Peer review status:

This is a non-peer-reviewed preprint submitted to EarthArXiv.

On the seasonal predictability of the 2020 North Atlantic tropical cyclone season

E. L. Levin[a] M. Chien,[b] E. A. Barnes,[b, c] H. He,[d] G. A. Vecchi,[d,e] W. Yang,[e]

[a] *Program in Atmospheric and Oceanic Sciences, Princeton University*

[b] *Faculty of Computing and Data Sciences, Boston University*

[c] *Department of Earth and Environment, Boston University*

[d] *High Meadows Environmental Institute, Princeton University*

[e] *Department of Geosciences, Princeton University*

Corresponding author: Emma Levin, emma.levin@princeton.edu

ABSTRACT: The 2020 Atlantic tropical cyclone (TC) season was exceptionally active, producing over twenty named storms, yet several seasonal forecasts failed to predict such extreme activity across their ensemble spread. Even when forced with the observed 2020 sea surface temperatures (SSTs), physics-based models simulated only a moderately active season across their ensemble members. Using observations and statistical, dynamical, and deep learning (DL) models, we evaluate several hypothesis regarding why the observed hyperactive outcome fell outside the ensemble range of the physics-based models forced with observed SSTs. Analysis of observed large-scale conditions indicates that 2020 did not exhibit favorable predictors of hyperactivity, indicating that the moderate activity in the models should not be unexpected. We also find support for a role in subseasonal atmospheric variability in enhancing the 2020 activity relative to predictions based on monthly and seasonal characteristics. To comprehensively characterize the range of outcomes for the 2020 season, we construct a 1,000-member ensemble using an DL emulator forced with observed SSTs. The observed hyperactivity corresponds to a 0.5 percent event within this ensemble. Although highly unlikely in any single year, such an outcome has roughly a 20 percent chance of occurring at least once in a 45-year period. Taken together, our findings support the interpretation that 2020 represents an unlikely but possible outcome, potentially enhanced by subseasonal atmospheric variability, given current understanding, models, and observations. These results serve to remind us that rare events will occur in a chaotic climate system, and large ensembles are one approach to sample them.

SIGNIFICANCE STATEMENT: The 2020 Atlantic tropical cyclone (TC) season was exceptionally active, producing more than twenty named storms, including Hurricanes Laura and Sally, which caused severe impacts. Yet several state-of-the-art weather and climate models, spanning both physics-based and deep learning (DL) approaches, failed to anticipate this heightened activity. Our study evaluates several hypotheses to explain the models' inability to realistically simulate the 2020 season as hyperactive. Our findings indicate that the 2020 season was an unlikely outcome given the large-scale environmental conditions observed that year, rather than a fundamental failure by the modeling or observational systems used to study TCs. We use a novel approach of simulating 1,000 possible 2020 scenarios with a DL-based model to quantify just how unlikely this event was.

1. Introduction

Tropical cyclones (TCs) that form in the North Atlantic Ocean, typically between June and November, pose severe hazards to communities across the United States and the Caribbean (Young and Hsiang 2024; Weinkle et al. 2018; Klotzbach et al. 2018; Pielke et al. 2008). Accurate forecasts are therefore essential for preparedness and risk mitigation. Short-term weather forecasts, issued several days in advance, provide critical guidance for evacuation decisions and emergency response. In contrast, seasonal forecasts, issued months before the onset of the hurricane season, aim to anticipate the overall characteristics of the upcoming season. Such forecasts are valuable to emergency management agencies for resource allocation and logistical planning (Roberts and Wernstedt 2016), to utilities for anticipating changes in demand (Molina and Rudik 2024), and to insurance and real estate sectors for modeling potential losses and adjusting reinsurance or mortgage contracts (Gete et al. 2025).

Seasonal forecasts can be generated by extending pre-season climate anomalies through the forecast period using dynamical models (e.g., Murakami et al. 2025; Vecchi et al. 2014; Vitart and Stockdale 2001; Zhao et al. 2010), statistical models (e.g., Gray 1984; Klotzbach and Gray 2003), hybrid statistical-dynamical models (e.g., Vecchi et al. 2011) or deep learning (DL)-based models (e.g., Zhang et al. 2025a). In addition, retrospective analyses can be performed to investigate the relationship between large-scale environmental conditions and past TC activity. This can be achieved using atmosphere-only dynamical models forced with observed monthly or daily historical sea surface temperatures (SSTs) (Delworth et al. 2012; Zhao and Held 2010; Zhao et al. 2009; Chen

and Lin 2011) or statistical–dynamical models with observed monthly SSTs as predictors (e.g., Vecchi et al. 2011). While these experiments do not reproduce the specific storms that occurred, they enable the generation of plausible realizations of past seasons and provide insight into the factors that contributed to periods of enhanced or reduced activity.

Evaluated over the extended period 1871-2024, two atmosphere-only models developed at the Geophysical Fluid Dynamics Laboratory (GFDL), HIRAM and AM2.5-C360, have demonstrated considerable skill in simulating both interannual and multidecadal variability of Atlantic TC activity when forced with observed monthly SSTs (Levin et al. 2025). Nonetheless, as expected for storm-season simulations (Vecchi and Villarini 2014), these models exhibit biases in certain years.

Since 1980, both models have generally captured the observed interannual variability in Atlantic TC counts, yet they verified poorly against the exceptionally active 2020 season. As shown in Fig. 1a, between 1980 and 2024, the observed storm counts computed using the methodology and duration thresholds of Landsea et al. (2010) typically fall within the ensemble spread of both models. The 2020 season represents a striking exception: while the models simulated a relatively moderate number of storms, with the ensemble means of both models anticipating roughly 8 storms, 23 were observed. The observed activity in 2020 also lay at the extreme upper tail of the ensemble distributions from several coupled seasonal prediction systems (e.g., Murakami et al. 2025) and from a deep-learning-based seasonal forecast (Zhang et al. 2025a). As shown in Fig. 1b and c, the ensemble spread poorly represents the 2020 storm count for more intense storms, including hurricanes and major hurricanes. Although the 2020 season's activity was remarkable, such hyperactivity is not without precedent. For instance, the 2005 Atlantic hurricane season produced 20 named storms, 15 hurricanes, and 7 major hurricanes (Fig. 1). In that case, however, the observed TC, hurricane, and major hurricane counts fell within the ensemble range of the AM2.5-C360 model, and HIRAM also simulated enhanced activity with a local maximum in TC and hurricane frequency.

The origin of the discrepancy between the observed storm count and the model-simulated activity during the 2020 Atlantic hurricane season remains unclear and has not been comprehensively examined. Here we propose and evaluate a set of hypotheses to explain why the 2020 season, which was exceptionally active in the observations, was not captured by the ensemble spread of several state-of-the-art models. We group these hypotheses into two broad families and systematically

4
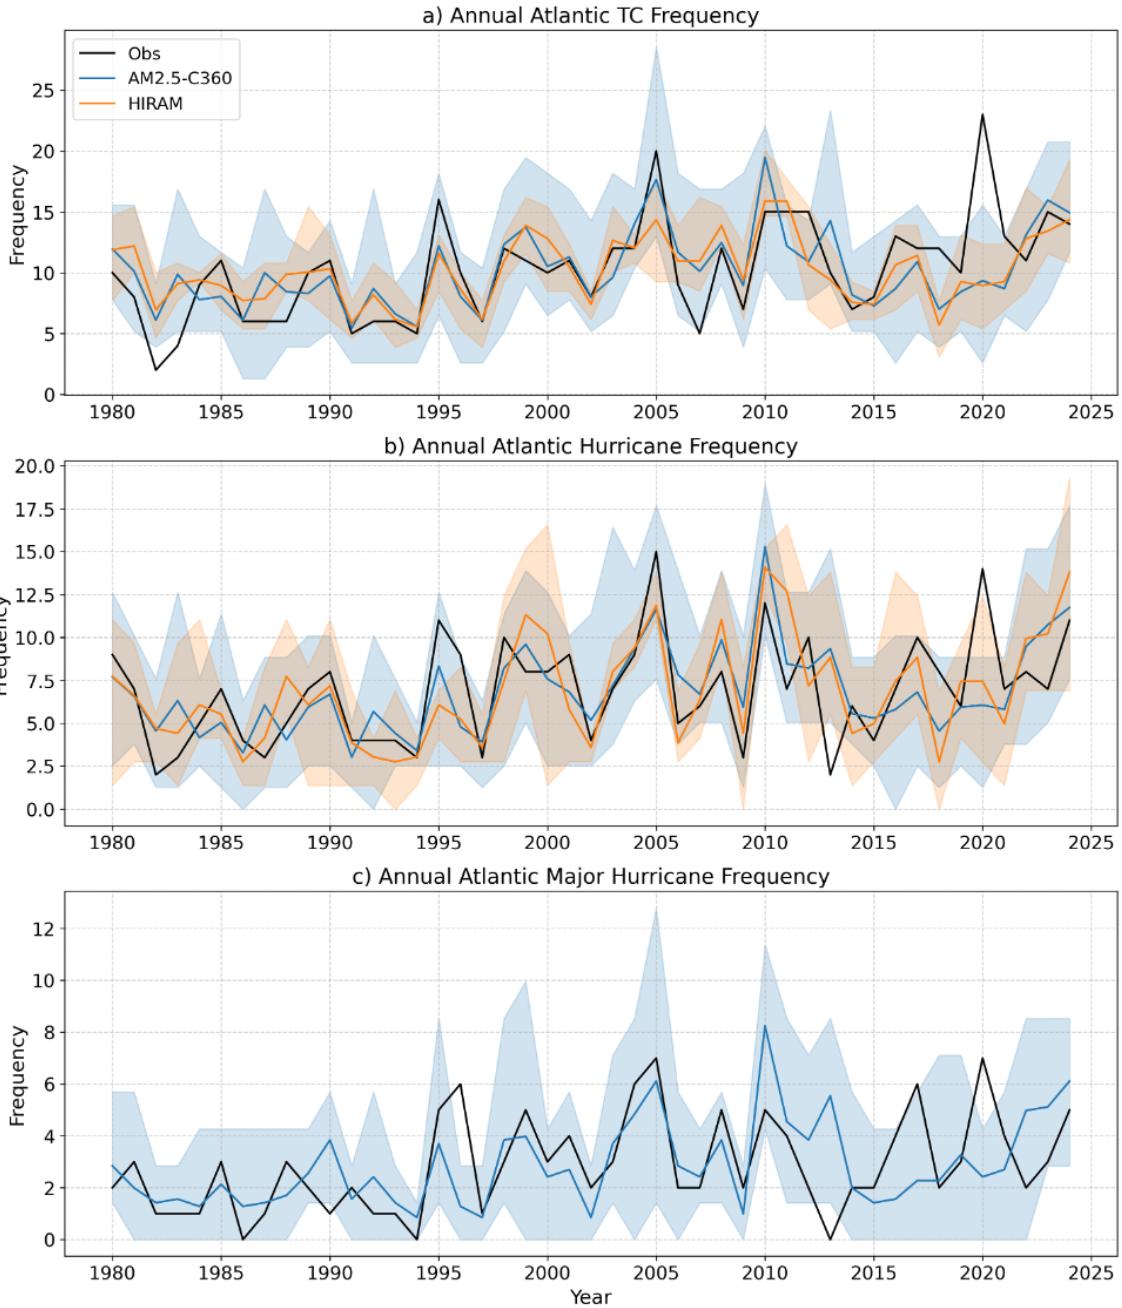
FIG. 1. Annual frequency of (a) TCs, (b) hurricanes, and (c) major hurricanes from 1980–2024. The observational record (black) is computed using the methodology and duration thresholds of Landsea et al. (2010) to estimate the observed TC count in panel (a). Model simulations include AM2.5-C360 forced with observed and bias-corrected SSTs from Chan et al. (2021) (blue) and HIRAM forced with the same SSTs (orange). Shading denotes the full ensemble range (minimum–maximum) for each model (ten members for AM2.5-C360; five members for HIRAM), and solid lines indicate ensemble means. Panel (c) shows only AM2.5-C360 results, as HIRAM does not produce storms of sufficient intensity to be classified as major hurricanes.

test each using a combination of models, theoretical arguments, and observational and reanalysis datasets.

