégoire, Jean-Claude Chatard, and Didier Chollet. 1999. “Effects of kicking, performance, and drag during draft swimming in elite triathletes.” In *Les Cahiers de l’INSEP*, 209–216. 24. Proceedings of the 2nd INSEP International Triathlon Congress.
- Modak Chowdhury, Subhasish M., and Roman M. Sheremeta. 2011. “Multiple equilibria in Tullock contests.” *Economics Letters* 112 (2): 216–219. <https://doi.org/10.1016/j.econlet.2011.05.002>.
- Moldovanu, Benny, and Aner Sela. 2001. “Optimal allocation of prizes in contests.” *American Economic Review* 91 (3): 542–558.
- Nishihata, Masaya. 2022. “Heterogeneous peer effects by gender, task, and monetary incentive: Evidence from speed skating.” *Southern Economic Journal* 89 (1): 62–89. <https://doi.org/10.1002/soej.12585>.
- Ohmichi, Hiroshi, Masayuki Takamoto, and Mitsuo Miyashita. 1983. “Measurement of the waves caused by swimmers.” In *Biomechanics and Medicine in Swimming V*, edited by Arend P. Hollander, Peter A. Huijing, and Gert de Groot, 103–107. Champaign, IL: Human Kinetics.
- Olbrecht, Jan. 2011. “Triathlon: Swimming for winning.” *Journal of Human Sport and Exercise* 6 (2): 233–246.
- Österreichischer Triathlonverband (ÖTRV). 2020. *ALOHA TRI Steyregg feiert Premiere*. Accessed February 12, 2025. <https://www.triathlon-austria.at/de/newsshow-alle----aloha-tri-steyregg-feiert-premiere?return=240>.
- Pearl, Judea. 2009. *Causality: Models, reasoning, and inference*. 2nd ed. Cambridge: Cambridge University Press. ISBN: 9780521895606.
- Petrongolo, Barbara, and Christopher A. Pissarides. 2001. “Looking into the black box: A survey of the matching function.” *Journal of Economic Literature* 39 (2): 390–431. <https://doi.org/10.1257/jel.39.2.390>.
- Pissarides, Christopher A. 2000. *Equilibrium unemployment theory*. 2nd ed. Cambridge, MA: MIT Press. ISBN: 9780262161879. <https://mitpress.mit.edu/9780262161879/equilibrium-unemployment-theory/>.
- Ravallion, Martin. 2004. *Pro-poor growth: A primer*. Policy Research Working Paper 3242. World Bank. <https://documents1.worldbank.org/curated/en/358321468761705849/pdf/wps3242growth.pdf>.
- Ravallion, Martin, and Shaohua Chen. 2003. “Measuring pro-poor growth.” *Economics Letters* 78 (1): 93–99. [https://doi.org/10.1016/S0165-1765\(02\)00205-7](https://doi.org/10.1016/S0165-1765(02)00205-7).
- Reichel, Felix. 2025a. “High-Performance Variance–Covariance Matrix Construction Using an Uncentered Gram Formulation.” Revised version, *arXiv preprint arXiv:2511.08223* (December). <https://arxiv.org/abs/2511.08223>.
- . 2025b. *Strategic effort and bandwagon effects in finite multi-stage games with non-linear externalities: Evidence from triathlon*. arXiv: 2505.03247 [econ.GN].
- . 2025c. *Two-Stage Asymmetric Tullock Contests with Cost Shifters and Endogenous Continuation Decision*. arXiv: 2510.00349 [econ.GN]. <https://arxiv.org/abs/2510.00349>.
- . 2025d. “Using Covid-19 response policy to estimate open water swim drafting.” *Scientific Journal of Sport and Performance* 4 (3): 443–453.
- . 2026. “A Novel Grouped-Gram-Based Algorithm for Fast and Memory-Efficient Fixed Effects Estimation.” *Journal of Engineering and Artificial Intelligence* 2 (1): 1–6. <https://doi.org/10.64142/jelai.2.1.40>. <https://doi.org/10.64142/jelai.2.1.40>.

