ATMOSPHERIC SCIENCE

Atlantic cyclone intensification across comparable starting conditions

A 42-season comparison of 24-hour wind gains in the Atlantic best track.

Abstract

Rapid intensification is often expressed as a threshold applied to a change in maximum wind. A comparison across eras can therefore depend on the threshold and on where and when the storms being compared begin their intensification windows. We examined NOAA Atlantic best-track records for 1982–2002 and 2003–2023. Stratifying on starting wind, longitude, latitude and season phase retained 6,017 of 6,355 eligible 24-hour windows. At a threshold of 30 knots, 7.57% of early windows and 8.38% of late windows qualified, a difference of +0.81 pp (95% season-bootstrap interval, −1.69 to +3.43 pp). Changes within starting contexts contributed +0.66 pp, and changes in the distribution of starts contributed +0.15 pp. Both intervals included zero. The total point estimate was −0.93 pp at 25 knots and +1.36 pp at 35 knots, with intervals spanning zero at all three thresholds.

Introduction

A tropical cyclone can gain enough wind speed in a day to change landfall hazards in the time available to prepare. Kaplan and DeMaria defined rapid intensification using a 30-knot increase over 24 hours, approximately the upper 5% of the over-water changes in their North Atlantic development sample [1]. That threshold provides a reference point, but the probability of crossing it can shift when the starting wind and locations represented in a sample change.

Best-track analyses offer several ways to study that shift. Bhatia and colleagues studied trends in a per-window rapid-intensification fraction in two archives for 1982–2009, alongside climate simulations [2]. Garner quantified changes in the fastest intensification experienced by each storm [3]. A storm maximum answers a question about a cyclone's most extreme episode. An all-window statistic describes the fraction of recorded starts followed by a large change and gives longer tracks more opportunities to contribute.

Study question

We extend a window-level comparison through 2023 using one Atlantic archive [4, 5]. We ask how the rapid-intensification fraction differs across two equal-length eras and how much of that difference accompanies a shift in the starting conditions. The estimand combines starts from tropical storms and hurricanes that stay in those states at the sampled fixes. Comparisons are restricted to starting contexts represented in both eras, and the same set of contexts is used to examine several thresholds.

The analysis has three parts. First, we locate retained starts and measure how the two upper tails differ. Next, we separate the difference at 30 knots into changes within contexts and shifts in the proportion of starts contributed by those contexts. Finally, we examine how those comparisons respond to changes in window eligibility and the omission of individual seasons. This design keeps the denominator visible at each step.