The first family of hypotheses considers whether deficiencies in models or observational datasets contributed to the poor simulation of 2020 Atlantic TC counts. Within this framework, we assess the following possibilities:

1. Atmosphere-only models may exhibit structural limitations in their representation of TC formation and evolution, reducing their reliability in simulating interannual Atlantic TC variability.

2. Observed storm counts in recent decades, as shown in Fig. 1 (black lines), may be inaccurate.

3. Errors may be present in the observed monthly SST datasets used to force the models in 2020.

4. SST-forced models may not adequately represent the influence of observed aerosol changes on TC activity. Two major events in 2020 likely altered aerosol concentrations over the Atlantic basin. First, International Maritime Organization (IMO) regulations on the sulfur content of shipping fuel took effect on 1 January 2020, reducing maximum sulfur content from 3.5% to 0.5%. Implemented to improve air quality, this policy reduced aerosol loading over the oceans and altered regional radiative fluxes (Diamond 2023; Jordan and Henry 2024; Zhang et al. 2025b). Second, the onset of the COVID-19 pandemic led to substantial reductions in global emissions. Previous work suggests that decreased aerosol concentrations over North America and Europe can be associated with enhanced Atlantic TC activity (Murakami 2022, 2024). Because the SST-forced simulations analyzed here employ prescribed CMIP5 aerosol forcings rather than observed 2020 aerosol fields, the combined radiative and dynamical effects of these aerosol perturbations may not be represented.

The second family of hypotheses considers the possibility that the hyperactive 2020 season was a highly unlikely, but dynamically plausible, realization of the internal variability of the climate system. Under this interpretation, the large-scale environmental conditions in 2020 may have permitted a wide range of outcomes, including a low-probability hyperactive season. Subseasonal variability and internal atmospheric fluctuations could have amplified activity beyond the ensemble mean response, leading to an outcome that was not well captured by the models.

6

In the following sections, we apply statistical analyses, dynamical diagnostics, DL–based climate simulations, and a theoretical framework to evaluate these competing explanations. Although we do not identify a single definitive cause for the 2020 discrepancy, we argue subseasonal internal variability could have played a role. The evidence presented in this paper supports the claim that the 2020 season was a low-probability outcome, rather than a structural failure in the models or observational record. At the same time, we cannot exclude the possibility that additional mechanisms not examined in this paper influenced the 2020 Atlantic TC season.

The remainder of the paper is organized as follows. Section 2 describes the data, models, and methodological framework used to test each hypothesis. Section 3 presents the results for both families of hypotheses. Sections 4 and 5 summarize our results and discuss the implications for understanding the predictability of the 2020 Atlantic TC season and for quantifying uncertainty in weather and climate prediction.

2. Methods and data

We employ fully dynamical simulations, a statistical–dynamical model, and a deep learning–based climate model, in conjunction with a theoretical seed–probability framework in observational and reanalysis datasets, to evaluate why the 2020 Atlantic TC season was inadequately represented by the ensemble spread of several modeling systems.

a. Dynamical Models

To place the 2020 Atlantic TC season in the context of recent historical variability, we generate a multi-ensemble record of TC activity over 1980–2024 using two global high-resolution atmospheric models developed at the Geophysical Fluid Dynamics Laboratory: AM2.5-C360 and HIRAM (Zhao et al. 2009). These models share the same dynamical core but differ in their convection parameterizations and horizontal resolutions. AM2.5-C360 uses the relaxed Arakawa–Schubert convection scheme, while HIRAM employs the parameterization developed by Bretherton et al. (2004). HIRAM, with a horizontal resolution of approximately 50 km, has demonstrated strong skill in reproducing historical TC variability (Chen and Lin 2011; Zhao et al. 2009; Zhao and Held 2010). AM2.5-C360, with a finer horizontal resolution of 25 km, has been used in previous studies to examine the influence of SSTs on TC activity (Chan et al. 2021), the seasonal cycle of TCs (Yang

et al. 2021), and the evolution of storm tracks under warming conditions (Kortum et al. 2024). Together, these models provide an independent dynamical framework for assessing the extent to which 2020 lies within the range of SST-forced variability.

To generate a historical record of Atlantic TC counts, we conducted multi-ensemble experiments with both models. We generated ten ensemble members for AM2.5-C360 and five for HIRAM, each forced with bias-corrected observed monthly sea surface temperatures (Chan et al. 2021) from the Hadley Centre Sea Ice and Sea Surface Temperature (HadISST) dataset over the period 1980–2024. Each ensemble member was initialized with distinct atmospheric conditions but used identical historical SST forcing. Aerosol concentrations are prescribed from the Coupled Model Intercomparison Project Phase 5 (CMIP5) projections. From these simulations, we compute annual counts of named TCs, hurricanes, and major hurricanes to construct a modeled range of historical Atlantic TC activity.

To evaluate hypothesis 4 from Section 1, which proposes that the 2020 reduction in atmospheric aerosols was not represented in the SST-forced simulations, we conduct additional experiments. Specifically, we generate a five-member ensemble of AM2.5-C180 (a slightly lower-resolution configuration of AM2.5-C360 with approximately 50-km grid spacing) spanning 1980–2024. In these simulations, the model is forced with observed aerosol fields from the MERRA-2 reanalysis (Gelaro et al. 2017), rather than the CMIP5 aerosol reconstructions used in the historical integrations. This experiment isolates the effect of realistic interannual aerosol variability on modeled TC counts.

As a second approach, because only several ensembles of the 1980–2024 record were conducted with observed aerosols, we perform an idealized modeling experiment to obtain additional samples and better quantify the relationship between sulfate and black carbon aerosols and Atlantic TC activity. Using the HIRAM model, we perform four simulations to isolate the effects of sulfate aerosols, which declined markedly in 2020 following the implementation of the IMO shipping regulation (Diamond 2023; Jordan and Henry 2024; Zhang et al. 2025b), and black carbon, which also decreased in 2020 due to reduced industrial activity and travel during the COVID-19 pandemic (Gupta et al. 2023).

We perform a control experiment (cntl) in which the model is forced by the annual cycle of SSTs averaged over the period 1986–2005, following the methodology of Vecchi et al. (2019).

This climatological SST cycle is repeated for 200 years following a 10-year spin-up, such that the 200-year mean represents the model's steady-state climatological baseline. We then perform three additional experiments using the same SST forcing as the control. In the first sensitivity experiment, sulfate aerosols are removed over the North Atlantic (zeroSO4NA), and the model is integrated for 50 years. In the second experiment, sulfate aerosols are removed globally (zeroSO4global), and the model is again integrated for 50 years. The two zeroSO4 experiments allow us to isolate the effects of reduced sulfates in the oceans, largely brought on in 2020 by the IMO shipping regulations. In the final experiment black carbon emissions are removed globally (zeroBCglobal), and the model is integrated over 50 years. This serves as an exaggerated analogue to the global reduction in black carbon brought on by the COVID-19 pandemic. Our experimental design allows us to directly assess the influence of reduced aerosol loading on Atlantic TC activity under otherwise identical boundary conditions, providing an idealized parallel for the aerosol reductions that occurred in 2020.

To identify TCs in the AM2.5-C360 simulations, we use a tracking algorithm described by Harris et al. (2016) that applies thresholds for wind speed, minimum sea level pressure, storm lifetime, and warm-core characteristics as in Chan et al. (2021). The algorithm operates on 6-hourly fields of sea level pressure, 850 hPa relative vorticity, 10 m wind speed, and mid-tropospheric (300–500 hPa) temperature. The detection procedure begins by locating local minima in sea level pressure and then applying an 850 hPa vorticity threshold of 1.5×10^{-4} s^{-1} to exclude weak or poorly organized systems. For the remaining candidates, additional criteria ensure that only robust and long-lived storms are retained. Each storm must have a minimum lifetime of 72 hours (Villarini et al. 2011a), maintain a warm core for at least 48 hours (defined as a local maximum in 300–500 hPa temperature enclosed by a 2° C contour within 500 km of the minimum sea level pressure), and exhibit at least 36 consecutive hours during which both the warm core and maximum 10 m winds exceed 15.75 m s^{-1}. In addition, at least one time step along the storm's trajectory must feature maximum winds exceeding 17 m s^{-1}.

For the HIRAM model, we implemented minor modifications to the algorithm to ensure that the tracked TC frequency in the reference experiment closely matches the globally averaged TC frequency. These are the same adjustments for HIRAM are used in Yang et al. (2021). The warm-core temperature contour is increased from 2° C to 2.5° C, and the required proximity of the

9

warm core to the storm center is reduced from 500 km to 110 km. All other criteria are identical to those used for AM2.5-C360.

Hurricanes are defined in both models as storms that attain maximum winds exceeding 33 m s^{-1} at least once during their lifetime. Major hurricanes are identified in AM2.5-C360 as storms whose maximum winds exceed 45 m s^{-1}, a threshold slightly lower than the observed 50 m s^{-1} standard because the model rarely produces storms at the observed major hurricane intensity. HIRAM does not reliably simulate storms of sufficient intensity to meet this major hurricane threshold, and therefore major hurricane statistics are reported only for AM2.5-C360.

b. Statistical-Dynamical Model

To characterize the range of plausible seasonal Atlantic TC counts over the recent historical period (1980–2024), including 2020, we employ the statistical–dynamical model of Vecchi et al. (2011). This framework estimates the expected annual Atlantic TC count, λ, using a Poisson regression in which λ depends solely on SST. λ is modeled as a function of the mean seasonal SST anomaly in the Atlantic main development region (MDR) and the mean seasonal SST anomaly across the tropics, since these two predictors together have been shown to be a robust indicator of Atlantic TC activity in the present climate (Vecchi and Soden 2007; Villarini et al. 2011b; Eusebi et al. 2025) and throughout the last millennium (Yang et al. 2024). Specifically, the logarithm of λ is modeled as the following linear function:

$$\lambda = \exp\left(1.707 + 1.388, SST_{MDR} - 1.521, SST_{TROP}\right), \tag{1}$$

where SST_{MDR} and SST_{TROP} are defined relative to the 1982–2005 climatology. The MDR is defined as 10°–25°N, 80°–20°W, and the tropical mean spans 30°S–30°N.

This formulation provides a full probabilistic distribution of seasonal outcomes conditional on the SST state, assuming the seasonal TC distribution follows a Poisson distribution. Thus, for each year in 1980–2024, the model yields a distribution of possible TC counts. We use this distribution to quantify the percentile rank of the observed 2020 season relative to SST-forced expectations. This approach allows us to assess whether the 23 observed storms in 2020 constitute an extreme but statistically plausible realization under SST forcing.

We further use this model to evaluate hypothesis 3 from Section 1, which posits that discrepancies between observed and simulated 2020 TC counts may arise from errors or inconsistencies in the observed SST datasets used to force atmosphere-only models. The AM2.5-C360 and HIRAM simulations analyzed here (Section a) are forced with a corrected version of the HadISST dataset (Chan et al. 2021). To assess the sensitivity of the inferred TC distribution to the choice of SST product, we apply the statistical–dynamical model separately using three independent SST datasets: the Optimum Interpolation Sea Surface Temperature dataset (OISST; Huang et al. (2021)), the Hadley Centre Global Sea Ice and Sea Surface Temperature dataset (HadISST; Schneider et al. (2013)), and the Extended Reconstructed Sea Surface Temperature dataset (ERSST; Huang et al. (2017)).

Comparing the resulting modeled TC distributions across SST products serves two purposes. First, it enables us to evaluate the robustness of the estimated 2020 percentile to observational uncertainty in SST. Second, if substantial differences emerge across SST products, this would suggest that biases in the dataset used to force the atmosphere-only models could have contributed to their underestimation of 2020 activity. Conversely, consistency across SST products would indicate that SST dataset choice is unlikely to explain the model–observation discrepancy.

c. DL-based Model

Although the multi-ensemble historical integrations with AM2.5-C360 and HIRAM provide a dynamical range of Atlantic TC outcomes and require substantial computational resources, their ensemble sizes of ten and five members, respectively, are too small to fully characterize the probabilistic distribution of plausible 2020 seasonal outcomes. In particular, such limited ensemble sizes constrain our ability to robustly estimate the tail behavior of the distribution and to quantify the percentile rank at which the observed 2020 season verifies relative to the SST-forced model spread and observed count.