- Rubin, Donald B. 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology* 66 (5): 688–701. <https://doi.org/10.1037/h0037350>.
- Rybin, Dmitry, Yushun Zhang, and Zhi-Quan Luo. 2025. XX^t Can Be Faster. arXiv: [2505.09814](https://arxiv.org/abs/2505.09814) [cs.DS]. <https://arxiv.org/abs/2505.09814>.
- Sakalidis, Kandianos, Gerrit Paul Stein, Marije Titia Menting, Marije Titia Elferink-Gemser, and Florentina Johanna Hettinga. 2022. "The Role of the Social Environment in Pacing and Sports Performance: A Narrative Review from a Self-Regulatory Perspective." *International Journal of Environmental Research and Public Health* 19 (23): 16131. <https://doi.org/10.3390/ijerph192316131>.
- Santos Silva, J. M. C., and Silvana Tenreyro. 2006. "The log of gravity." *The Review of Economics and Statistics* 88 (4): 641–658. <https://doi.org/10.1162/rest.88.4.641>.
- Skaperdas, Stergios. 1996. "Contest success functions." *Economic Theory* 7 (2): 283–290.
- Song, Jian, and Daniel Houser. 2021. "Non-exclusive group contests: An experimental analysis." *Journal of Economic Psychology* 87:102444. <https://doi.org/10.1016/j.joep.2021.102444>.
- Szidarovszky, Ferenc, and Koji Okuguchi. 1997. "On the existence and uniqueness of pure Nash equilibrium in rent-seeking games." *Games and Economic Behavior* 18 (1): 135–140.
- Tadelis, Steven. 2013. *Game theory: An introduction*. Princeton University Press.
- Theil, Henri. 1967. *Economics and information theory*. Amsterdam: North-Holland. ISBN: 9780444102829.
- Toussaint, Huub M., Renée Bruinink, Ronald Coster, et al. 1989. "Effect of triathlon wet suit on drag during swimming." *Medicine & Science in Sports & Exercise* 21 (3): 325–328.
- Tullock, Gordon. 1980. *Efficient rent seeking*. College Station: Texas A&M University Press.
- Vives, Xavier. 2005. "Strategic complementarity in multi-stage games." *IESE Research Papers*, no. 58, <https://blog.iese.edu/xvives/files/2011/06/58-Strategic-Complementarity-in-Multistage-Games.pdf>.
- Wooldridge, Jeffrey M. 2010. *Econometric analysis of cross section and panel data*. 2nd ed. Cambridge, MA: MIT Press. ISBN: 9780262232586. <https://mitpress.mit.edu/9780262232586/econometric-analysis-of-cross-section-and-panel-data/>.

Appendix / Supplementary Materials

A.1. Data Dictionary

This appendix lists (A) variables read from *data/preprocessed.csv* and (B) variables computed by the analysis pipeline. “Level” indicates the unit at which the variable is defined (event, athlete, or athlete×event).

A.1.1 Original variables (from *data/preprocessed.csv*)

Variable	Type	Unit/Scale	Level	Description
X	integer	–	row	Source row/index (not used analytically).
athlete_id	string	–	athlete	Pseudonymous athlete identifier.
event_id	integer	–	event	Event identifier (race instance).
total	integer	seconds or rank	athlete×event	Overall finish time or overall rank as provided.
swimming	integer	seconds	athlete×event	Swim split time (lower is faster).
rank	integer	rank	athlete×event	Overall finishing position (1 = best).
male	integer (0/1)	–	athlete×event	Sex indicator (1 = male, 0 = female).
year	integer	YYYY	event	Calendar year (later overwritten from date).
name	string	–	event	Event name.
date	date	ISO	event	Event date (used to derive year).
event_category	factor	category	event	Race category (e.g., sprint/OLY/70.3, etc.).
location	string	–	event	Event location.
period	factor	category	athlete×event	Period label from source.
covid	integer (0/1)	–	athlete×event	Indicator as in source.
post_covid	integer (0/1)	–	athlete×event	Indicator as in source.
week_running	integer	index	event	Running week index or week-of-year from source.
age	integer	years	athlete×event	Athlete age at event.
age_squared	integer	years ²	athlete×event	Age squared from source.
eventYear	integer	YYYY	event	Event year (duplicate of year in many cases).
eventYearMonth	string	YYYY–MM	event	Year–month of event date.
cluster	integer	1 . . K	athlete×event	Source cluster tag (not the k-means below).
leader	integer (0/1)	–	athlete×event	Leader indicator (1 = leading swimmer/pack leader).
drafter	integer (0/1)	–	athlete×event	Drafter indicator (bike).