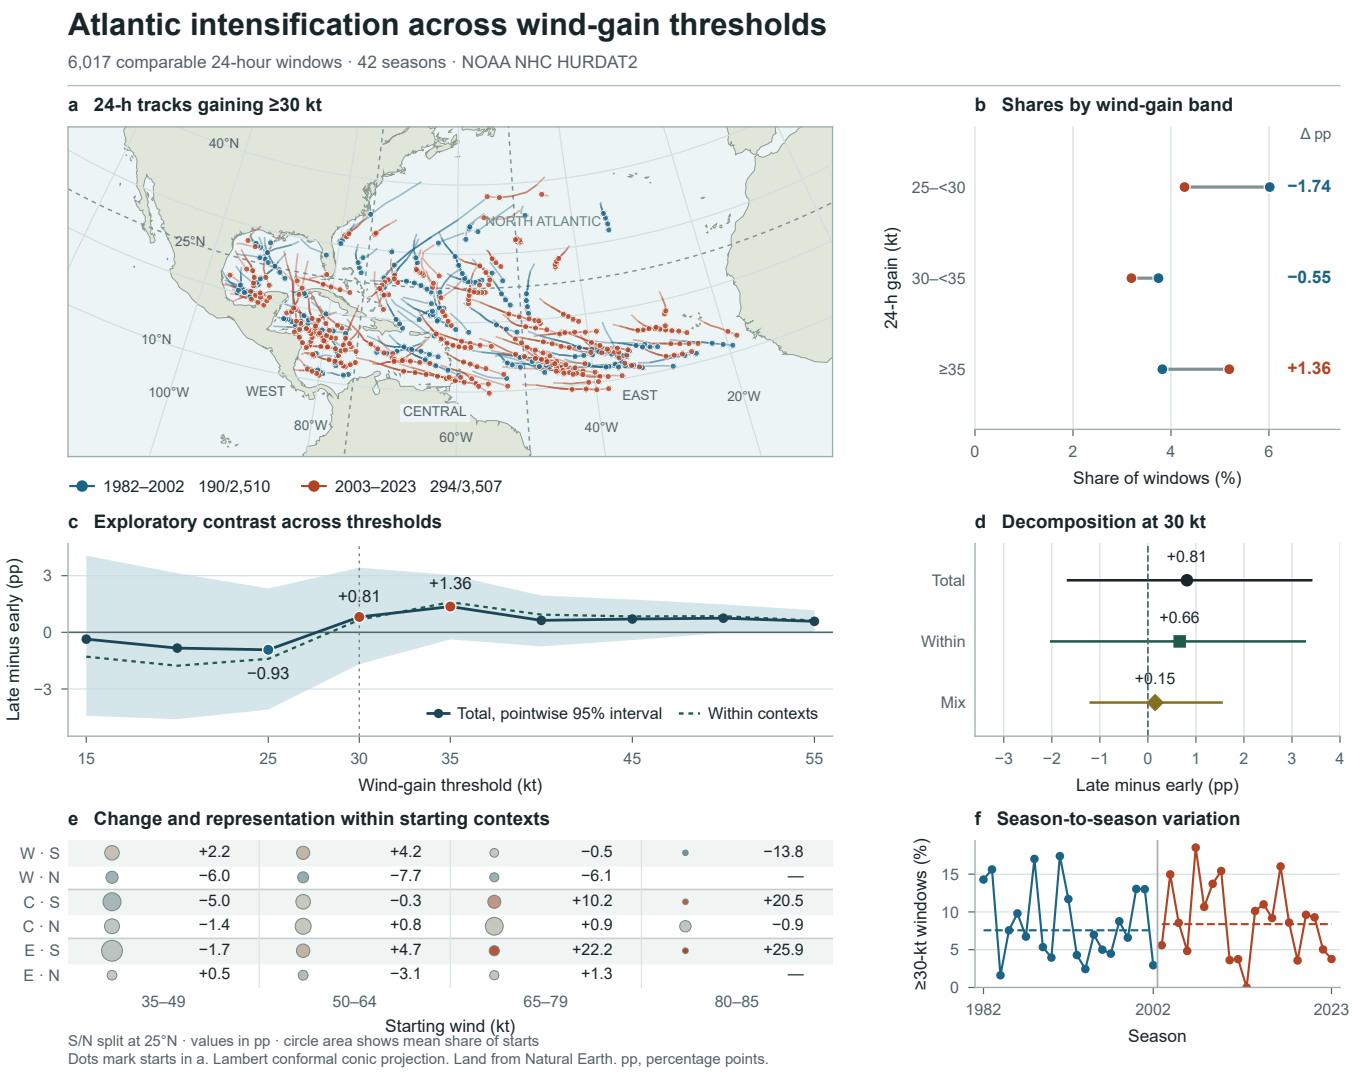


Figure 1. Window-level intensification across 42 Atlantic seasons. **a**, Start markers and tracks joining the five best-track fixes for every window gaining at least 30 knots. Consecutive windows can overlap. The two fractions give the qualifying windows out of all retained starts for each era. **b**, Shares of all retained starts in three disjoint gain bands and their late-minus-early differences. **c**, Total and within-context differences across nine thresholds. Shading joins the pointwise 95% season-bootstrap intervals of the total. **d**, The 30-knot contrast and its two components with 95% intervals. **e**, Context differences in percentage points. Circle area encodes the mean share of starts across eras. W, C and E refer to the longitude bands in a. Dashes indicate missing shared support. **f**, Annual rates and dashed pooled rates by era. pp, percentage points.

Coverage and annual variability

We retained 37 of 48 contexts and 94.7% of the eligible windows. The 2,510 early and 3,507 late windows came from 196 and 277 storms. At least 30 knots of gain occurred in 190 early and 294 late windows (Fig. 1a). Season denominators ranged from 28 to 315, and annual rates varied more than the pooled contrast (Fig. 1f).

