

The Reach of the World Cup Distraction Effect: Evidence from Global Trading Venues

Daniel Gatto*¹

¹Universidade Paulista (UNIP)

June 2026

Abstract

The FIFA World Cup is a clean, globally synchronized attention shock. Yet the markets that now carry most of the world's trading do not visibly respond to it in aggregate, and the headline evidence that they do turns out to be an artifact of how the effect is measured. Ehrmann and Jansen show that trades on a country's national exchange roughly halve when its team plays, carried by a local trader who is also the viewer. Two questions sit conflated on that mechanism: whether the global attention shock moves aggregate trading in the continuous, global, derivative venues where no team maps to the marginal trader, and whether the local effect generalizes beyond its first exchange. We separate them.

Our most transferable result is that two natural ways of measuring the effect manufacture it from no underlying response: a cross-market comovement decline (-0.086 , $p = 0.005$) that vanishes once the set of open instruments is held fixed, and an all-session cash-equity volume increase ($+19\%$) that is extended-hours sampling. Each reproduces the prior finding from a process with no match effect, the comovement case shown general by a Monte Carlo.

On one-minute data for eleven instruments across five market structures over 2018 and 2022, the aggregate channel is bounded near zero: equivalence tests reject a 30% effect in every group (declines beyond about 9% excluded in crypto spot, 6% pooled), and an injected 30–45% decline is recovered with full power, so these are genuine nulls, not weak tests. The local channel we bound rather than refute. No mapped venue we can reach shows a decline and the daily home response is bounded below about 11%; but the exact Ehrmann–Jansen object, a minute-level own-exchange trade count whose trader is the viewer, is left untested for want of free foreign intraday data. The next-day loser effect is absent at daily frequency in 20 national indices and 16 currencies.

We do not overturn the distraction effect in its own setting; the contribution is the bound on the aggregate channel, the evidence that it does not travel to the venues that now carry most trading, and the two mechanisms that show how readily it is mismeasured.

*Correspondence: daniel@daru.finance.

Keywords: investor inattention, limited attention, FIFA World Cup, trading volume, sports sentiment, loser effect, composition bias, equivalence testing, market microstructure.

JEL classification: G14, G12, G15, G41, Z21.

Contents

1	Introduction	3
2	Related Work	4
3	Data and Sample	5
4	Identification	6
5	The attention shock is real	7
6	The aggregate global-attention channel	8
7	The “comovement decline” is a composition artifact	10
8	The local distraction channel	11
9	The next-day loser effect	13
10	Country and cross-sectional heterogeneity	14
11	Robustness	15
11.1	Few-cluster cells under the wild-cluster bootstrap	16
11.2	Could the design have caught the effect? A positive control	16
11.3	Futures bounds on the full 2010–2022 history	18
12	Discussion	18
13	Threats to validity	19
14	What would change the reading	20
15	Conclusion	20
A	Pre-registration	21
B	Reproducibility	21
C	Figure and table provenance	22
D	Results inventory	22

1 Introduction

The FIFA World Cup is among the largest synchronized attention events on earth, and the 2022 final drew a live audience in the hundreds of millions. Limited-attention theory predicts that an event this large pulls attention from markets and depresses trading while it runs [5, 8, 7]. The sharpest evidence is Ehrmann and Jansen [1], who match minute-level trading on national stock exchanges to the World Cup schedule: when a country’s team plays, the number of trades on that country’s exchange drops sharply, by figures around 45 to 55% at the most attentive moments, with market activity “following the ball.” A separate strand documents sentiment rather than attention, where Edmans, García and Norli [2] find a national index underperforms the day after the country loses an important match, and related work reports World-Cup-linked effects in US index returns and single-name sentiment [3, 4].

That mechanism is specific. It runs through the marginal trader of a national cash-equity exchange, who in local daytime hours is disproportionately a local resident and therefore also the marginal football viewer. The same person sits on both sides of the distraction. The venue counts discrete trades, a measure sensitive to a thinning of small retail orders, and no substitute venue is open during the match. None of those features is shared by the markets that now carry most of the world’s trading.

Two hypotheses. That specificity splits the question in two, and the two need different tests. One concerns aggregate attention rather than any local trader:

H_{global} : the synchronized global attention shock depresses aggregate trading in continuous, global, derivative, or retail-heavy markets, where no national team maps to the marginal participant.

A null in crypto, index, commodity, or currency futures bears on H_{global} , not on Ehrmann–Jansen, because a São Paulo viewer who stops trading during Brazil’s match is a negligible share of global Bitcoin or S&P-futures flow. The second hypothesis is Ehrmann–Jansen’s own, carried to new venues:

H_{local} : the local distraction effect generalizes beyond the single national cash exchange to other mapped settings, where a team’s own currency or home market is traded while the team plays.

H_{local} is the structurally comparable test, and it is the hard one. The genuinely matching object, a minute-level own-exchange trade count whose marginal trader is the local viewer, needs foreign intraday data that is not freely available. Figure 1 lays out the split and where each channel lands.

What we find. The attention shock is real and large; Wikipedia tournament-article views run $5.7\times$ and $8.0\times$ their baseline (Section 5). The aggregate channel does not carry it. Across one-minute data on eleven instruments in five market structures, during-match volume is economically small and statistically indistinguishable from zero in every group, and equivalence tests reject a 30% effect in every group while the deepest venues exclude declines beyond about 9% in crypto spot and 6% pooled (Section 6). A positive-control injection confirms the design recovers a 30–45% effect with essentially full power at these cluster counts, so the nulls are informative rather than dead (Section 11). On the local




Figure 1: One attention shock, two channels, two marginal traders. The synchronized World Cup shock splits into an aggregate-attention hypothesis in globally intermediated venues (H_{global}) and the local Ehrmann–Jansen mechanism carried to mapped venues (H_{local}), which need different tests and resolve differently. The bottom panel is the transferable warning: two ordinary specifications reproduce the prior finding from a process with no underlying response.

channel the bound reaches only as far as the available data. CME own-currency futures, US-listed country ETFs, and daily home equities all show no decline; the daily home response is bounded below about 11%; and the minute-level local-trader effect is left untested (Section 8). At daily frequency the next-day loser effect is absent in 20 national indices and 16 currencies (Section 9). Two specifications reproduce the prior result as a composition artifact, the set of open instruments doing the work rather than any market response (Sections 7 and 9).

Contributions. We bound H_{global} : the aggregate global attention shock does not move aggregate trading in continuous global markets, and we report how large an effect each venue excludes rather than asserting a bare null. For H_{local} we offer a calibrated reading, since the distraction effect does not appear in any mapped venue we can reach but the precise Ehrmann–Jansen object is left untested with power; we therefore bound the daily home response and name the intraday foreign-exchange test as the open question rather than claiming a refutation in the original setting. Alongside the two hypotheses, we isolate two composition mechanisms by which natural specifications produce the prior finding from no underlying effect, the comovement case shown general with a Monte Carlo in which a zero-effect process produces a significant decline in 100% of simulations.

2 Related Work

Three literatures meet at this paper. On *limited attention*, the evidence is that attention is a scarce resource whose diversion moves markets: Barber and Odean [5] show retail buying concentrates in attention-grabbing stocks, DellaVigna and Pollet [7] find Friday earnings

news is underreacted to because attention is low, Hirshleifer, Lim and Teoh [8] document distraction by competing same-day news, and Da, Engelberg and Gao [6] validate search volume as a direct attention proxy. The World Cup is this literature’s cleanest natural experiment, a pre-scheduled attention shock with no confounding fundamental news.