(cont’d on next page)

Variable	Type	Unit/Scale	Level	Description
drafter_position	integer	ordinal	athlete×event	Within-group drafter position (1 = closest).
first_drafter	integer (0/1)	–	athlete×event	Indicator as named in source.
second_drafter	integer (0/1)	–	athlete×event	Indicator as named in source.
third_drafter	integer (0/1)	–	athlete×event	Indicator as named in source.
fourth_drafter	integer (0/1)	–	athlete×event	Indicator as named in source.
fift(h)_drafter	integer (0/1)	–	athlete×event	Indicator as named in source (typo preserved).
last_drafter	integer (0/1)	–	athlete×event	Indicator as named in source.

A.1.2 Computed variables (analysis pipeline)

Variable	Type	Unit/Scale	Level	Construction / Notes
year (overwritten)	integer	YYYY	event	Recomputed as year(date). Restrict years to 2015–2024.
<i>Sample restriction</i>	–	–	dataset	Keep athletes with at least one 2020 appearance; drop rows with rank > 500, leader = 1, drafter_position > 20.
swim_rank	numeric	rank (ties averaged)	athlete×event	Within each event_id: rank(swimming) with averaged ties.
rank_improvement	numeric	ranks	athlete×event	rank – swim_rank. Positive = worse overall than swim-only rank.
period_group	factor	{pre, 2020, post}	athlete×event	pre: 2015–2019; 2020; post: 2021–2024.
age_bin	factor	5-year bins	athlete×event	Bins over observed range using cut (include lowest).
group_size	integer	count	event	Number of unique athletes per event_id.
group_size_bin	factor	{small, medium, large}	event	small < 5; medium 5–10; large > 10.
best_swim	numeric	seconds	athlete×category	For each athlete×event_category: min(swimming).
cluster_id	integer	1–3	athlete	k-means on scaled best_swim across categories (k = 3, nstart = 25).
swim_skill	factor	{good, average, bad}	athlete	Labels for cluster_id ordered by mean time (fastest = good).
free_riding	numeric	ranks	athlete×event	$\frac{\text{rank_improvement}}{\text{rank_improvement}_{(\text{year}=2020)}}$
gender	factor	{Male, Female, Unknown}	athlete×event	From male: 1→Male, 0→Female, else Unknown.
<i>Aggregated outputs (separate summary tables):</i>				
fr_gender.mean_fr	numeric	ranks	year×gender	Mean of free_riding by (year, gender).
fr_gsize.mean_fr	numeric	ranks	year×group_size_bin	Mean of free_riding by (year, group_size_bin).
fr_age.mean_fr	numeric	ranks	year×age_bin	Mean of free_riding by (year, age_bin).

Files written. *athlete_swim_clusters.csv* (athlete_id, cluster_id, swim_skill); *stratified_working_dataset.csv* (row-level working data with computed variables); *free_riding_by_year_gender.csv*, *free_riding_by_year_groupsize.csv*, *free_riding_by_year_agebin.csv*; *t_tests_rank_improvement.csv*.

A.2. Non-parametric Construction of Swimmer-Ability Labels (K-means Clustering)

This appendix documents the construction of the swimmer-ability labels *good* / *average* / *bad* used throughout the analysis.