The difference within starting contexts

At 30 knots, the difference was +0.81 pp. Changes within contexts contributed +0.66 pp (95% interval, −2.04 to +3.29 pp); shifts in context shares contributed +0.15 pp (−1.22 to +1.56 pp; Fig. 1d). Individual cells varied in direction and share of starts (Fig. 1e). Cell differences are descriptive.

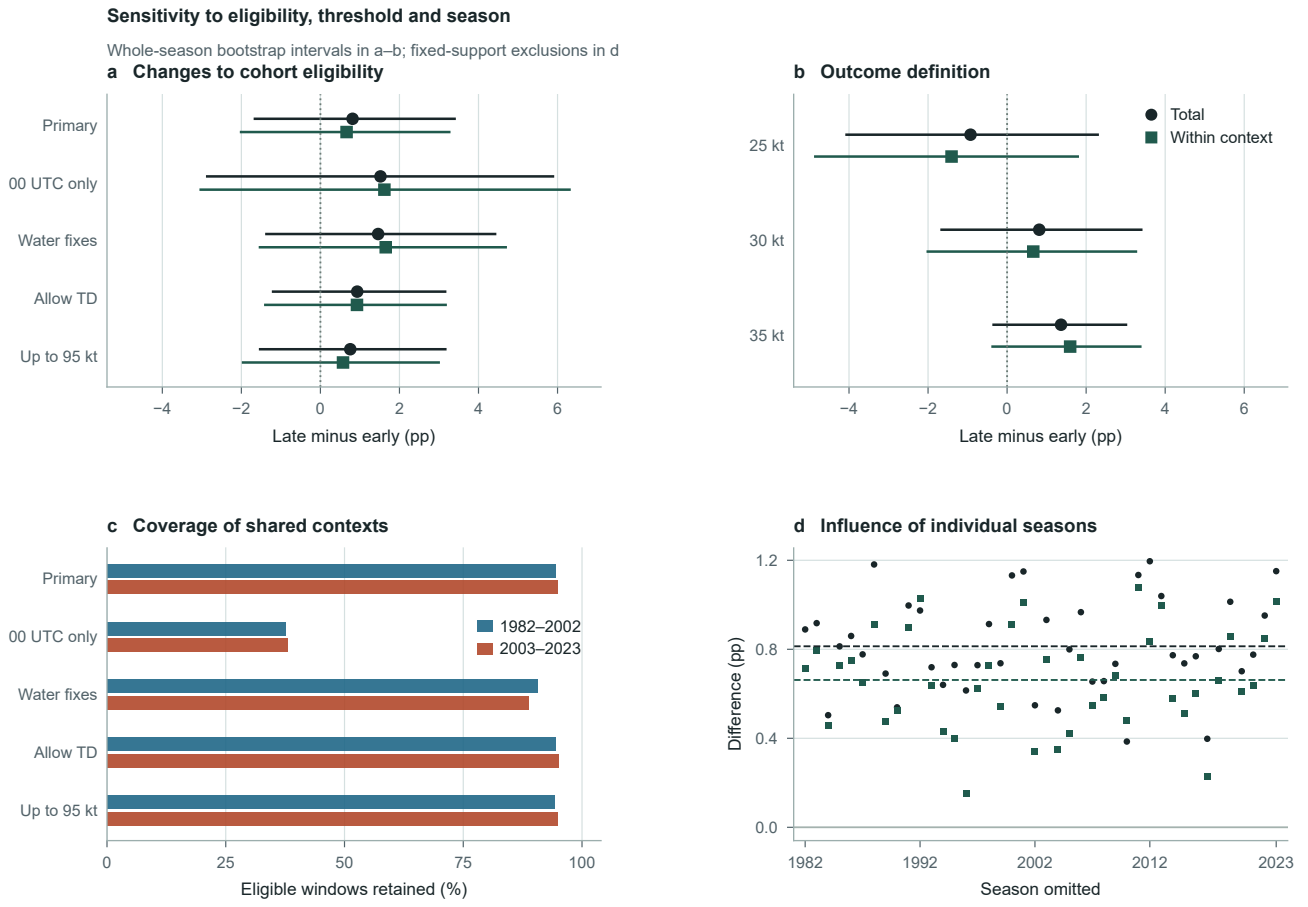


Figure 2. Sensitivity of the era comparison. **a**, Total (circles) and within-context (squares) differences under alternative eligibility rules. Water fixes require that all five recorded coordinates fall outside the land mask. Allow TD permits tropical depression status after the start. Each alternative recomputes support. **b**, Change in the point estimates when the required 24-hour wind gain changes, using the primary cohort. Lines in a and b show 95% season-bootstrap intervals. **c**, Fraction of eligible windows retained under each support rule. **d**, Re-estimation after omitting each season and keeping the primary support set. Dashed lines show the corresponding estimates with all seasons included.