A second strand applies that lens to football. Ehrmann and Jansen [1] are the direct antecedent and the result this paper tests: matching minute-level national-exchange trading to the World Cup schedule, they find trade counts on a country’s own exchange fall by figures around 45 to 55% while its team plays, with attention “following the ball.” Edmans, García and Norli [2] document the sentiment counterpart, a next-day loser effect in the losing nation’s index, and Kaplanski and Levy [3] and Bernile and Lyandres [4] report World-Cup-linked effects in US index returns and in single-name sentiment. What this literature has not done is separate the aggregate-attention channel from the local-trader mechanism, or test either outside the single national cash-equity exchange. Both findings rest on one venue, the centralized daytime national exchange whose marginal trader is the local viewer, and neither has been carried to the continuous, global, derivative, and retail-heavy venues that now intermediate most volume.

Methodologically the paper draws on an *inference* toolkit the original event studies predate. We bound each null with the two-one-sided-tests equivalence procedure [12] rather than reading a non-rejection as evidence of absence. The many-tests problem is handled with false-discovery control [10], and the count of nominal rejections is read against the multiple-testing critique of empirical asset pricing [11]. Where a cell rests on few match-day clusters we replace asymptotic cluster-robust inference, which over-rejects in that regime, with the wild-cluster restricted bootstrap of Cameron, Gelbach and Miller [13]. The contribution relative to all three is to split the World Cup question into a bounded aggregate channel and a power-calibrated local channel, and to show that two natural specifications manufacture the prior finding through composition.

3 Data and Sample

Markets. The core panel is one-minute data on eleven instruments in five market structures (Table 1). Crypto bars come from `data.binance.vision` and carry dollar volume, trade count, and taker-buy volume; spot covers both tournaments and the perpetual enters in 2022. Futures one-minute trade bars come from AlgoSeek, front contract by daily volume, with quote-bearing TAQ for the liquidity tests and an end-of-day options-greeks feed from the same source. The symmetric cross-market core spans the 2018 and 2022 tournaments, the two for which liquid one-minute crypto data exist, so the comovement and pooled equivalence tests run over one common instrument set rather than a futures set that widens as crypto enters; the futures venues, whose AlgoSeek history runs back to 2010, are additionally re-estimated on four tournaments (2010–2022) in Section 11.3 to bound the futures cells on all the data that exists. All core series are restricted to ± 21 days around each tournament so control minutes share the instrument’s own regime. Because many matches fall on weekends, when futures are closed and crypto is the only open venue, the futures during-match sample concentrates in weekday and Sunday-Globex matches.

Home markets. The local-channel arm (Section 8) adds three daily panels, the finest resolution freely available for foreign venues. Home equities for 26 nations (main index

Table 1: The eleven core instruments span five market structures, chosen to differ from a national cash-equity exchange on the dimensions limited-attention theory says should matter.

Group	Instruments	Structure
Crypto spot	BTC, ETH spot (Binance)	continuous, global, retail-heavy
Crypto perpetual	BTC, ETH USD _S -M perpetual	continuous, leveraged, retail-heavy
Equity futures	ES, NQ (CME)	near-24h, institutional
Commodity futures	GC, CL (CME)	near-24h, global
FX futures	6E, 6B, 6J (CME)	near-24h, automated
Cash equities	SPY, QQQ, IWM, 9 country ETFs	exchange-hours, the literature’s venue

plus top-capitalization names, about 100 series, 2009–2023) come from the Yahoo Finance daily feed. The top-capitalization constituents are chosen on current market cap rather than point-in-time membership, so the single-name cross-section is mildly survivorship-tilted; because the home-equity result is a match-day activity null and survivorship bears no direction on whether a fixture depresses trading, the tilt does not affect the bound, but we note it. The tilt touches only the volume null; the next-day loser-effect returns (Section 9) are measured on the national indices, not the single names, so that test is survivorship-clean. Twenty national stock indices with deep history (1998–2022) and sixteen home currencies (each gated at its float or redenomination date: 1998 for the five developed units and South Korea, 2000–2007 for the rest, 2015 for the rouble) come from TradingView/ICE daily exports. Argentina’s 1991–2001 dollar peg and the pre-2005 Turkish lira are excluded as untestable. Index returns are adjusted by the MSCI world return and currency returns are normalized to the dollar value of the local unit, so a negative return denotes home-currency depreciation.

Schedule, attention, controls. Match results, kickoff times, and eliminations back to 1998 come from the Fjelstul World Cup database, with goal-level timing for 2010–2022. Attention is the daily Wikipedia tournament-article pageview series, with Google Trends and GDELT news volume as cross-checks. Estimation uses the completed 2010–2022 tournaments; the ongoing 2026 World Cup is excluded, though a live collector records it for a later out-of-sample extension.

4 Identification

For instrument i and minute t we estimate

$$Y_{it} = \beta \text{Match}_t + \alpha_{i,w(t)} + \gamma_{d(t)} + \varepsilon_{it}, \quad (1)$$

where Match_t is one during the in-play window of any match, $\alpha_{i,w(t)}$ is a fixed effect for instrument interacted with minute-of-week $w(t)$, and $\gamma_{d(t)}$ is a calendar-date fixed effect. The minute-of-week interaction compares a match minute to the same instrument at the same minute of the same weekday on non-match days, removing each instrument’s intraday and weekly seasonality; the date effect absorbs whole-day shocks common across instruments. With the date effect included, β is identified *within* a calendar date, from the contrast between match and non-match minutes of the same day net of the minute-of-week seasonal: non-match days enter through the seasonal baseline they help estimate rather

than as a direct cross-day comparison, so a quiet trading hour cannot masquerade as a match effect and a whole-day volatility or macro regime cannot leak into β . Standard errors cluster by calendar date, so the effective sample size for any treatment is the number of distinct match-days, not the number of minutes. The local-channel tests (Section 8) replace Match_t with an own-team indicator that is one only when the mapped nation's own team is playing. Outcomes are log volume, log trade count, absolute return, and the Parkinson high-low range [14].

Power and bounds. A null is informative only if the test could have caught a real effect. Each cell is reported with its 95% confidence interval, the largest decline still consistent with the data, and the minimum detectable effect at 80% power, then checked against a 30% equivalence margin by two one-sided tests [12]. Where a single market rests on very few match-day clusters, asymptotic cluster-robust inference over-rejects [13], so those cells are not flagged by assertion: each is re-tested with a wild-cluster restricted bootstrap (Frisch–Waugh residualization on the fixed effects and the foreign-match control, the null imposed, Rademacher cluster weights over match-days, $B = 4999$), its bootstrap p reported beside the asymptotic one (Section 11). A well-powered cell's two p -values agree; a few-cluster cell's bootstrap p inflates.

Pre-registration of the design. The estimating equation, the clustering level, the 30% equivalence margin, the primary-versus-exploratory split, and the rule that few-cluster cells are read through the bootstrap rather than the asymptotic p are fixed before the cells are read. The 2026 tournament is held out of every estimate. Departures from a bare null, where they occur, are reported with the mechanism that produces them rather than corrected away.

5 The attention shock is real

A null market response means nothing unless the attention shock is large, so we fix that first. Tournament-article Wikipedia pageviews run $5.7\times$ their off-tournament baseline in 2018 and $8.0\times$ in 2022, and the daily series tracks the match calendar, correlating 0.34 and 0.70 with the number of matches played that day. Figure 2 sets that spike beside the during-match volume response in the four no-mapping market groups: the left panel rises several-fold, the right sits on zero. Reconciling that gap occupies the rest of the paper.