Feature matrix. For athlete $\mathbf{a}$ and event category $\mathbf{c} \in \{\text{Sprint, Kurz, Mittel, Lang}\}$, define $x_{\mathbf{ac}}$ as the best (minimum) recorded swim-out time in seconds. Missing values of $x_{\mathbf{ac}}$ are imputed by the median of the corresponding category, and each category column is standardized. Specifically, let

$$z_{\mathbf{ac}} = \frac{x_{\mathbf{ac}} - \text{median}(x_{\bullet\mathbf{c}})}{\text{sd}(x_{\bullet\mathbf{c}})}, \quad z_{\mathbf{a}} = (z_{\mathbf{a}1}, \dots, z_{\mathbf{a}C}),$$

so that each athlete is represented by a standardized C -dimensional feature vector $z_{\mathbf{a}}$.

Clustering algorithm. We apply k-means clustering for the non-parametric estimation of principal points with $k = 3$ clusters, using $n_{\text{start}} = 25$ random initializations and retaining the solution with minimal within-cluster dispersion. The estimator solves

$$\min_{\{\mu_j\}_{j=1}^k, c(\cdot)} \sum_{\mathbf{a}=1}^A \|z_{\mathbf{a}} - \mu_{c(\mathbf{a})}\|_2^2, \quad c(\mathbf{a}) \in \{1, \dots, k\},$$

where μ_j is the centroid of cluster j , and $c(\mathbf{a})$ assigns athlete $\mathbf{a}$ to a cluster.

Choice of k . To determine the number of clusters, we evaluate three diagnostics:

1. *Elbow method:* Plot $W_k = \sum_{\mathbf{a}} \|z_{\mathbf{a}} - \mu_{c(\mathbf{a})}\|_2^2$ against k and identify a point of diminishing returns.
2. *Silhouette score:* For athlete $\mathbf{a}$, compute

$$s_{\mathbf{a}} = \frac{b_{\mathbf{a}} - a_{\mathbf{a}}}{\max\{a_{\mathbf{a}}, b_{\mathbf{a}}\}},$$

where $a_{\mathbf{a}}$ is the average distance to athletes in the same cluster and $b_{\mathbf{a}}$ is the minimum average distance to athletes in another cluster. The diagnostic uses $\bar{s}(k) = A^{-1} \sum_{\mathbf{a}} s_{\mathbf{a}}$.

3. *Gap statistic:*

$$\text{Gap}(k) = \mathbb{E}^*[\log W_k^*] - \log W_k,$$

where W_k^* is the within-cluster dispersion under a reference distribution and the expectation is approximated by simulation. We adopt the one-standard-error rule, choosing the smallest k such that

$$\text{Gap}(k) \geq \text{Gap}(k+1) - s_{k+1}.$$

All three diagnostics favor $k = 3$.

Labeling clusters by ability. After clustering, we order the clusters by mean unscaled swim time. For cluster j with membership set C_j , define

$$\theta_j = \frac{1}{|C_j|} \sum_{\mathbf{a} \in C_j} \frac{1}{C} \sum_{\mathbf{c}} x_{\mathbf{ac}}.$$

The cluster with the smallest θ_j is labeled *good*, the next *average*, and the largest *bad*.




Figure 40: Elbow plot: within-cluster dispersion W_k as a function of k .




Figure 41: Average silhouette $\bar{s}(k)$.

Validity. Figures 40–42 support $k = 3$ as the optimal choice. Figures 43–44 demonstrate clear geometric separation and coherent cluster centers. Figures 45 and 47 confirm that the clusters correspond to a monotone ordering of swim ability (*good* < *average* < *bad*) and exhibit high within-cluster cohesion. This validates the construction of the swimmer-ability variable employed in the main analysis.




Figure 42: Gap statistic $\text{Gap}(k)$ with one-SE rule.




Figure 43: PCA projection of athletes with $k = 3$ clusters and 95% ellipses.




Figure 44: Cluster centroids across event categories (z-scored).




Figure 45: Swim times by cluster and event category (IQR boxes, $1.5 \times \text{IQR}$ whiskers). Median times follow *good* < *average* < *bad*.



Figure 46: Cluster sizes (athlete counts).




Figure 47: Silhouette widths by athlete for $k = 3$.