Thresholds give different point estimates

With identical denominators, 341 early and 444 late windows met the 25-knot threshold, giving a total difference of −0.93 pp (−4.09 to +2.32 pp). The counts at 35 knots were 96 and 182, giving +1.36 pp (−0.37 to +3.04 pp). The threshold profile in Fig. 1c shows that both the total and within-context point estimates change sign. The total intervals at all displayed thresholds include zero.

Eligibility and influential seasons

At 30 knots, all four alternative eligibility rules gave positive point estimates with intervals spanning zero. The 00 UTC restriction reduces overlapping starts and changes which contexts meet support. It gave +1.52 pp (−2.90 to +5.92 pp). Eight contexts retained 246 of 653 early and 344 of 904 late starts, about 38% in both eras. Omitting individual seasons gave total estimates from +0.39 to +1.20 pp and within-context estimates from +0.15 to +1.08 pp.

Methods

Data and 24-hour windows

We used the NOAA National Hurricane Center Atlantic HURDAT2 archive covering 1851–2025, revision dated 23 September 2026 [4, 5]. Storms were assigned to seasons using the year in the HURDAT2 storm identifier. We compared 1982–2002 with 2003–2023. Candidate starts were recorded at 00, 06, 12 and 18 UTC in June through November, at $5° \le \phi < 40°$N and 100°W inclusive to 15°W exclusive. We required recorded wind at all five consecutive six-hour fixes and tropical storm (TS) or hurricane (HU) status at each fix. Start wind was 35–85 knots inclusive. The outcome was the final-minus-initial maximum sustained wind, with a threshold of at least 30 knots in the main analysis. Each eligible start counted once.

Shared starting contexts

We crossed four starting wind bins (35–49, 50–64, 65–79 and 80–85 knots) with three longitude bands (W, west of 75°W; C, 75°W to west of 50°W; E, 50°W and farther east); two latitude bands (south of 25°N and 25°N or farther north); and two phases of the season (June–August and September–November). A context entered the comparison if it had at least 20 windows and appearances in at least five seasons of each era. The total rate difference uses the observed era-specific denominators in these retained contexts. For the matrix in Fig. 1e we collapse the two season-phase strata at each displayed wind-by-geography combination using normalized mean shares from the decomposition. A cell can therefore summarize one or two supported strata.

Accounting for shifts in the mix of starts

Let w_{es} be the fraction of retained windows in era e assigned to context s, and r_{es} the corresponding event fraction. We use the symmetric Kitagawa decomposition [6]:

$$\underbrace{\sum_s w_{1s}r_{1s} - \sum_s w_{0s}r_{0s}}_{\text{total difference}} = \underbrace{\sum_s \tfrac{1}{2}(w_{1s} + w_{0s})(r_{1s} - r_{0s})}_{\text{within-context difference}} + \underbrace{\sum_s \tfrac{1}{2}(r_{1s} + r_{0s})(w_{1s} - w_{0s})}_{\text{composition difference}}.$$

Era 0 is 1982–2002 and era 1 is 2003–2023. Multiplying by 100 gives percentage points. The within-context component compares the era-specific event fractions under a common weighting. The composition term holds each context at its mean event fraction while its share of starts changes. We generated 2,400 bootstrap replicates by independently resampling 21 entire seasons per era, preserving within-season dependence from shared storms and overlapping starts. Support was fixed at the observed set. We excluded any replicate with an empty supported context in at least one era. The main analysis retained 2,397 replicates and reports the 2.5th and 97.5th percentiles.

Table 1. Sample retention and total 30-knot contrast by eligibility rule.