The spike is intraday, not just a match-day average. The daily Wikipedia ratio could reflect heightened interest across the whole match-day rather than during the match itself, which would leave open whether attention surges in the trading minutes the treatment isolates. The Wikimedia hourly pageview dumps settle it at the during-match resolution, and we check ten matches spanning every round from group stage to final across both tournaments rather than rest the point on one. Reconstructing each match hour by hour, the two competing national-team articles and the tournament article jump in the kickoff hour against the same articles three hours earlier. For the 2022 final (Argentina–France, kickoff 15:00 UTC, 2022-12-18) the tournament article runs $2.9\times$ ($5,245 \rightarrow 15,219$ views), Argentina's $3.3\times$ ($1,038 \rightarrow 3,399$), France's $3.1\times$ ($893 \rightarrow 2,775$), and the Lionel Messi article $3.8\times$ ($1,984 \rightarrow 7,582$). The pattern holds across all ten: the team-article kickoff-hour spike averages $3.4\times$ the pre-match baseline, with per-match means from 1.7




Figure 2: Attention spikes, the aggregate channel does not respond. Left: Wikipedia tournament-article pageviews as a multiple of the off-tournament baseline. Right: during-match change in trading volume for the four market groups with no national-team mapping, 95% confidence intervals, all bracketing zero.

to $5.7\times$ and a clear rise in every match. The largest jumps are the group-stage surprises (Argentina–Saudi Arabia $5.0\times$, Germany–Japan $5.2\times$, Argentina–Netherlands $5.7\times$); the smallest single-team cell is Brazil at $1.4\times$, a perennially high-traffic article with little headroom rather than a weak response. The tournament article, already elevated throughout the event, moves less. The attention shock the rest of the paper conditions on is present at the hour of play, not merely on the calendar day, so the market nulls below are a non-response to a shock that is real at their own resolution. Figure 3 shows the kickoff-hour spike for all ten reconstructed matches.




Figure 3: Attention is present at the hour of play. For ten matches across every round of both tournaments, the competing teams’ Wikipedia articles draw more views in the kickoff hour than three hours earlier, by 1.7 to $5.7\times$. The largest jumps are the group-stage surprises; the smallest cell is a perennially high-traffic article with little headroom, not a weak response.

6 The aggregate global-attention channel

H_{global} asks whether the global attention shock depresses aggregate trading where no team maps to the marginal participant. It does not. Estimating Eq. (1) with log volume,

the during-match change is -2.7% in crypto spot, -3.9% in crypto perpetuals, -2.9% in equity-index futures, $+7.2\%$ in commodity futures, and $+11.0\%$ in currency futures; pooled across all eleven instruments it is $+2.3\%$. None is significant, and the estimates are not consistently signed, as compatible with a small increase as a small decrease, never the sharp decline distraction predicts. Table 2 reports each cell with its interval and the decline it excludes. The bound is tightest in the deepest market: crypto spot, on 48 match-day clusters, rules out during-match volume declines beyond about 9%, and the pooled estimate rules out declines beyond about 6%. Equivalence tests against a 30% margin reject an effect that large in every group. The bounds are not uniform, and we keep the gradient explicit rather than averaging it into one headline: crypto spot and the pooled estimate exclude declines beyond 9 and 6% and the 2.2-million-minute quoted spread is pinned to $\pm 1\%$, whereas the weekday-only futures cells are looser, ruling out only declines beyond about 15% (equity-index futures) at a minimum detectable effect near 22%. The two futures cells with positive point estimates make this explicit. Their decline-excluded figure in Table 2 is mechanically small, with FX futures excluding declines beyond only about 2%, but that is because the estimate sits above zero, not because the null is tight, since the largest decline still consistent with the data is read off the lower confidence bound. The cell’s true precision is its interval width, summarized by the minimum detectable effect, and on that measure these futures cells remain the loose ones. Re-estimated on the full 2010–2022 futures history (Section 11.3), the futures cells rest on 65 match-day clusters and still reject the 30% margin; commodity and FX tighten while equity-index futures stays the loosest, and the tightest bounds in the study remain crypto spot and the pooled cell. The deep, continuously traded venues deliver precise nulls; the thinner cells deliver bounded ones. Figure 4 plots the same estimates against the -30 to -50% band the cash-equity literature reports; no interval reaches it.

Table 2: The aggregate global-attention channel (H_{global}). During-match change in trading volume by market group, with the largest decline still consistent with the data at 95%. Every interval brackets zero and the pooled estimate is positive; the marginal trader in each venue is global, not a national viewer.

Market group	Effect (%)	95% CI (%)	Decline excl.	Clusters
Crypto spot (BTC, ETH)	-2.7	$[-9.4, 4.5]$	9%	48
Crypto perpetual	-3.9	$[-16.2, 10.1]$	16%	23
Equity-index futures (ES, NQ)	-2.9	$[-15.3, 11.3]$	15%	31
Commodity futures (GC, CL)	$+7.2$	$[-5.9, 22.2]$	6%	31
FX futures, average match	$+11.0$	$[-2.4, 26.3]$	2%	31
All instruments pooled	$+2.3$	$[-6.1, 11.5]$	6%	48

A minute-by-minute event study reaches the same place without the modeling. Figure 5 averages log volume in event time around kickoff across all matches: the series is flat through the opening whistle and across the 105-minute in-play window, with no step at kickoff and no recovery at the close. A formal pre-trend check agrees, with eight pre-kickoff bins all insignificant (smallest $p = 0.19$).

Beyond volume, the liquidity and risk-pricing channels are equally flat. Across 2.2 million instrument-minutes of bid/ask data, quoted spreads are unchanged ($+0.2\%$, $p = 0.66$) and displayed depth is essentially flat (-2.7% , $p = 0.08$). At-the-money implied volatility and the 25-delta put-call skew do not move, option volume is flat, and variance




Figure 4: During-match volume by market group, 95% confidence intervals, with the prior cash-equity literature’s reported -30 to -50% decline shaded. No interval reaches the band; the pooled estimate is mildly positive.




Figure 5: Trading volume in event time around kickoff, pooled across matches, relative to the same-minute-of-week baseline. The in-play window is shaded. Nothing steps at the opening whistle and nothing recovers at the close.

ratios, the spot-futures and spot-perp bases, and lead-lag measures show no change in how quickly prices update.

7 The “comovement decline” is a composition artifact

Where every direct channel came up flat, one specification reproduces the prior literature’s effect, for a reason that generalizes. Computing the mean pairwise return correlation across all available instruments in match versus control windows, comovement falls during matches by -0.086 ($p = 0.005$), the attention-driven decline the theory predicts. What moves is the set of instruments open during matches, not their correlation. Because many matches fall on weekends, the match-window correlation set tilts toward crypto, a higher-correlation cluster, relative to the weekday control set, so the mean correlation shifts on composition alone. The bias is an identity, not a coincidence of our sample. The naive estimator is a difference of set-averages, $\hat{\Delta} = \bar{\rho}_{S_m} - \bar{\rho}_{S_c}$, where $\bar{\rho}_S = \binom{|S|}{2}^{-1} \sum_{i < j \in S} \rho_{ij}$

averages the pairwise correlations over the instruments S open in a window. When the open set differs across window types, $S_m \neq S_c$, this has nonzero expectation $\mathbb{E} \bar{\rho}_{S_m} - \mathbb{E} \bar{\rho}_{S_c}$ even with every pairwise ρ_{ij} held *identical* across windows, because the two averages run over different pairs; the weekend closures make S_m crypto-heavy, a high- ρ cluster, so $\hat{\Delta} < 0$ by construction. Conditioning on the instrument set, or adding an instrument-count control, equalizes S_m and S_c and annihilates the bias term, which is precisely what the two checks that follow do.