Recent advances in deep-learning weather models have shown substantial promise for simulating TCs. Once trained, these models can be integrated efficiently with significantly lower computational costs compared to traditional dynamical models like AM2.5-C360 and HIRAM (e.g. (Bi et al. 2023; Chen et al. 2023; Lam et al. 2023; Lang et al. 2024; Kochkov et al. 2024). This efficiency enables the generation of large ensembles, providing an opportunity to characterize the full probabilistic

distribution of potential TC season outcomes. In this study, we employ the Ai2 Climate Emulator version 2 (ACE2; Watt-Meyer et al. (2025)) to perform multi-ensemble annual simulations of historical Atlantic TC seasons, including 2020. We use a version of ACE2 trained on the ERA5 reanalysis (Hersbach et al. 2020). ACE2 predicts key atmospheric variables, including temperature, humidity, zonal and meridional winds, precipitation, and surface pressure, at six-hour intervals through autoregressive integration. The model is forced by pentad (five-day-averaged) observed sea surface temperatures from 1979-2007 and daily observed sea surface temperatures from 2007 onwards. The model is also forced by annually varying greenhouse gas concentrations. The ACE2 architecture is based on a spherical Fourier neural operator framework (Bonev et al. 2023), with a horizontal resolution of 1° and eight vertical levels defined on a hybrid sigma–pressure coordinate. Previous studies have demonstrated that ACE2 is numerically stable over simulations spanning hundreds to thousands of years (Watt-Meyer et al. 2025) and can realistically reproduce TC characteristics across subseasonal to interannual timescales (Chien et al. 2025). These properties make ACE2 a robust platform for generating large ensembles of Atlantic TC seasons and for investigating the range of plausible outcomes for extreme events such as the 2020 season.

We generate a 1,000-member SST-forced ensemble for each year during 2005–2020 (including the anomalously active 2005 and 2010 seasons) and for the anomalously inactive 1982 season. Following Chien et al. (2025), ensemble members for a given year are generated by repeatedly integrating the ACE2 model under identically fixed SST boundary forcing. For a given target year (e.g., 2020), observed SSTs from that year are prescribed and repeated cyclically. The model is then integrated autoregressively for 100 consecutive years under these repeating SST boundary conditions. Each simulated year represents one dynamically distinct realization of the same SST-forced year and is treated as an independent ensemble member.

To obtain 1,000 members efficiently, we perform ten parallel integrations ("chunks"), each producing 100 ensemble members via the 100-year repeated-SST integration described above. Distinct initial atmospheric states are used for each chunk. Specifically, we initialize ACE2 from publicly available initial conditions and integrate the model forward for one year using SST boundary conditions from the year preceding the target year (e.g., 2019 SSTs when constructing the 2020 ensemble). The final model state from this spin-up integration is then used to initialize one 100-year repeated-SST chunk. This procedure is repeated to generate ten distinct initial

12

conditions, yielding 10 × 100 = 1,000 ensemble members. Because SST boundary conditions are held fixed within each ensemble while atmospheric initial conditions differ, the resulting spread reflects internal atmospheric variability under identical large-scale SST forcing. This design allows us to quantify the distribution of plausible seasonal TC outcomes conditioned on the observed SST state of a given year.

To detect TCs in ACE2, we follow the methodology of Chien et al. (2025) and Watt-Meyer et al. (2025), utilizing the TempestExtremes framework (Ullrich and Zarzycki 2017). We introduce a minor modification to the original TC tracking procedure. TC centers are identified at locations where sea-level pressure exhibits a local minimum accompanied by a warm core aloft. The original TempestExtremes approach defines a warm core as a region where geopotential heights between 300 and 500 hPa decrease by 58.8 $m^2\,s^{-2}$ within a 6.5° radial distance from the TC center. Because ACE2 does not output 300 hPa geopotential height, we substitute temperature at the third model level (T3) in the hybrid vertical coordinate (corresponding approximately to 250–400 hPa) to diagnose the warm core structure. Specifically, we require that T3 decreases by at least 0.4 K over a 6.5° radial distance from the TC center, which is equivalent to the geopotential criterion of the original TempestExtremes method under hydrostatic balance (see Supplementary Section S1 of Chien et al. (2025) for details).

After identifying TC centers, we link consecutive centers into a single TC track if their positions are within 8° of each other and occur within a 6-hour interval. To minimize spurious detections associated with surface pressure adjustments over land, two additional constraints are applied: (1) the surface elevation at the storm center must remain below 150 m for at least 60 h, and (2) the storm center must remain within 50°S–50°N for at least 60 h. Furthermore, we restrict tropical cyclogenesis events to occur within 30°S–30°N to exclude extratropical systems.

By generating a 1,000-member ensemble for the 2020 season, we obtain a probabilistic distribution of possible TC activity outcomes under the observed boundary conditions. This large ensemble enables us to robustly estimate the percentile rank of the observed 2020 storm count within the modeled distribution and to characterize the tail behavior of seasonal variability. In doing so, we can assess whether the observed hyperactivity represents a statistically rare but dynamically plausible realization, consistent with the second family of hypotheses outlined in Section 1.

13

d. Theoretical Framework

To assess whether the large-scale environmental conditions in 2020 were favorable for a hyperactive Atlantic season, and to evaluate whether the observed activity may represent a low-probability but dynamically plausible outcome consistent with the second family of hypotheses outlined in Section 1, we apply the theoretical framework of Hsieh et al. (2020) to ERA5 reanalysis data (Hersbach et al. 2020) over the period 1980–2024. This framework provides a physically grounded method for linking large-scale environmental variability to annual TC counts. It has previously been used to investigate the seasonal cycle of TCs (Yang et al. 2021), the sensitivity of TC activity to distinct SST patterns (Hsieh et al. 2022), the response of TCs to radiative forcing (Hsieh et al. 2023), and the long-term evolution of Atlantic TC variability (Levin et al. 2025). In these studies, monthly mean environmental fields were used to diagnose year-to-year variations in basin-wide TC activity.

The framework conceptualizes TC formation as a two-stage process involving precursor seed disturbances and subsequent development into fully formed tropical cyclones. Under this formulation, the annual number of North Atlantic TCs, N_{TC}, is modeled as a binomial random variable,

$$N_{TC} \sim \text{Binomial}(N_s, P), \tag{2}$$

where N_s is the total number of first stage rotating seeds present in the basin for a given year, which has units of storm count, and P is the dimensionless basin-aggregated probability that a first-stage seed transitions into a second-stage TC. Consequently, the expected value of N_{TC}, which has dimensions of storm count, is given by

$$N_{TC} = N_s \times P. \tag{3}$$

For instance, if a given season produces 75 seeds ($N_s = 75$), of which 15 transition into TCs ($N_{TC=15}$), the basin-aggregated transition probability from the first stage to the second stage for that season would be $P = 15/75 = 0.2$.

Hsieh et al. (2020)'s ansatz parametrizes N_s and P as functions of large-scale environmental conditions. First, they develop a proxy for N_s, known as the seed propensity index (SPI):

$$N_s \approx SPI = (-\omega) \cdot \frac{1}{1 + Z^{-1/\sigma}}, \tag{4}$$

where ω is the monthly mean 500 hPa vertical velocity in pressure coordinates (in units of Pa s^{-1} where $\omega > 0$ is for downward motion), $\sigma = 0.69$ is a constant nondimensional fitting parameter, and Z is nondimensional parameter that represents the ability of the low-level vorticity to spinup and is a function of low level vorticity (Ikehata and Satoh (2021)),

$$Z = \frac{f + \zeta}{\sqrt{|\beta + \partial_y \zeta| U}}, \tag{5}$$

where f represents the Coriolis parameter and β represents its meridional gradient, ζ is the monthly mean relative vorticity at 850 hPa, and U is assumed to be a constant wind speed of 20 m/s which is empirically fit using aquaplanet model simulations.

Next, the probability that a weakly rotating seed develops into a strongly rotating TC (P; a dimensionless value between 0 and 1) is parameterized as a probability index ($P(\Lambda)$):

$$P \approx P(\Lambda) = \frac{1}{1 + (\Lambda_0/\Lambda)^{1/\gamma}}, \tag{6}$$

where $\Lambda_0 = 0.014$ and $\gamma = -0.9$ are constant dimensionless fitting parameters, and Λ is the ventilation index defined by Tang and Emanuel (2010) and Tang and Emanuel (2012), which measures the degree to which the influx and circulation of cold dry air into the storm's convective plume can inhibit the storm's strength. The ventilation index is a non-dimensional metric:

$$\Lambda = \frac{v_s \cdot \chi}{PI}, \tag{7}$$

where v_s is the vertical wind shear between 850 hPa and 250 hPa in units of m s^{-1}, PI is potential intensity (the theoretical upper limit of a TC's wind speed based on temperature contrasts between the sea surface and upper troposphere in units of m s^{-1}), and χ is a dimensionless parameter

15

representing moist entropy deficit. This parameter is defined as

$$\chi = \frac{s_m^* - s_m}{s_0^* - s_b},\tag{8}$$

where s_m^* is the saturation moist entropies at 600 hPa in the inner core of a TC, s_m is the environmental entropy at 600 hPA, s_0^* are the saturation moist entropy at the sea surface, and s_b is the entropy of the boundary layer. The numerator of the ventilation index equation (7) represents the difference in midlevel entropy between the TC and its environment, while the denominator represents the air-sea disequilibrium.

Equations (4) and (6) serve as proxies for N_s and P expressed as functions of large-scale environmental factors that can be derived from climate model simulations. Thus, we can approximate the annual TC count using large-scale environemental factors as

$$N_{TC} \approx SPI \times P(\Lambda).\tag{9}$$

In this study, we extend the existing framework to investigate both seasonal and subseasonal variability by incorporating monthly mean and daily environmental data from ERA5, averaged over the tropical Atlantic (10°–30° N) during June through November, to derive a proxy for TC activity. Basin-wide means are calculated over a region extending beyond the Main Development Region (10°–25° N, 80°–20° W) to include the Gulf of Mexico, in order to account for the high concentration of storms that occurred there in 2020 (Fig. B1). Including the Gulf of Mexico ensures that differences in TC track locations among years are appropriately represented in the analysis. To our knowledge, this TC proxy ($N_{TC} \approx SPI \times P(\Lambda)$) has not been computed on subseasonal timescales to study subseasonal TC variability.

We also construct an alternative proxy for annual Atlantic TC counts that explicitly incorporates tracked precursor seeds. In this formulation, annual TC activity is approximated as

$$N_{TC} \approx N_s \times P(\Lambda),\tag{10}$$

where $P(\Lambda)$ (Equation (6)) represents the large-scale, basin-aggregated probability that a seed disturbance transitions into a TC, and N_s denotes the observed annual count of explicitly tracked

16

TC seeds. This formulation differs from the full proxy given in Equation 9, in which seed counts are estimated indirectly using the SPI (Equation 4). Here, seed frequency is diagnosed directly from reanalysis data, allowing us to more explicitly separate variability in seed frequency from variability in development probability.

To detect seeds in the ERA5 data we follow the methods described in Yang et al. (2021). Candidate points are identified as local pressure minima that exceed a vorticity threshold of 1.5×10^{-4} s^{-1}. Each seed must cover a radius of at least 50 km and reach a maximum 850 hPa relative vorticity of at least 4×10^{-4} s^{-1} during its lifetime.

3. Results

a. Testing the first hypothesis family

In this section, we present the results to evaluate the first family of hypotheses as to why the models poorly represented the hyperactive 2020 Atlantic TC season: whether there were fundamental deficiencies in the models or observational datasets.

1) HYPOTHESIS NUMBER ONE

This hypothesis poses that the dynamical atmosphere-only models used in this study to generate the historical TC record in Figure 1 (AM2.5-C360 and HIRAM) may be limited in their ability to represent TC genesis, and thus, are unreliable to assess Atlantic TC interannual variability.