Eligibility	Contexts	Retained early	Retained late	Total (pp)	95% interval (pp)
Primary	37	2,510	3,507	+0.81	−1.69, +3.43
00 UTC only	8	246	344	+1.52	−2.90, +5.92
All five fixes over water	31	2,209	2,908	+1.46	−1.40, +4.45
Allow TD after start	37	2,863	3,885	+0.93	−1.23, +3.19
Start wind at most 95 kt	39	2,734	3,753	+0.76	−1.56, +3.19

Alternative definitions and season exclusions

We repeated the analysis at 25 and 35 knots with the same primary windows and support. An exploratory display (Fig. 1c) additionally shows cutoffs every five knots from 15 through 55, with the primary support and bootstrap draws held constant. To evaluate dependence on eligibility, we restricted start hour to 00 UTC, restricted all five sampled fixes to water, allowed tropical depression (TD) status after the start, and expanded the upper starting wind bin to 80–95 knots, one change at a time. We used Natural Earth's 1:50m land polygons [7] to classify the five coordinates and retained windows with no sampled fix on land or its polygon boundary. This mask cannot capture between-fix land crossings. Each alternative re-applied the two support requirements to the subset it defined, so each has its own target cohort (Table 1). For the leave-one-season-out analysis, we retained the original 37 contexts and recomputed the point estimates after removing each season in turn.

Discussion

Among the starts represented in both eras, the mix of observed starting contexts explains little of the 30-knot point estimate. The uncertainty in the within-context component, however, remains large relative to that estimate. The result is compatible with increases or decreases of a few percentage points. It applies to a particular per-window estimand, and can coexist with an increase in per-storm maxima or with differences in the intensity distributions examined in previous work [2, 3]. Conditioning on shared starting contexts makes it possible to identify which part of the observed difference accompanies a change in the distribution of starts. Coarse strata alone cannot account for changes in atmospheric or ocean conditions within them.

Gains from 25 knots up to, but excluding, 35 knots accounted for 245/2,510 early windows (9.76%) and 262/3,507 late windows (7.47%). That 2.29 pp decline accounts for the gap between the two threshold contrasts. Figure 1b attributes 1.74 pp of the decline to gains of 25–< 30 knots and 0.55 pp to gains of 30–< 35 knots. Best-track winds are reported in five-knot increments, so shifting the threshold by one step can include or exclude a substantial fraction of windows. The threshold profile reuses the same windows at every cutoff.

HURDAT2 combines evidence from sources whose availability and uncertainty vary over time, including aircraft and satellites [8]. Estimated intensities near each cutoff inherit uncertainties that the season bootstrap does not measure. The bootstrap treats seasons within each era as exchangeable and cannot quantify low-frequency climate variability. The persistence and common-support rules select windows by observed track histories and represented conditions. The estimates describe that population of starts, with more weight on longer recorded tracks, and are descriptive with respect to physical causes.

Data and code availability

HURDAT2 and the land polygons are public. The PaperSpine source bundle includes immutable input digests, retrieval instructions, the protocol, all analysis and figure code, the resulting summary statistics and this editable manuscript. The study's display and vector exports are available from the PaperSpine research collection.

References

[1] Kaplan, J. & DeMaria, M. Large-scale characteristics of rapidly intensifying tropical cyclones in the North Atlantic basin. *Weather Forecast.* **18**, 1093–1108 (2003). Journal article.

[2] Bhatia, K. T. et al. Recent increases in tropical cyclone intensification rates. *Nat. Commun.* **10**, 635 (2019). doi.

[3] Garner, A. J. Observed increases in North Atlantic tropical cyclone peak intensification rates. *Sci. Rep.* **13**, 16299 (2023). doi.

[4] NOAA National Hurricane Center. Atlantic HURDAT2, 1851–2025 (revision 23 September 2026).

[5] NOAA National Hurricane Center. Atlantic Hurricane Database (HURDAT2) Format (2022).

[6] Kitagawa, E. M. Components of a difference between two rates. *J. Am. Stat. Assoc.* **50**, 1168–1194 (1955). doi.

[7] Natural Earth. 1:50m land, layer 4.1.0, vector release 5.1.2.

[8] Landsea, C. W. & Franklin, J. L. Atlantic hurricane database uncertainty and presentation of a new database format. *Mon. Weather Rev.* **141**, 3576–3592 (2013). doi.