Two checks isolate the mechanism. Restricting to the constant crypto-only set that trades in both window types, the effect vanishes (+0.012, $p = 0.54$); adding an instrument-count control to the full-sample regression does the same (+0.002, $p = 0.93$). Figure 6, left, shows only the naive estimate clears significance. The right panel establishes generality: in a Monte Carlo with an identical return process in match and control blocks and no match effect anywhere, letting only the set of contributing instruments differ across block types produces a spuriously significant comovement decline in 100% of 400 simulations, which an instrument-count control returns to the nominal 5%. The spurious-rejection rate scales with the size of the composition shift; the 100% shown is for a shift in the direction and rough magnitude of the weekend futures closures that drive the effect in our data, so it illustrates the mechanism’s generality rather than fixing a single calibrated number. Holding composition fixed strips the raw-data confirmation back to nothing.




Figure 6: The cross-market comovement decline is composition, not attention. Left: the naive raw estimate is significant (red); holding the instrument set fixed or adding an instrument-count control returns it to zero. Right: a Monte Carlo with zero true effect and only a composition shift produces a significant decline in 100% of simulations, and the control restores the nominal 5% false-positive rate.

8 The local distraction channel

Everything so far has tested the aggregate channel, where no national team maps to the marginal trader. We now turn to H_{local} , Ehrmann–Jansen’s own hypothesis carried to new venues, the structurally comparable test and the hard one. The Ehrmann–Jansen object is a minute-level trade count on a national cash exchange whose marginal trader is the local viewer. No freely available data match all three of those features at once, so we test the cells that match two of them and read the result accordingly (Table 3). CME own-currency

futures during the mapped nation’s matches are minute-level but trade on a global, largely automated book rather than to local viewers. US-listed country ETFs during the home nation’s matches are minute-level but trade to US investors. Home equities carry the right marginal trader but only at daily frequency.

Table 3: The local distraction channel (H_{local}): own-team effects in mapped venues. No cell matches all three features of the Ehrmann–Jansen object (minute-level, own home exchange, local marginal trader). Where the test is powered the estimate is positive or flat; the genuinely matching object is left untested.

Mapped own-team cell	Effect (%)	95% CI (%)	Clust.	Resol.	Marginal trader
Euro futures (6E), euro-area matches	+12.7	[0.6, 26.2]	60	1-min	global
Sterling fut. (6B), England matches [†]	−30.0	[−44.0, −12.5]	10	1-min	global
Yen futures (6J), Japan matches	+15.1	[−7.3, 42.8]	11	1-min	global
Own-team futures, pooled (2 tourn.)	+11.0	[−4.0, 28.2]	28	1-min	global
Own-team futures, pooled (4 tourn.)	+8.9	[−3.9, 23.6]	60	1-min	global
Country ETFs, own-country matches	−4.7	[−17.2, 9.7]	55	1-min	US
Home equities, own-match day	−0.6	[−10.7, 10.7]	63	daily	home (local)

Marginal trader: *global* = CME book, *US* = US-listed ETF, *home (local)* = the resident viewer (the Ehrmann–Jansen trader); only the last matches the original mechanism, and it is observed only at daily resolution. [†] Ten match-days; date-clustered inference over-rejects at this cluster count [13], so this cell is not read as a finding.

Read across the table, no mapped cell shows a during-match decline, and the two best-powered cells are positive with intervals that exclude one. The euro-futures cell, on 60 own-match days, is +12.7% and its interval excludes any decline; the four-tournament own-team pool, also 60 clusters, is +8.9% and rules out declines beyond about 4%. Both powered cells lean on the euro: euro-area fixtures supply 24 of the 28 treated own-team dates in the two-tournament pool and remain the majority in the four-tournament extension, so the pooled estimate is close to the euro effect rather than independent corroboration across four currencies, a gain in power and not in breadth, with the yen and sterling the only non-euro own-currency tests and both individually underpowered. The one large negative cell, sterling at −30%, rests on ten match-days, exactly the regime where cluster-robust inference over-rejects,¹ and the per-nation cells conflict in sign (the yen +15%, the peso −21%, a meaningless Canadian-dollar +91% on two match-days). The home-equity cell carries the right marginal trader and shows −0.6% on 63 clusters, excluding a day-level home decline beyond about 11%.

What the table cannot do is see the Ehrmann–Jansen effect itself. That effect is a thinning of small trades during the match minutes on the team’s own exchange. It can be present at minute resolution while averaging to nothing over a daily bar. Our minute-level mapped cells trade to global or US participants who do not watch the match, and our right-trader cell is daily. The minute-level own-exchange trade count with a local marginal trader is therefore the one test we cannot run, and we do not claim to have run it. Closing that gap needs intraday data for foreign cash exchanges, which is not freely distributed; it is the natural target for a paid-data extension.

¹Three sterling cluster counts appear in the paper, one per specification: the 2018–2022 own-team *interaction* coefficient (six match-days, Section 11), the four-tournament own-match *level* cell reported here and in Table 3 (ten), and the bootstrap specification on a marginally wider window (eleven, Table 5). All three sit in the few-cluster regime and none is read as a finding.

9 The next-day loser effect

Where the local distraction test ran into a resolution gap, its sentiment counterpart does not. The Edmans–García–Norli loser effect lives at daily frequency, its own native resolution, so here the evidence is direct. Across 20 national stock indices back to 1998 (110 match-day clusters), the first-session index return after the nation’s own match, world-return adjusted, is -0.5 bp after a loss ($p = 0.96$); the developed-market subset shows -15 bp ($p = 0.14$) matched by an equal post-*win* dip (-10 bp, $p = 0.14$), so wins and losses are statistically indistinguishable and no loser effect is present. A placebo pins down what the common negative is not: the same nations’ tournament-season index-days that are *not* post-match sessions are flat ($+0.7$ bp, $p = 0.87$, 205 date-clusters), so the post-match dip is specific to the session after a fixture rather than a season-wide calendar drift, and being equal for wins and losses it is in neither form a sentiment response to the result. The 16-currency panel back to 2002 (74 clusters) gives a next-session loser effect of $+9.9$ bp ($p = 0.13$), the opposite of what distraction predicts. Figure 7 collects the home-market estimates; all bracket zero.




Figure 7: Home-market and home-currency effects, 95% confidence intervals. Match-day home-equity volume, the next-session loser effect in 20 national indices and 16 currencies, and the small-country volume tilt all bracket zero.

Behind that pooled null, whether any effect should concentrate in the smallest, most retail-dominated economies is testable by interacting the own-match term with an emerging-market indicator, and the three measures conflict. Emerging-market home volume is 15% lower on own-match days than the developed baseline ($p = 0.30$, correctly signed but insignificant). The equity loser-effect interaction reverses sign and reaches nominal significance with the sign distraction does not predict ($+45$ bp, $p = 0.05$, emerging markets falling *less* after a loss), while the currency interaction is flat ($+3$ bp, $p = 0.82$). A correctly signed but underpowered volume tilt in the smallest markets is the seam where a home effect could still hide, and only intraday foreign-exchange data could resolve it. A single-nation stand-alone test is not estimable in any case, which is the cluster-count point again: Brazil, the best case with 34 World Cup matches across six tournaments, has only six losses since 2002, and a return regression on six treated observations gives a degenerate standard error, so we report no single-nation loser-effect p -value.