A.3. Regression Results: Effect Heterogeneity

A.3.1 Heterogeneous Effects by Gender on Free Riding

Table 12: Heterogeneous Effects by Gender on Free Riding

	(1) MALE OLS, Dep. Var.: Free Riding	(2) FEMALE OLS, Dep. Var.: Free Riding
year = 2015	13.076*** (2.754)	26.205*** (2.780)
year = 2016	8.860*** (2.028)	25.116*** (2.414)
year = 2017	-1.026 (2.429)	13.292*** (2.392)
year = 2018	7.984*** (2.013)	5.737** (1.845)
year = 2020	-35.052*** (2.304)	-21.851*** (1.909)
year = 2021	-12.537*** (2.091)	-4.350* (1.881)
year = 2022	-5.127* (2.364)	-7.027*** (1.978)
year = 2023	-2.673 (2.631)	-4.571 (2.531)
year = 2024	9.997*** (2.555)	2.535 (2.897)
Observations	3,490	11,538
RMSE	20.9	40.2
Adj. R ²	0.588	0.682
Within R ²	0.182	0.057
Group size FE	✓	✓
Group Gender Composition mix FE	✓	✓
Group Drafter position FE	✓	✓
Event category FE	✓	✓
Event Name FE	✓	✓
Event Location FE	✓	✓
Athlete FE	✓	✓
Athlete Swim skill FE	✓	✓

Notes: Standard errors clustered at the athlete level are in parentheses.
Significance levels: ***p < 0.01, **p < 0.05, *p < 0.10, · p < 0.10.

A.3.2 Heterogeneous Effects by Swim Ability on Free Riding

Table 13: Heterogeneous Effects by Swim Ability on Free Riding

	(1) GOOD	(2) BAD	(3) AVERAGE
	OLS, Dep. Var.: Free Riding		
year = 2015	17.487*** (3.332)	45.087*** (8.087)	25.862*** (3.598)
year = 2016	19.124*** (3.213)	38.071*** (6.367)	20.414*** (2.793)
year = 2017	1.713 (2.954)	26.649*** (6.274)	12.976*** (2.902)
year = 2018	2.349 (2.620)	11.192** (3.474)	5.382* (2.153)
year = 2020	-7.957** (2.482)	-32.809*** (3.488)	-30.837*** (2.508)
year = 2021	-3.438 (2.851)	-7.420* (3.528)	-8.766*** (2.493)
year = 2022	-7.704** (2.758)	-1.229 (4.317)	-9.907*** (2.546)
year = 2023	-5.107 (3.351)	-9.941* (4.773)	-1.862 (3.291)
year = 2024	-2.529 (4.509)	10.474 (5.704)	5.061 (3.545)
Observations	5,945	2,798	6,285
RMSE	41.1	38.2	36.4
Adj. R ²	0.627	0.688	0.717
Within R ²	0.028	0.122	0.082
Group Size FE	✓	✓	✓
Group Gender Composition Mix FE	✓	✓	✓
Group Drafter Position FE	✓	✓	✓
Event Category FE	✓	✓	✓
Event Name FE	✓	✓	✓
Athlete FE	✓	✓	✓

Notes: Standard errors clustered at the athlete level are in parentheses.

Significance levels: ***p < 0.01, **p < 0.05, *p < 0.10, · p < 0.10.