This explanation appears unlikely. In addition to the demonstrated skill of AM2.5-C360 and HIRAM in reproducing the long-term historical variability of Atlantic TC activity (Levin et al. 2025), the well-established influence of SSTs on TC frequency and intensity (Chan et al. 2021), and the accurate simulation of the annual TC cycle (Yang et al. 2021), it is notable that the anomalous 2020 season also lies at or beyond the ensemble range in other independent modeling systems. Seasonal forecasts from physics-based coupled models such as GFDL SPEAR and FLOR (Murakami et al. 2025), as well as from DL-based models such as NeuralGCM (Zhang et al. 2025a), similarly place 2020 at the edge or outside their ensemble distributions. The fact that multiple, structurally distinct modeling frameworks fail to encompass the observed 2020 hyperactivity within their ensemble spread suggests that a model-specific deficiency is unlikely to be the sole explanation.

17

2) Hypothesis Number Two

This hypothesis posits that the observed storm counts in recent decades, shown in Fig. 1, may be inaccurate. Although it is not possible to retrospectively reanalyze past seasons to independently verify every aspect of the historical record, confidence in TC observations during the modern era is high. Since the 1960s, operational monitoring has relied on polar-orbiting and geostationary satellites, supplemented beginning in the 1980s by routine aircraft reconnaissance. In more recent decades, scatterometers and microwave radiometers have further improved the detection and characterization of TCs (see Supplemental Fig. SB1 of Klotzbach et al. (2022)).

Our analysis focuses on the post-1980 period shown in Fig. 1, when these modern observing systems were consistently in place. The launch of advanced geostationary satellites, including GOES-16 and GOES-17 beginning in 2016, further enhanced spatial and temporal resolution, increasing confidence in the completeness and intensity estimates of recent storms, including those in 2020. Although observational uncertainty is larger prior to the satellite era and may be somewhat higher before 2016 relative to the most recent years, the modern record is widely regarded as reliable for basin-wide TC counts.

Moreover, Figs. 1b and c show that the 2020 season was also anomalously active in terms of hurricanes and major hurricanes. These more intense systems are less susceptible to detection ambiguity than TCs, since high wind speeds, organized structure, and satellite signatures are more clearly identifiable. Given the robustness of the modern observing network and the particularly clear identification of intense storms, it is unlikely that errors in the observed 2020 record can explain the discrepancy between the observed activity and the ensemble range of the model simulations.

3) Hypothesis Number Three

The third hypothesis proposes that the discrepancy between the observed and simulated 2020 Atlantic TC season ensemble spread arises from errors in the SST dataset used to force the atmosphere-only models (HadISST). To evaluate this possibility, we apply the statistical–dynamical model of Vecchi et al. (2011), which uses seasonal SST anomalies as the sole predictor of annual Atlantic TC activity. We generate modeled seasonal TC distributions for 1980–2024 using three independent SST products as input: OISST, HadISST, and ERSST.

The resulting SST-based seasonal TC activity statistical reconstructions are shown in Fig. A1. Across the historical period, the model captures a substantial fraction of the observed interannual variability, with correlations between modeled and observed TC counts ranging from $r = 0.66$ to $r = 0.70$ across the three SST products. Historically active seasons such as 1995, 2005, and 2010, as well as relatively inactive seasons such as 1982 and 2009, fall within the 90% confidence interval of the modeled Poisson distribution.

For the 2020 season, two key findings emerge. First, although the three SST datasets exhibit differences in global mean trends and regional structure (Menemenlis et al. 2025), their predictions for 2020 are nearly identical when used in the statistical–dynamical TC model. The mean predicted TC count ranges from 10.9 to 11.8 storms across the three products, demonstrating close agreement. Based on this consistency, we have no evidence to support the presence of a fundamental error in the HadISST dataset, which was used to force the dynamical simulations. Second, all three SST products indicate that 2020 should have been only moderately active. In each case, the observed TC count lies well above the 90% confidence bound of the modeled Poisson distribution. Thus, even when accounting for uncertainty across independent SST datasets, seasonal mean SST conditions do not support an expectation of hyperactive activity comparable to what was observed.

Taken together, the agreement among the three SST products and their consistent underestimation of 2020 activity suggest that errors in the SST dataset used to force the atmosphere-only models are unlikely to explain the large discrepancy between the observed and simulated TC counts.

4) Hypothesis Number Four

Our fourth hypothesis posits that the failure of the observed 2020 TC count to fall within the ensemble range of the AM2.5-C360 and HIRAM simulations may reflect limitations of SST-forced, atmosphere-only models in capturing the impacts of the abrupt aerosol perturbations that occurred that year. In 2020, aerosol concentrations declined in association with international shipping regulations and the onset of the COVID-19 pandemic. If these changes altered the large-scale environment in a manner favorable to TC development, but were not adequately represented in the model forcing, they could have contributed to the model–observation discrepancy.

To evaluate this hypothesis, we employ two complementary approaches. First, we perform a five-member AM2.5-C360 ensemble integration over 1980–2024 using observed aerosol concentrations

19

derived from reanalysis, in contrast to the historical simulations shown in Fig. 1, which use prescribed CMIP5 aerosol forcing. Second, we conduct idealized sensitivity experiments in which sulfate and black carbon aerosols are substantially reduced to emulate an exaggerated analogue of the 2020 aerosol decline. The experimental design is described in Section 2.

For the first approach, the resulting Atlantic TC counts are shown in Fig. A2 alongside the observed record. The multi-member simulation reproduces interannual variability reasonably well, with a correlation of $r = 0.61$ relative to observations. The model captures several inactive seasons such as 2002 and 2009 and active seasons such as 2005 and 2010 as local minima and maxima, respectively. However, even when forced with observed aerosol concentrations, the model simulates 2020 as a relatively inactive season. The absence of a meaningful increase in simulated TC activity suggests that any radiative or dynamical effects associated with the 2020 aerosol reduction were either small or already reflected indirectly in the observed SST boundary conditions. This result indicates that aerosol forcing alone does not reconcile the simulated and observed TC counts for 2020.

The idealized experiments yield a similar conclusion. The distributions of annual Atlantic TC counts from the control and aerosol-reduction simulations are shown in Fig. A3. The distributions are broadly similar, with all three experiments yielding an average of approximately twelve Atlantic TCs per year. Although the cntl and zeroSO4global simulations exhibit somewhat greater interannual variability than the zeroSO4NA and zeroBC simulations, none of the differences in distribution are statistically significant. A Kolmogorov–Smirnov test comparing each experimental simulation to the control indicates that we cannot reject the null hypothesis that the annual TC counts from the zeroSO4NA and cntl simulations are drawn from the same distribution ($p = 0.74$). Similarly, we cannot reject the null hypothesis that the zeroSO4global and cntl simulations originate from the same distribution ($p = 0.28$), nor can we reject the null hypothesis that the zeroBCglobal and cntl simulations originate from the same distribution ($p = 0.99$). These results suggest that, within this modeling framework, the direct radiative effects of prescribed low-frequency (>monthly) aerosol variations on atmospheric heating do not exert a statistically robust influence on Atlantic TC activity. This conclusion does not preclude potential aerosol influences operating through SST-mediated pathways or synoptic-scale interactions, which are not explicitly represented in these experiments.

20

Taken together, the observed-aerosol ensemble and the idealized sensitivity experiments both indicate that aerosol reductions in 2020 were unlikely to be a primary driver of the discrepancy between observed and simulated TC activity. More broadly, across the first family of hypotheses, we find little support for explanations rooted in modeling or observational errors. Thus, the results presented thus far fail to uncover support for the conclusion that the discrepancy between the modeled and observed Atlantic TCs during the 2020 season arose principally from model or observational errors.

b. Testing the second hypothesis family

In this section, we evaluate the second hypothesis family, which considers the possibility that the hyperactive 2020 season was a highly unlikely but dynamically plausible outcome.

1) Statistical analysis of modeled TC activity

To evaluate this hypothesis, we first examine the historical ensembles from the dynamical AM2.5-C360 and HIRAM models shown in Fig. 1. Rather than treating 2020 in isolation, we quantify how often the observed storm count falls outside the modeled ensemble spread across the entire 45-year record.

Assuming the dynamical models are well calibrated, the observed record can be interpreted as an additional equally likely realization drawn from the same distribution as the model ensemble members. Under this assumption, the AM2.5-C360 system effectively contains 11 realizations per year, comprising 10 model members plus the observational record, while HIRAM contains 6 realizations per year, comprising 5 model members plus the observational record. To see the implication of this assumption, consider a simple analogy. Suppose we place 1 red ball (the observation) and $n-1$ blue balls (the ensemble members) in a line at random. Because all permutations are equally likely, the red ball is equally likely to occupy any of the n positions. The red ball lies at one of the two ends with probability $2/n$. In our application, the "ends" correspond to the smallest or largest value among the n realizations. Thus, if the system is well calibrated, the probability that the observation lies outside the ensemble range in any given year (i.e., is either the

21

minimum or maximum among the n realizations) is

$$P(\text{outside ensemble range}) = \frac{2}{n}.$$

528 For AM2.5-C360, $n = 11$, giving the probability that the observation lies outside the ensemble

529 range for in a given year as

$$\frac{2}{11} \approx 0.18.$$

530 For HIRAM, $n = 6$, giving the probability that the observation lies outside the ensemble range for

531 in a given year as

$$\frac{2}{6} = \frac{1}{3}.$$

532 Treating each of the 45 seasons from 1980-2024 as independent trials with probability $p = 2/n$,

533 the total number of years in which the observation falls outside the ensemble spread follows a

534 binomial distribution:

$$X \sim \text{Binomial}(45, 2/11)$$

535 for AM2.5-C360 and

$$X \sim \text{Binomial}(45, 1/3)$$

536 for HIRAM.

537 The expected values, or expected number of instances when the observational record lies outside

538 the ensemble range during the historical period, are therefore

$$\mathbb{E}[X] = 45 \times \frac{2}{11} \approx 8$$

539 for AM2.5-C360 and

$$\mathbb{E}[X] = 45 \times \frac{1}{3} = 15$$

540 for HIRAM.

541 In Fig. 1, the observed record lies outside the AM2.5-C360 ensemble spread 6 times for total TCs,

542 3 times for hurricanes, and 5 times for major hurricanes. Under the binomial model with $p = 2/11$,

543 these outcomes are well within sampling variability and are therefore consistent with a calibrated

system. Similarly, for HIRAM, the observed record lies outside the ensemble spread 15 times forx TCs and 12 times for hurricanes, both close or equal to the expectation of 15 exceedances.

Thus, even in a perfectly calibrated system with only five to ten ensemble members, it is expected that the observation will fall outside the ensemble range multiple times over a 45-year period. Although 2020 stands out visually because it lies outside the ensemble spread simultaneously for TCs, hurricanes, and major hurricanes, similar exceedances occur elsewhere in the historical record (e.g., 2013 and 1996). When viewed in this broader probabilistic context, 2020 does not constitute a uniquely unprecedented departure from model expectations but rather a low-probability realization that is statistically consistent with a small finite-ensemble sampling variability.

At the same time, the limited ensemble sizes of AM2.5-C360 and HIRAM constrain our ability to precisely estimate tail probabilities. With only ten and five members, respectively, the ensemble spread provides a coarse sampling of the underlying distribution, and the frequency of exceedances is sensitive to small-sample variability. For this reason, while the dynamical ensembles suggest that occasional out-of-range seasons are expected, they do not allow a precise quantification of how unlikely a hyperactive season such as 2020 may have been. This limitation motivates the use of the ACE2 DL–based model to generate a much larger ensemble of seasonal TC counts and more robustly characterize the distribution's tail behavior.

To determine the probability of a hyperactive 2020 Atlantic TC season under the observed 2020 SST forcing, we generate a huge 1,000-member ensemble for that year, each representing a plausible realization of the season (Section c). We use the ACE2 DL model due to its demonstrated ability to reproduce key TC characteristics (Chien et al. 2025) and its computational efficiency (Watt-Meyer et al. 2025). Specifically, we generate a 1,000-member ensemble for each year within the 2005-2020 period (which includes predictably active 2005 and 2010 seasons) and the 1982 season, which was anomalously inactive. This approach enables a direct comparison between the 2020 season, during which the observed hyperactive TC count fell outside the ensemble spread of smaller dynamical model ensembles forced with observed SSTs, and other hyperactive seasons (e.g., 2005 and 2010), for which the observed storm counts lay within the ensemble range and coincided with a local maximum in the ensemble mean. This comparison allows us to assess whether the behavior of the DL–based model diverges substantially from that of physics-based atmosphere-only models. This

23

large-ensemble technique also enables us to capture the full spectrum of possible 2020 TC season outcomes under the observed SST forcing.