Alongside the loser tests, the local-channel arm surfaces the second composition artifact. In US cash equities, the literature’s own venue, regular-hours volume during matches is -5.9% ($p = 0.19$) and trade count -2.5% , both insignificant. A specification that adds thin pre- and post-market bars instead yields a spurious $+19\%$ trade-count

increase ($p = 0.02$), positive where distraction predicts a fall, and driven by when the sampled bars fall rather than by any market response.

10 Country and cross-sectional heterogeneity

Those pooled home-market tests average over nations and results; five sharper cuts ask whether a reaction hides in a particular country, result type, or market culture. Each reorganizes the 1998–2022 index and 2002–2022 currency panels already built (no new data), and each is read through the same cluster-count discipline, with a wild-cluster restricted bootstrap on the cells below thirty match-day clusters (Table 4).

Brazil's own win. We estimate every nation's own-win reaction on its own series, not Brazil's alone. Of the ten nations with enough World Cup wins to test (≥ 8), Brazil has the largest positive index reaction (+74bp, rank first; the cross-nation median is -7 bp and several are negative, so the result is not a noisy-everywhere artifact in which many nations look large), and the real is the one significant currency cell (+55bp, $p = 0.005$, though only Brazil and Uruguay have enough wins to test in the currency panel). Two cautions temper it. On its own index series Brazil is only marginal ($p = 0.078$) and *no* single index nation clears significance (0 of 10); the cells that do reach significance are the pooled cross-section contrast (IBOV +91bp over the cross-section, bootstrap $p = 0.023$; real +49bp, bootstrap $p = 0.025$) and the currency, both of which borrow power from the panel. And the significance is partly carried by the 2002 championship (seven of Brazil's wins, during the real's election-crisis swings), since excluding 2002 the cross-section estimates stay positive but lose significance (+76 and +39bp, bootstrap $p = 0.07$ and 0.12). We read it as a sign-robust, significance-fragile win-sentiment lead concentrated in one nation: the optimistic mirror of the loser effect, the wrong direction for distraction, and a cell that does not survive the global correction below ($q = 0.32$).

Loss aversion. If losses bite harder than wins gratify, elimination should move markets more than advancement. Held in the same knockout context, the next session after a knockout win (advancing) and after a knockout loss (eliminated) are statistically identical and both null (indices +5bp versus -0 bp, currencies +22bp versus -11 bp, none significant), and the loss-times-elimination term adds nothing (+10bp, $p = 0.56$). No asymmetry appears, so no loss aversion at this resolution.

Developed versus emerging is the size interaction of Section 9, recomputed here and unchanged: the one nominal cell (emerging equities +45bp after a loss) is wrong-signed, with emerging markets falling *less*, and does not survive the bootstrap ($p = 0.047 \rightarrow 0.053$).

Host nations. A nation playing its own tournament has unique home exposure, yet host-team results move the next session no more than away results (host-win and host-loss interactions both insignificant; the one nominal cell, a -32 bp host-win currency reaction, dies under the bootstrap, $p = 0.024 \rightarrow 0.063$). Host-tournament match-days do carry -0.34% lower intraday index volatility (bootstrap $p = 0.033$), a bootstrap-robust but economically trivial cell, and one of twenty-six in this family.

Football culture. Coding nations whose dominant domestic spectator sport is a competing major league (United States, Japan, South Korea, Australia, Canada, New Zealand) against the football-centric remainder, the result-by-culture interactions are null in both return and volatility (all $p > 0.15$): markets in the football-mad economies react no more than the rest. The coding is a coarse author-assigned binary, so this cut is the softest of the

five, but it points the same way.

Folding these into the global pool (the twenty heterogeneity coefficients plus the twelve per-nation own-win cells, the full family rather than a chosen subset) lifts it to 190 during-match and own-team tests, of which 24 are nominally significant against 9.5 expected by chance and, as before, none survives a single Benjamini–Hochberg correction; the most significant of the new cells, Brazil’s win, is one nation and partly one tournament.

Table 4: Country and cross-sectional cuts of the home-market panels (no new data). Returns in basis points, volatility in percent; 95% CI date-clustered; bootstrap p shown for cells under thirty match-day clusters (wild-cluster restricted, $B = 4999$). Read the bootstrap p and the cluster count, not the point estimate, on the thin cells.

Cut	Headline cell	Estimate	95% CI	Clust.	p (asy/boot)
Brazil win	IBOV, own win vs cross-section	+91 bp	[18, 164]	27	.015/.023
	real, own win vs cross-section	+49 bp	[9, 88]	23	.016/.025
	real, excluding 2002	+39 bp	[−6, 83]	16	.087/.12
Loss aversion	eliminate (KO loss), idx	−0.1 bp	[−39, 38]	40	.996/−
	advance (KO win), idx	+4.7 bp	[−32, 41]	43	.80/−
Dev vs emerging	loss $\times$ emerging, idx	+45 bp	[1, 89]	28	.047/.053
Host	match-day vol $\times$ host, idx	−0.34%	[−0.63, −0.04]	24	.025/.033
	host $\times$ win (return), ccy	−32 bp	[−61, −4]	23	.024/.063
Football culture	win $\times$ football-centric, idx	+65 bp	[−24, 153]	104	.15/−

11 Robustness

Past the country cuts, the nulls survive the obvious stress tests. A placebo that shifts every match window forward one day is null, excluding minutes around scheduled macro releases leaves the estimates unchanged, and re-estimating at five- and fifteen-minute aggregation reproduces the one-minute figures. Widening the universe to 29 instruments across seven asset classes, adding the full Treasury complex, metals, energy, and more equity indices, produces two nominal rejections out of 29 against 1.5 expected by chance, with a median absolute effect of 2.8% and a dead-flat Treasury complex. High-stakes, knockout, and must-win matches show no larger effect than ordinary ones. Nor does the minute a goal is scored show a spike, against the in-play-betting intuition that informed flow should arrive there [9]; the goal-window estimate is a marginally negative $\sim -7\%$ that holds steady across ± 2 , ± 5 , ± 10 and ± 15 -minute windows and shifts by under half a point when the reconstructed goal timestamps are jittered by ± 1 to ± 3 minutes, so the absence of a spike is not an artifact of approximate goal timing.

The few single-market cells that look dramatic do not withstand inference, and each fails for a stated reason rather than a generic correction. Sterling-futures volume falls 34% when England or Scotland plays and Nasdaq volume rises 35% when the United States plays, but these rest on six and two match-days, the regime where date-clustered inference over-rejects [13]. These are the 2018–2022 domestic-interaction coefficients (the differential against a no-mapped-team match); the four-tournament own-match *level* cells in Tables 3 and 5 are smaller and rest on more match-days (sterling -30% on ten and -23.5% under the bootstrap on eleven, the one-day difference being the slightly wider window the bootstrap spec uses; Nasdaq $+16\%$ on eight, $p = 0.40$), and every spec is read

the same way. Pooled across mapped markets the own-team effect is a non-significant +11%, positive where distraction predicts a fall. The comovement decline is composition (Section 7), the all-session cash-equity jump is extended-hours sampling (Section 9), and the per-currency cells conflict in sign on a handful of games each.