A.3.3 Heterogeneous Effects by Age on Free Riding

Table 14: Year Fixed Effects on Free Riding by Age Bin

	(20,25]	(25,30]	(30,35]	(35,40]	(40,45]	(45,50]	(50,55]	(55,60]	(60,65]
year = 2015	23.382* (10.151)	2.185 (6.193)	15.656` (8.113)	21.082** (7.709)	12.624 (8.046)	31.260** (10.225)	39.769*** (9.601)	13.096 (15.271)	2.844 (47.737)
year = 2016	19.843* (8.273)	5.528 (7.966)	11.265` (6.386)	18.226* (8.106)	15.633` (9.255)	27.338** (9.080)	29.292*** (8.299)	13.975 (11.528)	6.497 (21.515)
year = 2017	30.764* (12.214)	-0.895 (7.023)	10.185 (6.802)	23.333* (10.108)	1.371 (9.002)	35.294*** (9.545)	55.286*** (14.231)	24.177` (13.346)	19.618 (14.961)
year = 2018	15.821* (6.883)	0.074 (6.560)	5.797 (6.558)	7.763 (6.910)	1.014 (7.909)	12.383` (6.929)	2.377 (7.780)	0.264 (11.158)	-10.817 (11.817)
year = 2020	-13.603** (4.926)	-21.158*** (6.165)	-13.456** (4.780)	-23.919** (7.491)	-28.724*** (8.617)	-32.547*** (7.259)	-19.706*** (5.143)	-24.452* (9.544)	-34.659*** (8.475)
year = 2021	-2.666 (5.846)	-3.815 (8.010)	7.007 (9.502)	-0.923 (11.574)	-7.862 (11.059)	-5.507 (7.475)	5.073 (7.627)	2.940 (8.895)	-8.151 (6.957)
year = 2022	2.963 (8.960)	-14.890` (7.954)	7.299 (8.253)	-8.173 (9.215)	-7.743 (9.514)	-16.141* (7.467)	6.075 (8.559)	-5.435 (11.054)	-14.879 (9.709)
year = 2023	-8.667 (7.314)	-7.707 (8.213)	8.820 (11.919)	-3.149 (12.086)	-6.402 (9.411)	-9.058 (12.138)	-0.670 (8.114)	-10.134 (9.276)	-4.676 (10.588)
year = 2024	-8.898 (7.665)	-9.682 (9.877)	9.598 (10.049)	-19.976** (7.039)	14.237 (15.572)	-1.064 (11.157)	6.314 (9.004)	2.669 (9.237)	8.331 (13.033)
Observations	1,008	1,732	2,066	2,218	2,377	2,074	1,632	883	319
RMSE	54.6	62.1	63.5	69.7	68.2	71.2	66.3	59.3	57.4
Adj. R ²	0.233	0.175	0.187	0.202	0.227	0.264	0.301	0.300	0.318
Within R ²	0.064	0.018	0.015	0.046	0.031	0.080	0.090	0.040	0.063
Event category FE	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Standard errors clustered at the event level in parentheses.
Significance levels: ***p < 0.01, **p < 0.05, *p < 0.10, `p < 0.10.

A.3.4 Heterogeneous Effects by Drafter Position on Free Riding

Table 15: Year Fixed Effects on Free Riding by Drafter Position

	Drafter position							
	1	2	3	4	5	6	7	8
year = 2015	26.001** (9.265)	29.542*** (7.511)	43.445*** (12.575)	43.764` (24.498)	24.663 (29.021)	3.055 (5.973)	39.820* (18.401)	23.071 (18.816)
year = 2016	23.525* (10.098)	22.178*** (6.525)	42.104*** (11.897)	31.698 (26.135)	49.701* (19.861)	13.684** (5.249)	27.024 (18.398)	19.967 (17.733)
year = 2017	9.580 (8.725)	7.896 (7.488)	0.158 (10.231)	29.526 (25.703)	38.878* (18.320)	2.967 (4.763)	21.477 (16.316)	5.076 (12.760)
year = 2018	-2.004 (7.194)	0.638 (5.974)	20.826* (8.764)	16.737 (24.052)	37.902* (14.773)	-2.855 (4.043)	25.200` (13.120)	8.532 (12.319)
year = 2020	-31.538*** (7.619)	-17.886*** (4.691)	-2.664 (8.972)	-53.006` (30.894)	-21.136 (23.088)	-23.429*** (4.157)	-6.425 (15.094)	0.522 (11.837)
year = 2021	-12.835` (6.937)	-6.190 (5.453)	19.606* (8.294)	-50.833` (27.820)	8.479 (19.562)	-11.937** (4.431)	-9.802 (15.422)	-3.358 (11.887)
year = 2022	-14.824* (7.438)	-5.273 (5.257)	7.385 (7.884)	-19.327 (23.884)	32.669` (16.937)	-14.551*** (4.056)	11.364 (16.343)	9.591 (14.767)
year = 2023	-4.447 (7.502)	-8.188 (6.386)	4.894 (14.480)	5.444 (27.278)	22.100 (18.389)	-11.782* (4.748)	7.574 (12.944)	7.089 (13.099)
year = 2024	-10.077 (12.173)	1.160 (7.992)	5.882 (13.970)	-9.080 (20.705)	72.760** (27.985)	-2.053 (6.035)	16.832 (13.291)	37.775** (14.480)
Observations	1,733	2,524	1,276	550	689	3,650	960	782
RMSE	25.8	28.8	22.2	8.6	13.8	31.7	17.5	13.1
Adj. R ²	0.543	0.610	0.414	1.077	1.589	0.593	-0.451	-9.244
Within R ²	0.075	0.055	0.092	0.378	0.203	0.035	0.077	0.087
Group size FE	✓	✓	✓	✓	✓	✓	✓	✓
Group Composition Gender mix FE	✓	✓	✓	✓	✓	✓	✓	✓
Event category FE	✓	✓	✓	✓	✓	✓	✓	✓
Event Name FE	✓	✓	✓	✓	✓	✓	✓	✓
Event Location FE	✓	✓	✓	✓	✓	✓	✓	✓
Athlete Swim skill FE	✓	✓	✓	✓	✓	✓	✓	✓
Athlete FE	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Standard errors (in parentheses) are clustered at the athlete level.