We include in our analysis several inactive seasons, such as 1982 and 2009, each of which produced eight or fewer TCs. The ensemble spread of both the AM2.5-C360 and HIRAM atmosphere-only models realistically simulated these seasons as inactive. In both cases, the ensemble means of both models exhibited local minima in activity, with ensemble means below nine TCs (Fig. 1). Across both models, no ensemble member exceeded sixteen storms, providing strong confidence in their classification as inactive. We also simulate the 2013 season, which produced ten TCs, only two of which reached hurricane intensity. The ensemble ranges of both AM2.5-C360 and HIRAM encompassed the observed activity of 2013 (Fig. 1), but the observed hurricane and major hurricane counts lie outside the ensemble spread (Fig. 1).

Additionally, our 1,000-member ACE2 ensemble record includes the active 2005 and 2010 seasons, which produced twenty and fifteen storms, respectively. In contrast to 2020, both dynamical models encompassed comparably active outcomes within their ensemble spread for these seasons, and their ensemble means exhibited local maxima in activity.

The observed historical annual TC counts and the distribution of simulated annual TC counts from the 1,000-member historical ACE2 ensemble during the 1982, 2005-2020 period are shown in Figure 2a. ACE2 exhibits a moderate positive correlation between the observed TC counts and the ensemble-mean simulated counts ($r = 0.62$), indicating skill in distinguishing between relatively active and inactive seasons due to SST and carbon dioxide external forcing alone. The ensemble mean peaks during active seasons such as 2005 and 2010, and shows local minima during inactive seasons such as 2009 and 2014, consistent with the observed variability.

Beyond reproducing the observed mean-state variability, the ensemble exhibits substantial inter-annual spread. For each of the 17 seasons shown, the simulated range of TC counts (the difference between the ensemble minimum and maximum) spans roughly 10 to 26 storms. For instance, in 2005 the ensemble indicates that between approximately five and 28 TCs were dynamically plausible, whereas in the relatively inactive 2014 season the simulated range extends from one to 19 TCs. Despite this considerable variability, the observed TC count falls within the ensemble distribution for most seasons in the historical record. Furthermore, the calibration Q–Q plot (Fig. 2b) shows the empirical cumulative distribution function closely tracking the 1:1 line, indicating

24
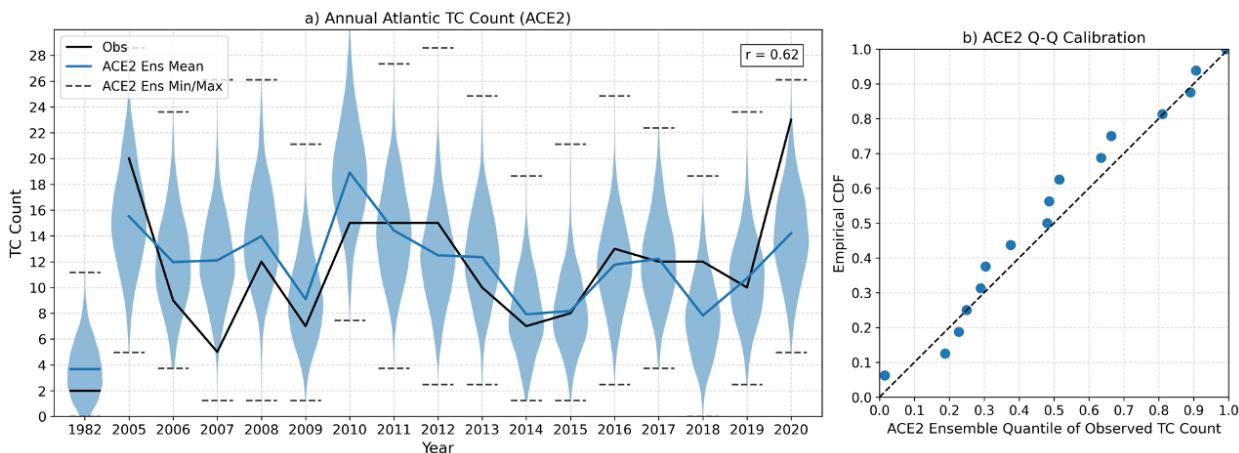
FIG. 2. (a) Annual Atlantic tropical cyclone (TC) counts for 1982 and 2005–2020. The adjusted observational record computed using the methods from Landsea et al. (2010) is shown in black. Blue violin plots depict the distribution of simulated TC counts from the 1,000-member ACE2 ensemble forced with observed SSTs, with the ensemble mean indicated by the solid blue line, and ensemble minimum and maximum indicated by the dashed dark gray lines. The correlation between the ensemble-mean TC frequency and the observed counts is shown in the upper-right corner. (b) Quantile–quantile (Q–Q) calibration diagnostic of ACE2 over the 17 verification years. For each year, we compute the percentile (quantile) at which the observed TC count falls within the 1,000-member ensemble distribution. This yields one quantile value per year (17 total). We then plot the empirical cumulative distribution function (CDF) of these 17 quantiles. If the ensemble is well calibrated, these verification quantiles should be uniformly distributed across years, and the empirical CDF should align with the 1:1 line (black dashed).

that the ACE2 ensemble spread is well calibrated. In other words, the verifying observational quantile is approximately uniformly distributed across years, suggesting that the ensemble dispersion realistically represents seasonal uncertainty.

Although ACE2 generally distinguishes between active and inactive years, visual inspection of Figure 2a indicates that discrepancies between the ensemble mean and observations are the largest prior to 2010 and in 2020, while agreement improves during the 2010–2019 period. This is consistent with the fact that these latter years overlap with the model's training period of 2011-2019 (Watt-Meyer et al. 2025).

25

TABLE 1. Modeled distribution of 2020 seasonal TC counts using a statistical-dynamical model with three SST products as inputs and the DL-based ACE2 1,000-member ensemble model. For each modelled TC count distribution we show the distribution's mean, quantile at which the observed 23 TC counts in 2020 verified against the modeled distribution, and the corresponding percent of such an event.

Model	Model Mean TC count	Verified Quantile	Event Percent
Statistical-Dynamical (HadISST)	11.0	0.999	0.1%
Statistical-Dynamical (ERSST)	11.8	0.997	0.3%
Statistical-Dynamical (OISST)	10.9	0.999	0.1%
ACE2 1,000 member ensemble	14.1	0.995	0.5%

There are several years in which the observed TC count falls near the extreme tails of the ACE2 ensemble distribution. One such case is 2007, an anomalously inactive season with only five observed TCs. The 1,000-member ACE2 ensemble, forced with observed 2007 SSTs, simulated a plausible range of approximately one to 26 storms, with an ensemble mean of 12. Only 15 of the 1,000 members produced five or fewer TCs, indicating that the realized season corresponds to a 1.5% lower-tail event within the modeled distribution. An even more pronounced tail event occurs in 2020. The ensemble simulates a mean of approximately 14 TCs (Table 1), compared to the 23 storms observed. Only 37 ensemble members produce at least 20 storms (3.7%), and just five members simulate 23 or more storms (0.5%). Thus, under the observed 2020 SST forcing, ACE2 characterizes the realized season as a rare but physically plausible outcome, verifying at the 99.5th percentile of the modeled distribution.

To place this result in broader context, Table 1 compares ACE2 with the statistical–dynamical Poisson model described in Eq. 1, forced separately with three SST datasets (HadISST, ERSST, and OISST). All models consistently indicate that the boundary conditions in 2020 were not expected to produce such an exceptional season, with event probabilities ranging from 0.1% to 0.5%.

Although a 0.5% event is highly unlikely in any single year, its occurrence over a multidecadal record is less implausible. Assuming independence between seasons, the probability of observing at least one 0.5% event over a 45-year period (e.g., 1980–2024) is

$$P(\text{At least one } 0.5\% \text{ event}) = 1 - (0.995)^{45} \approx 0.20. \tag{11}$$

TABLE 2. The percentage of 10,000 resampled realizations of annual Atlantic TC counts from the ACE2 model

661 that satisfy specific activity thresholds. Each realization is constructed by randomly selecting either 10 or 15

662 ensemble members from the 1,000 ACE2 simulations of the 2020 season, without replacement, and repeating

663 this sampling process 10,000 times.

Criteria across all samples (%)	10 ensembles	15 ensembles
≥1 ens with at least 20 TCs	32.99%	44.96%
≥2 ens with at least 20 TCs	5.86%	12.06%
≥1 ens with at least 23 TCs	8.03%	11.51%
≥2 ens with at least 23 TCs	0.24%	0.56%
mean > 12 TCs ('active season')	97.96%	99.62%
mean < 12 TCs ('inactive season')	2.04%	0.38%
mean > 17 TCs ('$+1\sigma$ season')	0.65%	0.16%
mean < 7 TCs ('-1σ season')	0.00%	0.00%

644 Thus, an event with a 0.5% annual probability has roughly a 20% chance of occurring at least once

645 in a 45-year historical record. While the 2020 season appears extraordinarily unlikely in isolation,

646 its occurrence within a multidecadal period is not statistically implausible.

647 To evaluate how ensemble size influences the ability to capture the full range of plausible 2020

648 Atlantic TC outcomes, and to demonstrate how reliance on the ensemble mean alone can obscure

649 tail risk, we conduct a resampling experiment using the 1,000-member ACE2 simulation of the

650 2020 season. From this large ensemble, we randomly draw either 10 or 15 members without

651 replacement and repeat this procedure 10,000 times. These sample sizes mimic the ensemble sizes

652 typically used in physics-based seasonal forecasting systems, which often consist of only 10–15

653 members (e.g., Murakami et al. (2025)).

654 For each resampled ensemble, we evaluate whether specific activity thresholds are met (Table

655 2). In particular, we assess the probability that at least one or at least two members simulate (i)

656 a season with ≥20 TCs (comparable to other hyperactive years such as 2005) and (ii) a season

657 with ≥23 TCs, matching the observed 2020 count. We also classify seasons using the ensemble

658 mean. An "active" season is defined as one with a mean ≥12 TCs, consistent with the 2005–2020

659 observed mean (approximately 12 storms).

664 A 10-member (15-member) ensemble includes at least one member with ≥23 TCs in approx-

665 imately 8% (12%) of realizations, suggesting that even with a limited ensemble, there is a non-

negligible chance that a hyperactive scenario would appear within the ensemble spread. At the same time, reliance on the ensemble mean would have masked much of this tail behavior. Nearly all resampled ensembles (98–100%) produce a mean above 12 storms, correctly indicating that 2020 was likely to be active. However, extremely active outcomes are rarely reflected in the mean: fewer than 1% of samples exceed the $+1\sigma$ threshold (≥ 17 storms), and none fall below -1σ.

2) LARGE-SCALE ENVIRONMENTAL CONDITIONS IN 2020

To further evaluate the second family of hypotheses, that the hyperactive 2020 season was a highly unlikely but dynamically plausible realization of the climate system, we examine the large-scale environmental conditions that prevailed during that year. Internal atmospheric variability operating on subseasonal timescales could have amplified activity beyond what would be anticipated from seasonal-mean conditions alone. If so, 2020 should appear only modestly favorable in large-scale seasonal metrics, while exhibiting enhanced sensitivity to subseasonal variability.

To test this idea, we leverage ERA5 reanalysis data in combination with the theoretical TC proxy introduced by Hsieh et al. (2020), defined as $N_{\text{TC}} \approx \text{SPI} \times P(\Lambda)$, which approximates seasonal TC counts as a function of the large-scale environment. We compute and plot the seasonally and tropical Atlantic ($10°$–$30°$ N) basin-averaged Atlantic TC proxy using monthly mean ERA5 inputs for the period 1980–2024 ($\text{SPI} \times P(\Lambda)$, Fig. 3). This framework reproduces a substantial portion of the observed interannual variability in Atlantic TC activity, yielding a correlation of $r = 0.65$ between observed and both proxy-derived TC counts. Thus, seasonal-mean large-scale environmental conditions explain much of the historical variability in Atlantic TC activity.