The multiple-testing discipline starts with the pre-registered primary family (during-match volume, trade count, absolute return, and the high–low range across the five market groups, twenty tests in all), in which *not one coefficient is nominally significant* (smallest $p = 0.12$), so a Benjamini–Hochberg correction restricted to the primary outcomes has nothing to remove. The nominal rejections that do appear in the study all sit in exploratory or robustness specifications outside this family.²

11.1 Few-cluster cells under the wild-cluster bootstrap

The claim that the dramatic single-market cells over-reject is itself testable. Table 5 re-runs the own-team cells with the wild-cluster restricted bootstrap of Section 4 beside the asymptotic cluster-robust p . The well-powered euro cell, on 59 own-match days (the bootstrap specification’s marginally wider window, against the 60 own-match days and the +12.7% level estimate in Table 3; this is the same one-day window difference that takes sterling from ten clusters to eleven, here moving the euro point estimate from +12.7 to +10.6%), is unmoved: its asymptotic and bootstrap p agree to the third decimal (0.089 versus 0.088), so the positive euro estimate is not a small-sample artifact. The sterling cell, on eleven own-match days, behaves as the few-cluster regime predicts. Its p roughly doubles under the bootstrap (0.016 to 0.031), the asymptotic figure understating the true uncertainty. The two US cells are null on both. The bootstrap validates the inference rule: the asymptotic p is trustworthy where clusters are plentiful and inflated where they are scarce. That is why the sterling and Nasdaq headline cells, even before the sign-conflict across currencies and the multiple-testing context, are read as low-cluster artifacts. The check is applied symmetrically, not only to the cells we dismiss: the pooled own-team cell we *do* interpret sits at 28 clusters, and its asymptotic and bootstrap p agree (0.161 versus 0.155 all-mapped, 0.102 versus 0.107 FX-only), so the pooled +11% is not a small-sample artifact in either direction.

11.2 Could the design have caught the effect? A positive control

A null is only evidence if the test could have rejected. We never observe the Ehrmann–Jansen effect where it is known to exist, so we inject it into our own panels and check recovery. For each headline cell we take the real outcome residuals, add a known during-match decline of –30% or –45% (the prior literature’s range), and re-estimate with the same date-clustered inference; a zero injection returns the empirical size. Table 6 reports the rejection rate over $B = 2000$ draws. At the realized match-day cluster counts the

²Beyond the primary family, the study reports 190 during-match and own-team coefficients across the activity, liquidity, options, efficiency, comovement, intensity, goal-window, sentiment, domestic, home-market, breadth, and country-heterogeneity specifications (the last including every per-nation own-win cell, so the family is not pruned to a chosen subset); 24 are nominally significant at the 5% level (against 9.5 expected by chance), every one outside the primary family above. We do not interpret this count either way. The coefficients are nested and cross-correlated, so neither the binomial chance benchmark nor a false-discovery correction speaks to the hypotheses we argue; a Benjamini–Hochberg correction [10] removes all 24, but false-discovery control protects against false positives and does no work for a null thesis. The bounds in Tables 2–3 and the per-cell mechanisms above carry the argument.

Table 5: Own-team volume cells under the wild-cluster restricted bootstrap ($B = 4999$, Rademacher weights over match-days) beside the asymptotic cluster-robust p . The powered euro cell is unchanged; the few-cluster sterling cell’s p inflates; the US cells are null throughout; and the interpreted pooled cells (28 clusters) are stable across the two p -values.

Own-team cell	Match-days	Effect (%)	p (asy.)	p (wild boot.)
Euro futures (6E), euro-area	59	+10.6	0.089	0.088
Sterling futures (6B), England	11	−23.5	0.016	0.031
Nasdaq futures (NQ), USA	8	+16.0	0.398	0.403
S&P futures (ES), USA	8	−1.1	0.950	0.951
Own-team pooled, all-mapped	28	+11.0	0.161	0.155
Own-team pooled, FX-only	28	+14.9	0.102	0.107

well-powered cells recover an injected -30% with power 1.00 and correct 5% size, and recover the coefficient without bias, so their during-match nulls are informative, not dead tests. The single-market sterling design, on six clusters, is the exception that proves the rule: its size is inflated to 0.11 and its power against -30% falls to 0.75, the same few-cluster fragility the bootstrap flags, which is why such cells are not read as findings. One caveat scopes the recovery rates: Table 6 injects a homogeneous shift on every match-minute, so it speaks to a *uniform* effect, whereas the Ehrmann–Jansen effect is largest “at the most attentive moments.” We therefore re-run the injection with the same *average* decline placed on only a fraction ϕ of each match-day’s in-play minutes, at per-minute magnitude δ/ϕ , so the concentration leaves more within-cluster residual and inflates the date-clustered standard error. Concentrating a -30% average into a quarter ($\phi = 0.25$) and then a tenth ($\phi = 0.10$) of the window leaves the well-powered cells at power 1.00 with the recovered coefficient still at the injected average (at $\phi = 0.10$: crypto spot -31.9% , equity futures -32.5% , FX futures -32.5% , pooled own-team -30.5%), and the -45% average behaves identically. At these cluster counts the detection headroom is large enough that concentration does not degrade power, so the nulls are informative against a concentrated literature-magnitude effect, not only a uniform one. The check holds the *average* fixed; an effect whose per-minute size matches the literature but whose *average* over the in-play window is single-digit is governed by each cell’s minimum detectable effect rather than by this control.

Table 6: Positive-control injection. Empirical rejection rate at the 5% level when a known during-match effect is injected into the real panel, $B = 2000$ wild-cluster draws at each cell’s realized match-day cluster count. The 0% column is the empirical size; the -30% and -45% columns are power against the prior literature’s effect. Well-powered cells detect the effect with near-certainty; the six-cluster sterling design loses size control.

Cell (design)	Clusters	Size (0%)	Power -30%	Power -45%
Crypto spot (Match)	48	0.05	1.00	1.00
Equity-index futures (Match)	31	0.05	1.00	1.00
FX futures (Match)	31	0.05	1.00	1.00
Own-team futures, pooled	28	0.05	1.00	1.00
Sterling 6B, England (own)	6	0.11	0.75	1.00




Figure 8: The nulls are powered, not dead tests. Rejection rate when a known during-match decline is injected into the real panel at each cell’s realized cluster count. By the prior literature’s -30 to -50% range the well-powered cells reject with probability one, while a zero injection returns the correct 5% size. The six-cluster sterling design is the exception: its size inflates and its power against -30% falls to 0.75 , which is why such cells are not read as findings.

11.3 Futures bounds on the full 2010–2022 history

The symmetric cross-market core holds the futures venues to 2018 and 2022 so the comovement and pooled tests run over one instrument set that includes crypto, which has no liquid one-minute history before 2017. The AlgoSeek futures run back to 2010, so the futures-only during-match cells can be re-estimated on four tournaments (2010, 2014, 2018, 2022) under the identical specification, roughly doubling the match-day clusters from 31 to 65 and using all the futures history that exists. Restricted to 2018+2022 the procedure reproduces the Table 2 cells exactly; carried to four tournaments (Table 7) every futures group still rejects the 30% margin, and the positive-control injection recovers an injected -30% and -45% with power 1.00 and nominal size at the deeper cluster count, so these are informative nulls. The deeper history does not tighten uniformly, and we keep that explicit. Commodity and FX futures sharpen, their minimum detectable effect falling from about 20% to roughly 15% , while equity-index futures does not: its 2010 and 2014 during-match contrast is noisier, the standard error does not shrink, the minimum detectable effect stays near 22% , and the point estimate moves to -5.7% so the decline it excludes widens to about 18% . Equity-index futures therefore remains the loosest cell on the deeper history as on the shorter, and the tightest bounds in the study, crypto spot at 9% and the pooled estimate at 6% , still rest on 2018+2022 because crypto cannot be carried back. Four tournaments of futures confirm the bound and the gradient without overturning either.