Significance: ***p < 0.01, **p < 0.05, *p < 0.10, `p < 0.10.

A.4. Identification Robustness: CausalWizard.app Results For the Covid-19 Policy Shock As Treatment

Framework and Estimand

We follow the potential outcomes framework (Rubin 1974; Imbens and Rubin 2015). Let $T \in \{0, 1\}$ denote treatment ($T=1$ if covid > 0 , $T=0$ if covid ≤ 0), and Y denote the outcome ($Y=\text{free_riding}$). For unit i the potential outcomes are $Y_i(1)$ and $Y_i(0)$, and the observed outcome is

$$Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0).$$

The target estimand is the Average Treatment Effect (ATE):

$$\text{ATE} = \mathbb{E}[Y(1) - Y(0)].$$

Identification assumes a set of covariates X such that $\{Y(0), Y(1)\} \perp T \mid X$ (backdoor criterion Pearl 2009), positivity, consistency, SUTVA, and correct model specification (here, linearity).

Main Result and Interpretation

Causal Wizard estimates the ATE using linear regression under backdoor adjustment:

$$\widehat{\text{ATE}} = -21.87, \quad 95\% \text{ CI} = [-24.47, -19.01].$$

Interpretation (simple): on average, treatment reduces Y by ≈ 21.87 units relative to control.

Counterfactual Summaries

Empirical counterfactual scenarios (means over re-assigned treatments or re-weightings):

$$\begin{aligned} \mathbb{E}[Y \mid \text{all treated}] &= 14.42, & \mathbb{E}[Y \mid \text{all controls}] &= 36.29, \\ \mathbb{E}[Y \mid T=1] &= 20.27, & \mathbb{E}[Y \mid T=0] &= 36.29. \end{aligned}$$

These values are coherent with a large negative effect.

Generalization (Out-of-Sample Prediction)

On a held-out test set of $n=1000$:

$$R^2 = 0.13, \quad \text{RMSE} = 68.65, \quad \text{MAE} = 46.46.$$

While the predictive model explains limited variance, the treatment effect estimate is stable in robustness checks below.

Robustness and Refutation Tests

We formalize the three tests used by Causal Wizard. Throughout, let $\hat{\tau}$ denote the estimated ATE on the observed data $\mathcal{D} = \{(Y_i, T_i, X_i)\}_{i=1}^N$.

(1) Bootstrap Outcome Permutation Test (Significance). Under the sharp null $H_0: \tau=0$, the joint distribution of outcomes does not depend on the assignment relative to the

estimator, so permuting Y breaks any causal signal while preserving its marginal distribution. For $b = 1, \dots, B$ (here $B=100$):

1. Draw a random permutation π_b of $\{1, \dots, N\}$ and set $Y_i^{(b)} := Y_{\pi_b(i)}$; keep (T_i, X_i) fixed.
2. Compute $\hat{\tau}^{(b)}$ on $\{(Y_i^{(b)}, T_i, X_i)\}$ using the same estimator as for $\hat{\tau}$.