However, a clear discrepancy emerges in 2020, when the observed TC count substantially exceeds the proxy estimate (blue line in Fig. 3). When the 2020 season is excluded, the correlation between the observed counts and the seasonal $\text{SPI} \times P(\Lambda)$ (blue line) increases to $r = 0.70$. Deviations between the proxy and observations are also evident in the surrounding years, particularly from 2018 to 2021. These results indicate that, based on seasonally averaged large-scale conditions alone, 2020 did not appear exceptionally favorable for hyperactivity.

To determine whether subseasonal atmospheric variability contributed to the inability of the HIRAM and AM2.5-C360 models to capture a hyperactive 2020 TC season within their ensemble range when forced with monthly observed forcing, we further analyze the TC proxy in ERA5 data

28

701   on daily and subseasonal timescales. In addition to computing the TC proxy using exclusively

702   monthly ERA5 input data, we also compute the proxy using daily ERA5 data. For the 2020, 2010,

703   and 2005 seasons, we calculate the monthly product of SPI × P(Λ) (our TC proxy) using four

704   distinct approaches:

705   1. **$\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Daily}}$**: Daily input data are used to compute daily SPI and daily $P(\Lambda)$. Monthly

706   means of these daily values are then computed, multiplied, and averaged over the tropical

707   Atlantic (10°–30° N). This approach uses only daily data.

708   2. **$\text{SPI}_{\text{Monthly}} \times P(\Lambda)_{\text{Monthly}}$**: Monthly input data are used to compute the mean monthly SPI

709   and the mean monthly $P(\Lambda)$. These fields are then multiplied and averaged over the tropical

710   Atlantic (10°–30° N). This approach uses only monthly data.

711   3. **$\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Monthly}}$**: Daily input data are used to compute daily SPI, which is then averaged

712   to monthly means. These values are multiplied by monthly $P(\Lambda)$ computed from monthly

mean input data and averaged over the tropical Atlantic (10°–30° N). This approach combines daily and monthly inputs and isolates the effects of subseasonal changes in SPI.

4. $\text{SPI}_{\textbf{Monthly}} \times P(\Lambda)_{\textbf{Daily}}$: Monthly input data are used to compute monthly SPI. Daily input data are used to compute daily $P(\Lambda)$, which is then averaged to monthly means. The monthly SPI and monthly $P(\Lambda)$ are multiplied and averaged over the tropical Atlantic (10°–30° N). This approach also combines daily and monthly inputs and isolates effects of subseasonal changes in $P(\Lambda)$.

The distinction among the four approaches lies in the order of operations: whether monthly means are taken prior to computing SPI and $P(\Lambda)$, or whether these quantities are first computed daily and then averaged to monthly scales. This framework enables us to assess the influence of subseasonal atmospheric variability on the TC proxy. When approaches (1) and (2) closely agree for a given season, it suggests that subseasonal atmospheric variability plays a limited role, indicating higher seasonal TC predictability from monthly mean boundary conditions.

We calculate the seasonally averaged ratio of approaches (1) and (2) for the 2005, 2010, and 2020 seasons. We select the 2010 and 2005 seasons because, despite their high activity, both were anticipated to be active by the the ensemble spread of the HIRAM and AM2.5-C360 models forced with monthly SSTs (Fig. 1).

The seasonally-averaged ratios are shown in Fig. 4. For both the full TC season (June–November) and the peak period (August–October), the 2020 season exhibits the greatest deviation from one, with the lowest values (0.57 and 0.67), indicating the largest disagreement between approaches (1) and (2) and weakest seasonal TC predictability from monthly mean boundary conditions. In contrast, the ratios for 2010 (0.63 and 0.77) and 2005 (0.64 and 0.74) are closer to one and more similar to each other than to those of 2020, suggesting higher predictability for seasonal TC activity from monthly mean conditions. These results are consistent with the consensus across the ensemble spread of both DL and dynamical models that the 2005 and 2010 seasons would be active. These results also suggest that monthly conditions data alone was insufficient to represent the conditions leading to the exceptional activity of 2020, which helps explain why models using monthly data underestimated that season (Figs. 1) and why the monthly TC proxy, despite capturing interannual variability, failed to reflect the unprecedented activity observed in 2020 (Fig. 3). While this

30

framework helps diagnose differences in seasonal TC predictability with monthly mean boundary conditions, it does not fully explain the physical drivers behind the extreme activity of 2020.

Together, these results reinforce the second family of hypotheses. First, seasonal-mean large-scale conditions in 2020 did not strongly favor a hyperactive season, placing the observed outcome in the upper tail of the distribution of plausible outcomes. Second, 2020 exhibited stronger sensitivity to subseasonal variability than other more predictably active seasons, implying reduced predictability from monthly-mean boundary conditions. In this framework, the hyperactivity of 2020 need not necessarily imply a fundamental model failure. Rather, it could reflect a dynamically plausible but low-probability amplification of activity arising from internal atmospheric variability on weather timescales superimposed on only moderately favorable large-scale conditions.
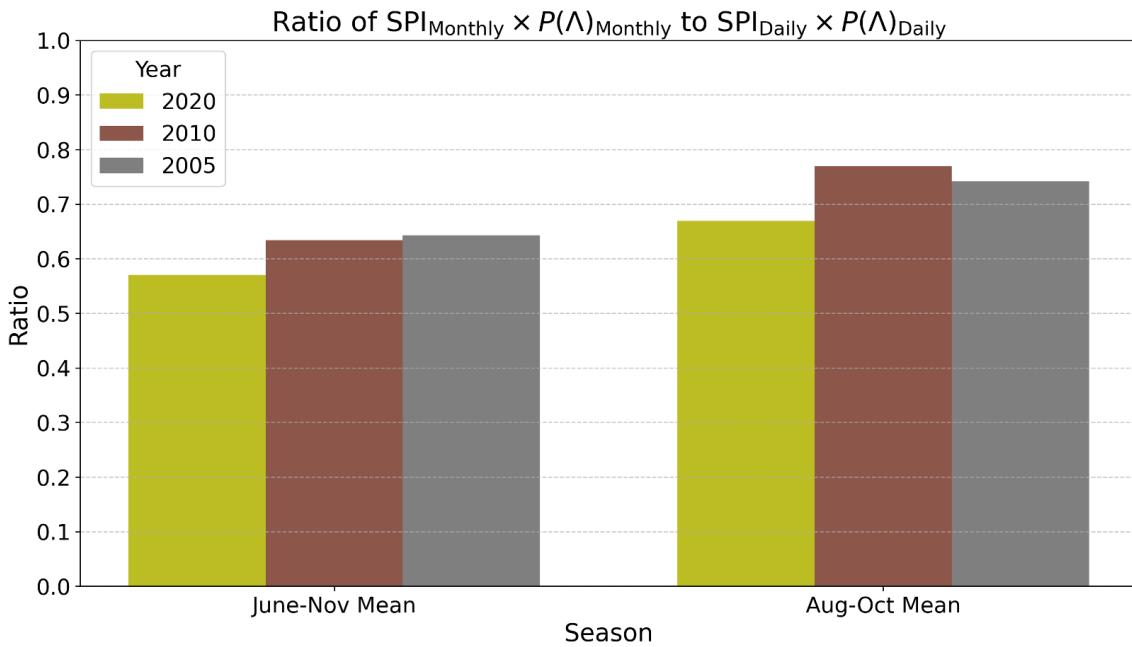
FIG. 4. The seasonally averaged ratio $\text{SPI}_{\text{Monthly}} \times P(\Lambda)_{\text{Monthly}}$ (approach 2) to $\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Daily}}$ (approach 1) is shown for the full TC season of June-November and the most active portion of the season, August-November for 2020 (olive bars), 2010 (brown bars), and 2005 (gray bars). A higher ratio closer to one implies greater predictability. We observe the ratio to be lower in 2020 than 2010 and 2005 for the full season and most active portion of the season.

The monthly and daily TC-proxy calculations for 2005, 2010, and 2020 (Fig. 5) provide further insight into the role of subseasonal variability in shaping seasonal outcomes. To isolate which

component of the proxy (either the seed component (SPI) or the environmental transition component ($P(\Lambda)$)) contributes most strongly to subseasonal sensitivity, we compare two hybrid formulations: $\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Monthly}}$ and $\text{SPI}_{\text{Monthly}} \times P(\Lambda)_{\text{Daily}}$.

Across all three active seasons, the orange (monthly-only) and red (monthly SPI, daily $P(\Lambda)$) curves remain relatively close, whereas the gray curve (daily SPI, monthly $P(\Lambda)$) departs more substantially, particularly during the peak months (July–October). This behavior indicates that subseasonal variability in the seed component (SPI), contributes more strongly to subseasonal variability than subseasonal fluctuations in the environmental transition probability $P(\Lambda)$.
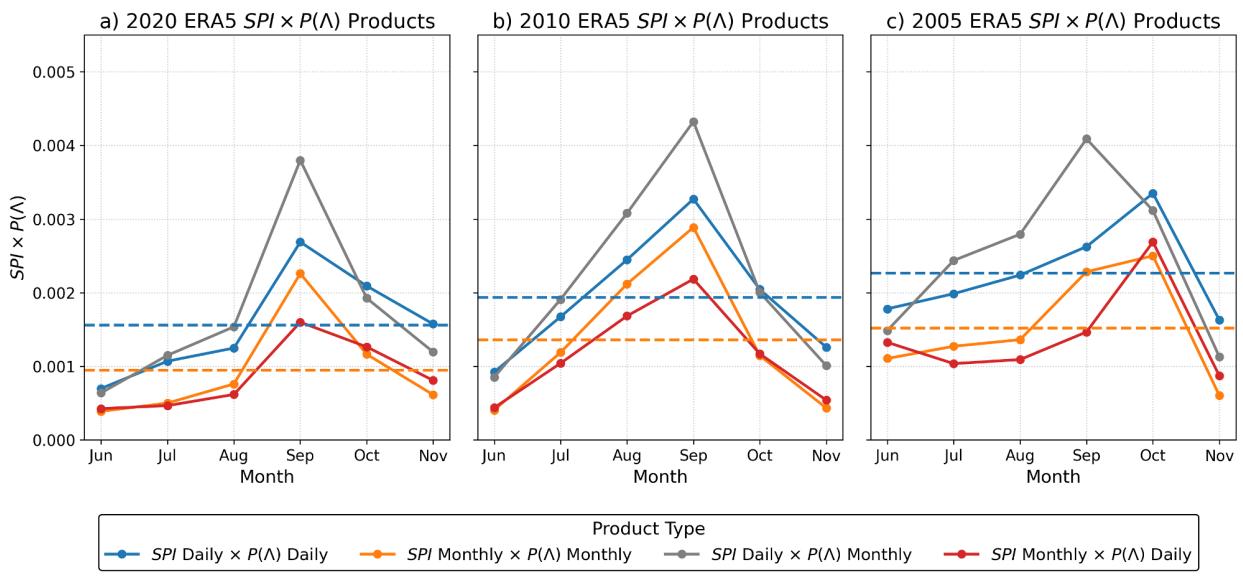
FIG. 5. Monthly mean products of $\text{SPI} \times P(\Lambda)$ computed using four methodological approaches applied to ERA5 data during the Atlantic tropical cyclone season (June–November) for (a) 2020, (b) 2010, and (c) 2005. The blue line corresponds to approach 1, $\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Daily}}$; the orange line corresponds to approach 2, $\text{SPI}_{\text{Monthly}} \times P(\Lambda)_{\text{Monthly}}$; the gray line corresponds to approach 3, $\text{SPI}_{\text{Daily}} \times P(\Lambda)_{\text{Monthly}}$; and the red line corresponds to approach 4, $\text{SPI}_{\text{Monthly}} \times P(\Lambda)_{\text{Daily}}$. The seasonal mean of the blue line, representing a proxy for TC activity derived solely from daily data, is indicated by the horitzontal blue dotted line. The seasonal mean of the orange line, representing a proxy for TC activity derived solely from monthly data, is indicated by the horizontal orange dotted line.