12 Discussion

The evidence answers the two hypotheses differently, and keeping them apart is what the data support. For H_{global} the verdict is a clean, bounded no: the synchronized global attention shock, demonstrably real in the Wikipedia series, does not depress aggregate trading in continuous global markets, and the bound excludes declines a third the size of the cash-equity literature’s in every venue. We do not oversell this. That the aggregate

Table 7: Futures during-match volume cells on four tournaments (2010–2022) beside the 2018+2022 core. Same specification (Eq. 1, log volume, date-clustered); “decline excl.” is the largest decline still consistent with the data at 95% and “MDE” the minimum detectable effect at 80% power. Clusters roughly double; all three groups reject the 30% margin on both windows; commodity and FX tighten while equity-index stays the loosest. Groups as in Table 2.

Futures group	Effect (%)	95% CI (%)	Decline excl.	MDE	Clusters
Equity-index, 2018+2022	−2.9	[−15.3, 11.3]	15%	22%	31
<i>four tournaments</i>	−5.7	[−17.9, 8.2]	18%	22%	65
Commodity, 2018+2022	+7.2	[−5.9, 22.2]	6%	20%	31
<i>four tournaments</i>	+2.2	[−7.5, 13.0]	8%	15%	65
FX, 2018+2022	+11.0	[−2.4, 26.3]	2%	20%	31
<i>four tournaments</i>	+6.3	[−3.5, 17.1]	4%	15%	65

global attention shock would move global trading was no one’s stated prior, so bounding it near zero is a useful negative rather than a surprise: a globally intermediated market does not depend on any one population’s attention to keep trading, and flow a São Paulo viewer drops is absorbed elsewhere. Limited attention binds at the level of a person and of a local market; aggregate global liquidity does not register it.

H_{local} admits a narrower reading. The distraction effect does not appear in any mapped venue we can reach, yet the venues that match Ehrmann–Jansen’s resolution carry a marginal trader who is not the viewer, and the venue with the right trader is daily. We can bound the daily home response and we can say the effect does not show up in CME currency futures or US country ETFs, neither of which is traded by the distracted viewer. The original minute-level effect on a foreign cash exchange stays beyond reach, and our daily home-equity null is silent on an intraday dip that averages out over the session. The powered part of that mapped evidence is essentially the euro, since the other own-currency cells are individually underpowered, so it is best read as one clean currency plus a set of bounds rather than four independent confirmations. What the data license is generalization-bounding, not refutation: the effect, real in its own setting, does not travel to the global and derivative venues that now dominate volume, and its survival at the minute level on a home exchange is left open.

Set beside those two verdicts on the football effect, the composition mechanisms carry a warning that outlasts it. Both the comovement and the full-session specifications are the kind of natural choice an author would make, and both produce the prior literature’s effect from no underlying response. The Monte Carlo shows the comovement case is general: any cross-sectional comovement measure compared across windows that sample different instruments will show a spurious shift, and event windows that differ in their session mix will move any volume measure. Such artifacts survive most readily in an event-study literature that predates routine multiple-testing and composition checks.

13 Threats to validity

Three limits frame the reading. Attention is validated at the daily level through Wikipedia and corroborated at the hour of play by the hourly Wikimedia reconstruction of ten matches across both tournaments and every round (Section 5); finer attention within the during-match minutes themselves is inferred from television audience figures rather than

measured directly. Because the exchanges close on Saturdays, the futures during-match sample tilts toward weekday and Sunday matches, and the cleanest weekend matches are observed mainly in crypto. The local-channel arm, in turn, is daily for the only cell with the right marginal trader, so it bounds the day-level home response and is silent on the intraday minutes the original cash-equity finding concerns. The first two limits qualify H_{global} without undoing it, since the bound holds in venues open during weekend matches. The daily-resolution limit binds on H_{local} , and it is the reason we bound rather than refute.

14 What would change the reading

A reported null is worth only as much as the conditions that would overturn it, so we set those out. Two channels could in principle weaken the H_{global} bound. One is a mis-timing of the during-match windows against the true kickoffs, but the minute-level event study (Figure 5) is flat through the opening whistle on the schedule we use, and the hourly attention reconstruction confirms the windows sit where the attention is. The other is contamination of the matched control, against which the placebo, macro- exclusion, and aggregation checks leave the estimates unchanged. The reading we cannot defend against is the resolution gap on H_{local} . An intraday trade-count dip on a foreign cash exchange, present at the minute and averaged away in our daily home bar, would be invisible to us. It is exactly the Ehrmann–Jansen object, which is why we bound rather than refute it and why the paid intraday pull is the decisive next step. Two things would change the headline if they appeared and have not. A powered own-team cell turning significantly negative would revive H_{local} in a mapped venue; instead the best-powered comparable cell (euro futures, 60 match-days) is positive and survives the bootstrap unchanged. A composition-robust comovement or session measure turning up the prior effect would undercut the artifact claim; instead both vanish the moment the instrument set or the session window is held fixed, and the Monte Carlo shows the naive versions manufacture the effect from nothing. The result that would most cheaply falsify the whole paper, a clean during-match volume decline in any of the eleven liquid global instruments, is not in the data.

15 Conclusion

The World Cup is a clean, repeated, globally synchronized attention shock, and the markets that now carry most of the world’s trading do not visibly respond to it in aggregate. During-match volume, volatility, liquidity, options pricing, and cross-market comovement are flat across crypto, index, commodity and currency futures, with equivalence bounds on volume and the quoted spread that exclude the declines the cash-equity literature reports (the volatility and options channels are reported as non-significant nulls without an equivalence margin), so the aggregate global attention channel is bounded near zero. The structurally comparable local test shows no decline in any mapped venue we can reach, the minute-level effect on a home exchange traded by local viewers is left untested with power, and the next-day loser effect is absent wherever daily data can test it. Both specifications that match the prior result turn out to be sampling which instruments or which session hours, not a market effect. The distraction effect, real in the national cash exchange where it was found, does not generalize to global, derivative and retail venues; whether it survives

at the minute level on its own exchange is the question that paid intraday foreign data would settle.

Funding and disclosure

AlgoSeek provided the author with access to its market data, and the one-minute futures trade bars, quote-bearing TAQ, and end-of-day options-greeks panels used here were obtained from that source under a paid commercial license; the crypto, index, currency, attention, and schedule data are from the public or separately licensed sources named in Section 3. Data access is the entirety of AlgoSeek’s involvement. The research was conceived, designed, and carried out independently: AlgoSeek did not propose or direct the study, had no say in the hypotheses, the analysis, or the decision to publish, and saw no draft before submission. The findings and any errors are the author’s own, and the author has no other financial interest in the outcomes reported.

A Pre-registration

Fixed before the cells were read: the estimating equation (1); clustering by calendar date; the two-hypothesis split of Section 1; log trade count as the primary H_{local} outcome and log volume as the primary H_{global} outcome; the 30% two-one-sided-tests equivalence margin; the rule that any cell on fewer than roughly thirty match-day clusters is read through the wild-cluster restricted bootstrap rather than the asymptotic p ; and the exclusion of the 2026 tournament from every estimate. The H_{local} intraday verdict rule (reject if the pooled own-team effect excludes -30% ; confirm if it is at or below -30% at $p < 0.05$; otherwise report the achieved bound) is registered in `HLOCAL_DATA_SCOPE.md` before any home-exchange data is acquired.