Define the (two-sided) permutation p-value

$$\hat{p} = \frac{1}{B} \sum_{b=1}^B \mathbf{1}\left\{ \left| \hat{\tau}^{(b)} \right| \geq \left| \hat{\tau} \right| \right\}.$$

Result reported: $\hat{p} \in [0, 0.01] \leq 0.05$ (“Pass”), providing evidence against H_0 .

(2) Placebo Treatment Test (No Spurious Association). Here we permute T to destroy any link between treatment and outcome while preserving the marginals of T and (Y, X) . For $b = 1, \dots, B$:

1. Draw a random permutation π_b and set $T_i^{(b)} := T_{\pi_b(i)}$; keep (Y_i, X_i) fixed.
2. Compute $\hat{\tau}_{pl}^{(b)}$ on $\{(Y_i, T_i^{(b)}, X_i)\}$.

We compare the observed $\hat{\tau}$ to the placebo distribution:

$$\hat{p}_{pl} = \frac{1}{B} \sum_{b=1}^B \mathbf{1}\left\{ \left| \hat{\tau}_{pl}^{(b)} \right| \geq \left| \hat{\tau} \right| \right\}.$$

Reported summary: original effect -21.87 ; placebo mean ≈ -0.18 with $p=0.96$ (“Pass”), consistent with no spurious effect under random assignment of T .

(3) Random Common Cause (Injected Confounder) Test (Stability). We add an artificial covariate $U^{(b)} \sim \mathcal{N}(0, 1)$ independent of (Y, T, X) to stress-test sensitivity to extra (irrelevant) common causes:

1. For $b = 1, \dots, B$, draw $U^{(b)} = \{U_i^{(b)}\}_{i=1}^N$ i.i.d. from $\mathcal{N}(0, 1)$.
2. Refit the estimator including $U^{(b)}$ as a covariate, yielding $\hat{\tau}_{aug}^{(b)}$.

Stability is assessed by the overlap of $\hat{\tau}$ with the sampling distribution of $\{\hat{\tau}_{aug}^{(b)}\}$:

$$\hat{p}_{aug} = \frac{1}{B} \sum_{b=1}^B \mathbf{1}\left\{ \left| \hat{\tau}_{aug}^{(b)} - \hat{\tau} \right| \geq \left| \hat{\tau} \right| \right\}.$$

Reported: original effect -21.87 , modified effect remains -21.87 , with $p=0.86$ (“Pass”). This suggests the estimate is insensitive to adding an irrelevant confounder.

Notes. The three procedures above are model-agnostic wrappers around the chosen estimator: they hold fixed the estimation recipe and only perturb the data in targeted ways. Under H_0 or under the null that the effect is spurious/unstable, the perturbed statistics should mimic or exceed the observed effect in magnitude with nontrivial probability; small empirical p (test 1) and large placebo/augmented p values (tests 2–3) support a genuine and stable effect.

Assumptions

- **Ignorability (Backdoor):** $\{Y(0), Y(1)\} \perp T \mid X$.
- **Positivity:** For all x in the support of X , $0 < \Pr(T=1 \mid X=x) < 1$.

- **Consistency:** If $T_i = t$, then $Y_i = Y_i(t)$.
- **SUTVA:** No interference; one unit's treatment does not affect another's outcome.
- **Model specification:** Linear model correctly captures the conditional mean $m(t, x) = \mathbb{E}[Y | T=t, X=x]$ used by the estimator.

Directed acyclic graph (DAG)




Figure 48: Directed acyclic graph (DAG) used for identification.

Conclusion

The estimated ATE of -21.87 with 95% CI $[-24.47, -19.01]$ indicates a strong, statistically significant, and robust negative causal effect of treatment on the outcome. The three refutation tests support significance, non-spuriousness, and stability of the estimate, respectively.