Therefore, we test whether seed disturbances contributed to the hyperactive 2020 TC season emerging as a rare realization within the ensemble spread of multiple models. To do so, we exam-

32

ine explicitly tracked Atlantic TC seeds from three independent ERA5-based datasets (Moon et al. (2025); Vishnu et al. (2020); Ikehata and Satoh (2021)), in addition to our own tracking methodology. We restrict all datasets to storms forming in the North Atlantic during June–November. Each method tracks storms using 6-hourly ERA5. The resulting seed counts are shown in Fig. 6. These approaches differ in methodology and yield different mean seed climatologies, yet they share broad interannual patterns. Notably, none of the datasets identifies 2020 as a local maximum in seed count (Fig. 6). If the hyperactivity of 2020 were driven by an anomalously large number of seeds, we would expect at least one tracking methodology to identify an extreme seed year.

We then incorporate explicitly tracked seeds into the TC proxy, replacing SPI with observed seed count to compute $N_{TC} \approx \text{Seeds} \times P(\Lambda)$, where $P(\Lambda)$ is derived from seasonal monthly ERA5 conditions. This modified proxy reproduces a substantial portion of observed interannual variability (correlations $r = 0.63$–0.69) and successfully captures both inactive and active years such as 2005 and 2010. Yet across all three seed-tracking methodologies, the proxy converges on a moderately active 2020 season of roughly twelve storms, below the observed 23.

Thus, even when incorporating subseasonally tracked seeds and accounting for observed atmospheric variability within the proxy framework, the 2020 atmospheric conditions do not appear conducive to a hyperactive TC season. While subseasonal variability appears more influential in 2020 than in other active seasons, suggesting reduced predictability from monthly mean boundary conditions, even this variability, as captured by the theoretical framework proposed by Hsieh et al. (2020), cannot fully explain the realized activity.

4. Discussion

This study examines the apparent unpredictability of the 2020 Atlantic TC season, during which observed activity across all intensity bins, including TCs, hurricanes, and major hurricanes, exceeded the ensemble spread of state-of-the-art GFDL atmosphere-only climate models forced with observed SSTs (Fig. 1). Seasonal forecasts from both dynamical systems (Murakami et al. 2025) and DL–based models (Zhang et al. 2025a) similarly underestimated the realized storm frequency across their ensemble spreads. Although Klotzbach et al. (2022) suggested that elevated late-season activity in 2020 may have been partially anticipated using large-scale indices, that work focused on intensity metrics such as Accumulated Cyclone Energy and rapid intensification. To
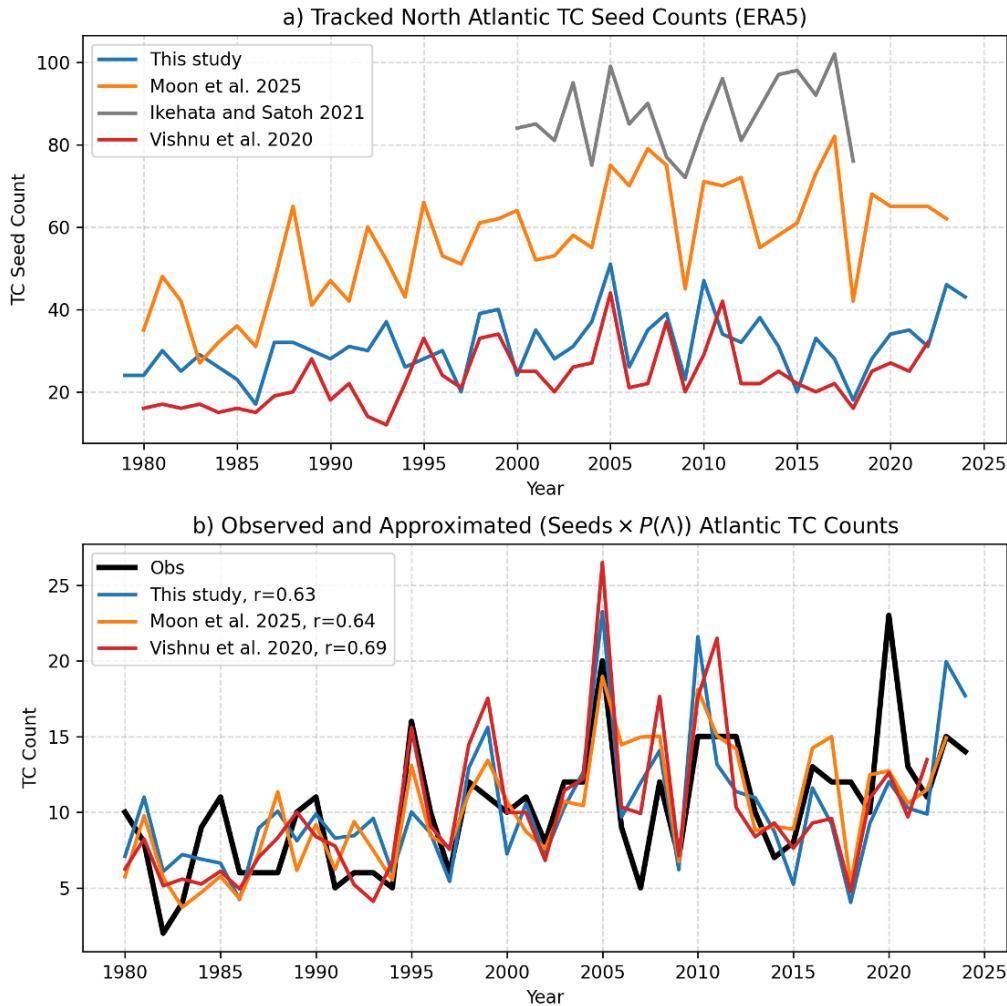
FIG. 6. (a) Tracked Atlantic tropical cyclone (TC) seed counts from the ERA5 dataset using the methods described in Section 2 (blue; 1979–2024), from Moon et al. (2025) (orange; 1980–2023), Ikehata and Satoh (2021) (gray; 2000–2018), and Vishnu et al. (2020) (red; 1980–2022).(b) Observed annual Atlantic TC counts (black) and approximated annual TC frequencies derived from seed counts (Seeds×$P(\Lambda)$), where $P(\Lambda)$ represents the transition probability estimated from seasonally and MDR-averaged monthly ERA5 conditions. Seed counts are obtained following the respective tracking methods in Section 2 (blue; 1980–2024), Moon et al. (2025) (orange; 1980–2023), and Vishnu et al. (2020) (red; 1980–2022). Approximated TC counts are scaled by a multiplicative factor so that their mean matches the observed mean. Correlation coefficients between observed and approximated TC counts for each seed dataset are shown in the upper-left corner of panel (b).

the best of our knowledge, the predictability of storm frequency on seasonal timescales has not previously been examined.

We proposed two broad families of hypotheses to explain the discrepancy between observations and model simulations. The first considers whether model or observational deficiencies prevented the ensembles from capturing a hyperactive outcome comparable to that observed. The second examines whether the hyperactive season represented a dynamically plausible but low-probability outcome given the large-scale conditions present during the season.

After systematically testing four hypotheses within the first family using dynamical and statistical–dynamical models, we find little evidence supporting model or observational failure. Independent SST datasets yield consistent predictions of only moderately active conditions in 2020 using a statistical-dynamical model (Fig. A1). Simulations forced with observed aerosols do not substantially increase activity relative to simulations using prescribed CMIP5 aerosols (Fig. A2). Idealized aerosol reduction experiments, acting as exaggerated analogues of the 2020 drop in aerosols) also show no statistically robust impact on basin-wide TC frequency (Fig. A3). Moreover, analysis of large-scale environmental conditions using the theoretical framework of Hsieh et al. (2020) indicates that 2020 did not exhibit environmental parameters strongly favoring a hyperactive season. Collectively, these results suggest that the SST-forced atmosphere-only models did not fundamentally fail in simulating moderate activity, as the large-scale state itself was not strongly conducive to a hyperactive TC season.

Using the theoretical framework, we further find that subseasonal variability in 2020 may have made it more difficult for the atmosphere-only GFDL models to simulate a hyperactive season given monthly boundary conditions alone (Fig. 4). In contrast, for other active seasons such as 2005 and 2010, the same models produced relatively elevated activity under identical monthly SST forcing, suggesting that the character of intraseasonal atmospheric variability in 2020 reduced the likelihood of a hyperactive outcome within the ensemble. Even after incorporating known sources of subseasonal variability within the framework (Fig. 6), 2020 remains an outlier. This suggests the 2020 season was either a result of an exceptionally rare realization of internal atmospheric variability or from variability mechanisms not fully captured by the framework.

Additional features of the 2020 season are consistent with enhanced subseasonal noise. Activity was disproportionately concentrated in the Gulf of Mexico (Fig. B1), a region where TC variability

35

is less tightly coupled to basin-scale predictors (Ng and Vecchi 2020). The season also included a large number of relatively weak storms compared with other hyperactive years such as 2005 and 2010 (Fig. B2). These characteristics are consistent with a scenario in which weather-scale variability amplified seasonal counts beyond what would be anticipated from large-scale environmental forcing alone.

We also assessed whether 2020 can be interpreted as a low-probability but possible outcome under the observed SST forcing. Within the limited five to ten member dynamical ensembles, the observed record of TC, hurricane, and major hurricane counts is expected to lie outside the ensemble range several times over a 45-year period. Therefore, exceedance of the ensemble range in 2020 across all storm intensity bins is not unprecedented when viewed in a multidecadal context. However, small ensemble sizes limit the ability to robustly estimate tail probabilities. To address this limitation, we generated a 1,000-member ensemble of the year 2020 using the ACE2 DL–based model. This large ensemble indicates that, under the observed 2020 SST forcing, a 23-TC season corresponds to approximately a 0.5% event. Although such an outcome is highly unlikely in any single year, the probability of at least one 0.5% event occurring within a 45-year period is approximately 20%. Thus, the hyperactive 2020 season can be interpreted as a rare but dynamically plausible realization rather than clear evidence of systematic model failure.

5. Conclusions

The evidence presented here supports the conclusion that 2020 was a result of merely an unlikely atmospheric state, not necessarily a model or observational failure. Further, the results from this paper support the conclusion that the hyperactive 2020 season can be understood as an unlikely outcomes that even a well-calibrated system, like the modeling systems used in this study, will exhibit. This interpretation does not exclude the possibility that additional mechanisms not examined here contributed to the observed activity. Future observations, modeling advances, or reanalysis improvements could alter this assessment. New evidence could emerge to support hypotheses one through four, for which we were not able to find compelling support in the results presented here. Moreover, if similar large discrepancies between observations and model ensembles become more frequent, reassessment of model calibration and structural assumptions would be warranted.

36

Our findings have several broader implications. First, the 2020 season occurs at the end of the recent historical record and is anomalously hyperactive, which makes it disproportionately influential in linear trend estimates (Fig. D1). Because trend calculations are especially sensitive to extreme values near the endpoints of a time series, the inclusion of 2020 amplifies the estimated 1980–2024 increase in Atlantic TC, hurricane, and major hurricane activity by approximately 6 storms per century for TCs, 2 storms per century for hurricanes, and 1 storm per century for major hurricanes. Although removing 2020 does not reverse the positive sign of the TC or major hurricane trends at the 95% confidence level, it substantially reduces their magnitudes. In contrast, for hurricanes, excluding 2020 causes the 95% confidence interval to cross zero, implying that the statistical evidence for a positive hurricane trend weakens considerably in its absence. This behavior indicates that 2020 acts as a high-leverage endpoint in the regression and exerts disproportionate influence on the inferred trends. This sensitivity is important because projected changes in TC frequency under anthropogenic forcing remain an open and actively debated research question (e.g., Knutson et al. 2020; Sobel et al. 2021). If the exceptional 2020 season primarily reflects internal variability, such as subseasonal fluctuations or weather-scale noise, then its strong influence on multidecadal trends complicates the interpretation of those trends as indicators of externally forced climate change. These results highlight the need for caution when drawing conclusions about long-term TC trends from relatively short observational records, particularly when extreme endpoint seasons exert substantial leverage on estimated changes. Our findings therefore stress the importance of employing statistical measures more resistant to the influence of outliers, such as the median of pairwise slopes (Lazante 1996).