B Reproducibility

The full pull-and-analysis pipeline is public at <https://github.com/DaruFinance/worldcup-attention-markets>; market data is not redistributed (the AlgoSeek panels are commercially licensed; see the funding and disclosure statement) but the scripts document every source and build step. The repository pins its environment (`requirements.txt`), and `run_all.py` regenerates every figure and table in dependency order, skipping any stage whose licensed input is absent so that a non-holder still reproduces the schedule, attention, control, and free-data (Binance crypto and daily home-equity) results; the AlgoSeek-derived futures, equity, quote (TAQ) and options panels are licensed and not redistributed, with their parse scripts provided so a license-holder can rebuild them. The committed inputs are fixed by SHA-256 in `DATA_MANIFEST.sha256`. Every stochastic procedure fixes its seed: the comovement Monte Carlo, the wild-cluster and pooled bootstraps, and the positive-control injection each pin a NumPy generator seed, so the bootstrap p -values, the 100% false-positive rate, and the injection power figures reproduce exactly on re-run. The control series for the loser-effect adjustment (the MSCI-world ACWI and the dollar index) are cached to `data/controls/` so the regressions reproduce bit-for-bit without a network call. A top-level `run_all.py` executes these stages in dependency order, skipping any whose licensed input is absent and

verifying the committed inputs against the manifest (`run_all.py--check`). Core analyses: the cross-market and aggregate-channel tests (`analysis/regress_crossmarket.py`, `regressions.py`), the equivalence and hypothesis split (`equivalence_tost.py`, `hypothesis_split.py`), the comovement artifact and its Monte Carlo (`comovement.py`, `artifact_montecarlo.py`), the home-market and loser-effect arms (`home_equities.py`, `equities_extended.py`, `fx_exotic_extended.py`), the five country and cross-sectional cuts of Section 10 (`country_effects.py`, reusing those panel loaders and emitting `country_effects_summary.csv` and the per-cut files), the wild-cluster bootstrap on the single-market and pooled own-team cells (`wild_bootstrap.py`, `pooled_bootstrap.py`), the positive-control injection of Section 11.2 (`injection_power.py` for the uniform shift and `injection_concentrated.py` for the concentrated-shock variant, emitting `injection_power.csv` and `injection_concentrated.csv`), the four-tournament futures extension of Section 11.3 (`xm_futures_4tourn.py`, with a built-in 2018+2022 validation gate against `xm_main.csv`, emitting `xm_futures_4tourn.csv`; its positive-control injection at the deeper cluster count is `injection_futures_4tourn.py`, emitting `injection_futures_4tourn.csv`), the loser-effect calendar placebo of Section 9 (`loser_calendar_placebo.py`), the goal-window timing robustness (`goal_timing.py`), the \$5 intraday-attention reconstruction of ten matches across both tournaments (`attention_final.py`, emitting `attention_hourly_final.csv` per article and `attention_hourly_matches.csv` per match from the Wikimedia hourly dumps; `attention_hourly.py` is the vendor-slow pooled multi-match version), and the global multiple-testing pool (`master_fdr.py`). The intraday H_{local} pipeline (`ingest/pull_home_intraday.py`, `analysis/hlocal_intraday.py`) and its pre-purchase power calculation (`hlocal_feasibility.py`) are vendor-agnostic and run on a data route when one is supplied.

C Figure and table provenance

Figures are vector PDFs emitted by `figures/make_figures.py` from the committed result CSVs in `analysis/out/`: Figure 2 from `wiki_validation.csv` and `equivalence_bounds.csv`; Figure 3 from `attention_hourly_matches.csv`; Figure 4 from `equivalence_bounds.csv`; Figure 5 from `kickoff_curve.csv`; Figure 6 from `comovement.csv` and `artifact_mc.csv`; Figure 8 from `injection_power.csv`; Figure 7 from the home-equity, extended-index, and extended-FX loser CSVs. Figure 1 is a schematic drawn in TikZ and carries no data. Table 2 is `equivalence_bounds.csv`; Table 3 is `fx_home_country.csv`, `pooled_domestic.csv`, `domestic_attention.csv`, and `home_equities_distraction.csv`; Table 5 is `wild_bootstrap.csv`.

D Results inventory

The pre-registered primary family (during-match volume, trade count, absolute return, and high–low range across the five market groups, the activity: rows of `master_fdr.csv`) holds twenty coefficients, none nominally significant (smallest $p = 0.12$), so the family-restricted Benjamini–Hochberg correction is vacuous. The full multiple-testing pool (`master_fdr.csv`) holds all 190 during-match and own-team coefficients across activity, liquidity, options, efficiency, comovement, intensity, goal windows, sentiment, domestic, home-equity, home-currency, breadth, and country-heterogeneity families;

24 are nominally significant, all outside the primary family; none survives the global correction, and the smallest- p cells are the mechanism-diagnosed artifacts named in Sections 11 and 10. Equivalence bounds for every primary cell are in `equivalence_tost.csv` and `equivalence_bounds.csv`; the two-hypothesis classification with per-cell bound, minimum detectable effect, resolution, and marginal trader is in `hypothesis_split.csv`; the intraday H_{local} feasibility table is `hlocal_feasibility.csv`.

References

- [1] Ehrmann, M., & Jansen, D.-J. (2017). The Pitch Rather Than the Pit: Investor Inattention, Trading Activity, and FIFA World Cup Matches. *Journal of Money, Credit and Banking* 49(4), 807–821. (Earlier as ECB Working Paper 1424, 2012.)
- [2] Edmans, A., García, D., & Norli, Ø. (2007). Sports Sentiment and Stock Returns. *Journal of Finance* 62(4), 1967–1998.
- [3] Kaplanski, G., & Levy, H. (2010). Exploitable Predictable Irrationality: The FIFA World Cup Effect on the U.S. Stock Market. *Journal of Financial and Quantitative Analysis* 45(2), 535–553.
- [4] Bernile, G., & Lyandres, E. (2011). Understanding Investor Sentiment: The Case of Soccer. *Financial Management* 40(2), 357–380.
- [5] Barber, B. M., & Odean, T. (2008). All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies* 21(2), 785–818.
- [6] Da, Z., Engelberg, J., & Gao, P. (2011). In Search of Attention. *Journal of Finance* 66(5), 1461–1499.
- [7] DellaVigna, S., & Pollet, J. M. (2009). Investor Inattention and Friday Earnings Announcements. *Journal of Finance* 64(2), 709–749.
- [8] Hirshleifer, D., Lim, S. S., & Teoh, S. H. (2009). Driven to Distraction: Extraneous Events and Underreaction to Earnings News. *Journal of Finance* 64(5), 2289–2325.
- [9] Croxson, K., & Reade, J. J. (2014). Information and Efficiency: Goal Arrival in Soccer Betting. *Economic Journal* 124(575), 62–91.
- [10] Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society B* 57(1), 289–300.
- [11] Harvey, C. R., Liu, Y., & Zhu, H. (2016). ... and the Cross-Section of Expected Returns. *Review of Financial Studies* 29(1), 5–68.
- [12] Lakens, D. (2017). Equivalence Tests: A Practical Primer for t Tests, Correlations, and Meta-Analyses. *Social Psychological and Personality Science* 8(4), 355–362.
- [13] Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics* 90(3), 414–427.

- [14] Parkinson, M. (1980). The Extreme Value Method for Estimating the Variance of the Rate of Return. *Journal of Business* 53(1), 61–65.