Second, this study underscores the importance of properly characterizing ensemble uncertainty, especially with a limited ensemble size (e.g., Deser et al. 2012; Deser 2020; Deser and Phillips 2023). Small ensemble sizes may underrepresent tail risk and can create the impression that out-of-range observations reflect model failure. Even with well-calibrated models, we still must plan for extreme departures from the ensemble mean, since they are still possible outcomes. Additionally, we stress that relying on the ensemble mean when making predictions can obscure plausible but extreme outcomes (Table 1).

Third, our results underscore the important role of subseasonal variability and weather-scale noise in shaping seasonal TC outcomes. The analysis suggests that 2020 may have been more

37

strongly influenced by subseasonal variability than other recent active seasons (Fig. 4). Such variability could reduce the ability of atmosphere-only models, when forced with monthly boundary conditions, to reproduce the observed level of hyperactivity. This finding reinforces the need to carefully separate externally forced climate signals from internally generated variability when interpreting extreme seasons. Weather-scale fluctuations can project onto seasonal statistics and thereby contaminate apparent climate signals (Kortum et al. 2024). Consequently, caution is warranted when attributing an extreme season such as 2020 to long-term climate variability without explicitly accounting for the contribution of internal atmospheric noise.

Lastly, this study is one of many to highlight the value of very large ensembles (e.g., Mahesh et al. 2024a,b). We found that the five to ten member size from the dynamical models used in this study is not sufficient to study a full distribution of potential outcomes, especially given that the observed record is expected to exceed the ensemble spread across a multidecadal historical record several times. Traditional modeling efforts leverage a relatively small ensemble size. For example, the seasonal forecasting systems in Murakami et al. (2025) and Zhang et al. (2025a) use ensembles of 10–15 members, and operational global meteorological agencies generally use no more than a few dozen. In contrast, this study employs 1,000 ensemble members of the ACE2 DL model to assess the full distribution of outcomes for 2020 Atlantic TC season. By leveraging the computational efficiency of DL–based weather and climate emulators, future forecast systems should continue to incorporate much larger ensembles and explicitly consider tail risks, rather than depending primarily on ensemble means, to prepare for low-probability but high-impact events.

In summary, the 2020 Atlantic TC season appears consistent with a rare realization within a probabilistic framework rather than evidence of systemic model failure. Even well-calibrated forecasting systems will occasionally encounter extreme outcomes like the 2020 season. Recognizing and quantifying this possibility is essential for interpreting past extremes and preparing for future risk.

Data availability statement. Source code of the HIRAM model is available from `https://www.gfdl.noaa.gov/hiram-quickstart`, and the source code of the ACE2 model is available from `https://github.com/ai2cm/ace`. The forcing data and initial conditions for the ACE2 model are available from `https://huggingface.co/allenai/ACE2-ERA5/tree/main/initial_conditions`. ERA5 data are available from `https://cds.climate.copernicus.eu/datasets/reanalysis-era5-pressure-levels-monthly-means?tab=overview`, and MERRA2 data are available from `https://gmao.gsfc.nasa.gov/gmao-products/merra-2/data-access_merra-2/`. The OISST data are available from `https://www.ncei.noaa.gov/products/optimum-interpolation-sst`, the HadISST data are available from `https://www.metoffice.gov.uk/hadobs/hadisst/`, and the ERSST data are available from `https://www.ncei.noaa.gov/products/extended-reconstructed-sst`. The code to compute the components of the tropical cyclone activity proxy are available from `https://github.com/tlhsieh/tropical_cyclone_seeds` and `https://github.com/wy2136/wython/tree/main/xtci/shared`. The TC seed dataset of Ikehata and Satoh (2021) is available at `https://doi.org/10.5281/zenodo.5136292`, the low pressure system dataset of Vishnu et al. (2020) is available at `https://portal.nersc.gov/cfs/m3310/VishnuEtAl_TrackDataset/`, and the seed dataset of Moon et al. (2025) is available at `https://doi.org/10.5281/zenodo.15227119`.

**Additional figures regarding hypotheses three and four**
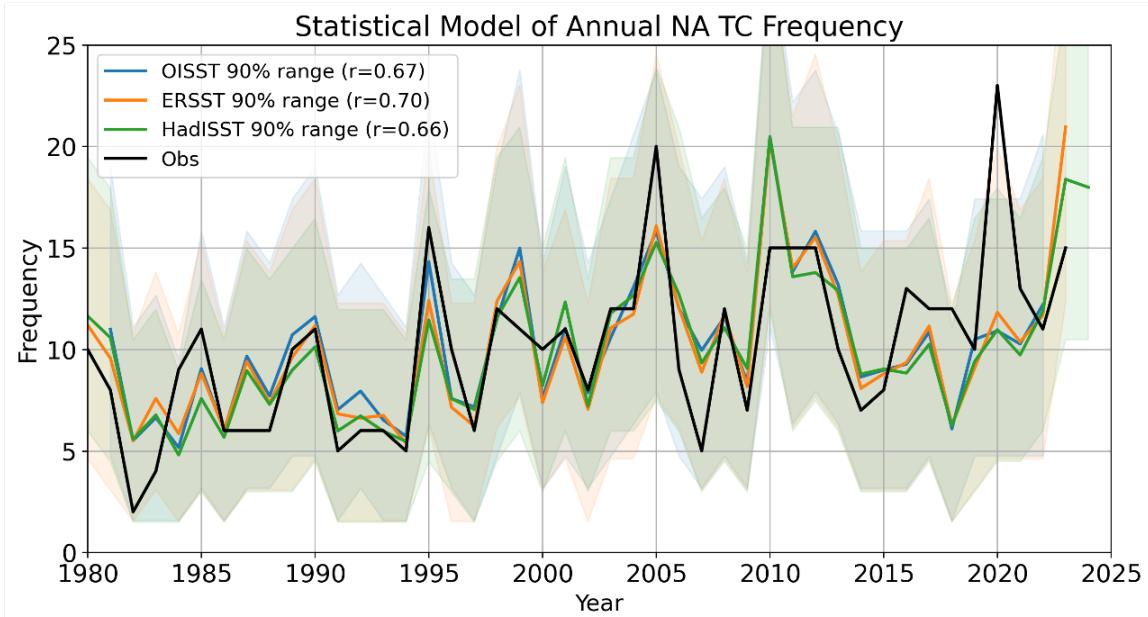
FIG. A1. Application of the statistical model for annual Atlantic TC counts to multiple observed sea surface temperature (SST) products. Expected TC counts from the model are shown in blue (OISST), orange (ERSST), and green (HadISST), with observed storm counts shown in black. Shaded regions indicate the 90% confidence intervals of the statistical fits. The correlation coefficients between modeled and observed TC counts are listed in the legend of the upper-left panel.
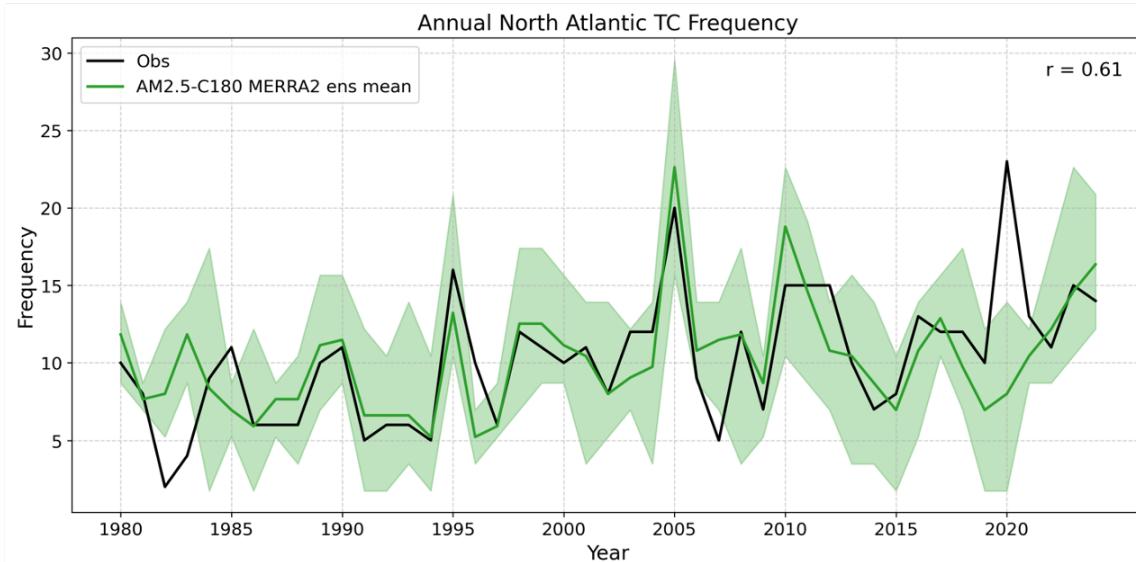
FIG. A2. Annual observed Atlantic TC frequency from 1980 to 2024 (black) and corresponding simulations from the AM2.5-C180 model forced with observed MERRA2 aerosols (green). Here, the green shading represents the minimum and maximum ensemble members for the model (5 total ensemble members), while the solid lines indicate the ensemble mean. The correlation coefficient between ensemble mean modeled and observed TC counts is shown in the upper-right corner of the panel.
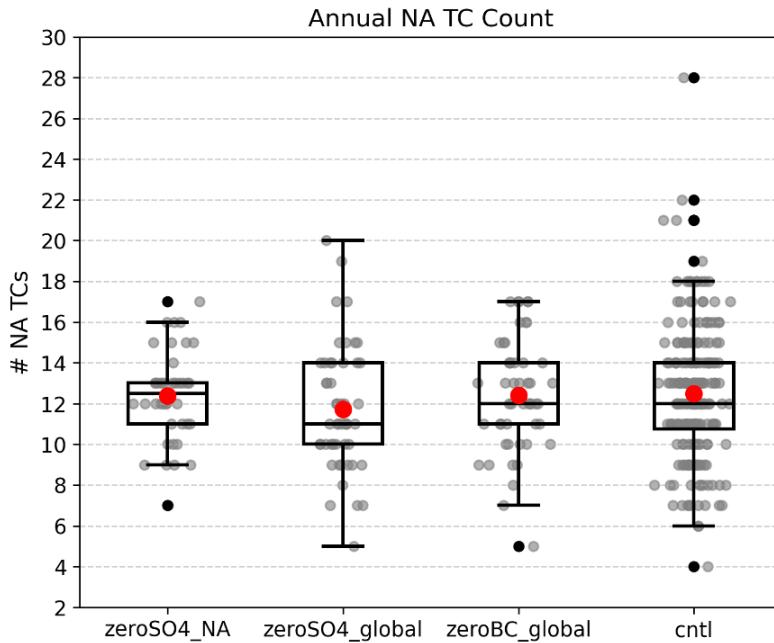
FIG. A3. Jittered data points and box-and-whisker plots showing the distribution of annual North Atlantic TC counts simulated by the zeroSO4NA, zeroSO4global, and cntl HIRAM experiments. For each box-and-whisker plot the horizontal line at the bottom of the box indicates the 25th percentile, the horizontal line in the middle of the box indicates the 50th percentile (median), the horizontal line at the top of the box indicates the 75th percentile, the whisker ends (upper and lower horizontal lines) extend to the furthest data point within 1.5 × the interquartile range, black dots beyond the whiskers indicate outliers beyond 1.5 × the interquartile range, the red dot indicates the mean value of each distribution.

APPENDIX B

Storm track density and intensity

We compare TC track density among the active seasons of 2020, 2010, and 2005 in Fig. B1. The most pronounced difference between the 2020 season, which was highly active yet difficult to predict, and the other two predictably active seasons is the greater storm density observed in the Gulf of Mexico. In 2020, there were at least 15 more TC days per year per $10° × 10°$ grid cell in this region compared to 2005 and 2010. These results indicate that environmental conditions in the Gulf of Mexico played an important role in distinguishing the 2020 season from the others. Accordingly, when computing basin-averaged quantities using ERA5 data in Section 2, we extend the analysis beyond the Main Development Region ($10°$–$25°$ N, $80°$–$20°$ W) to include the broader North Atlantic domain ($10°$–$30°$ N).
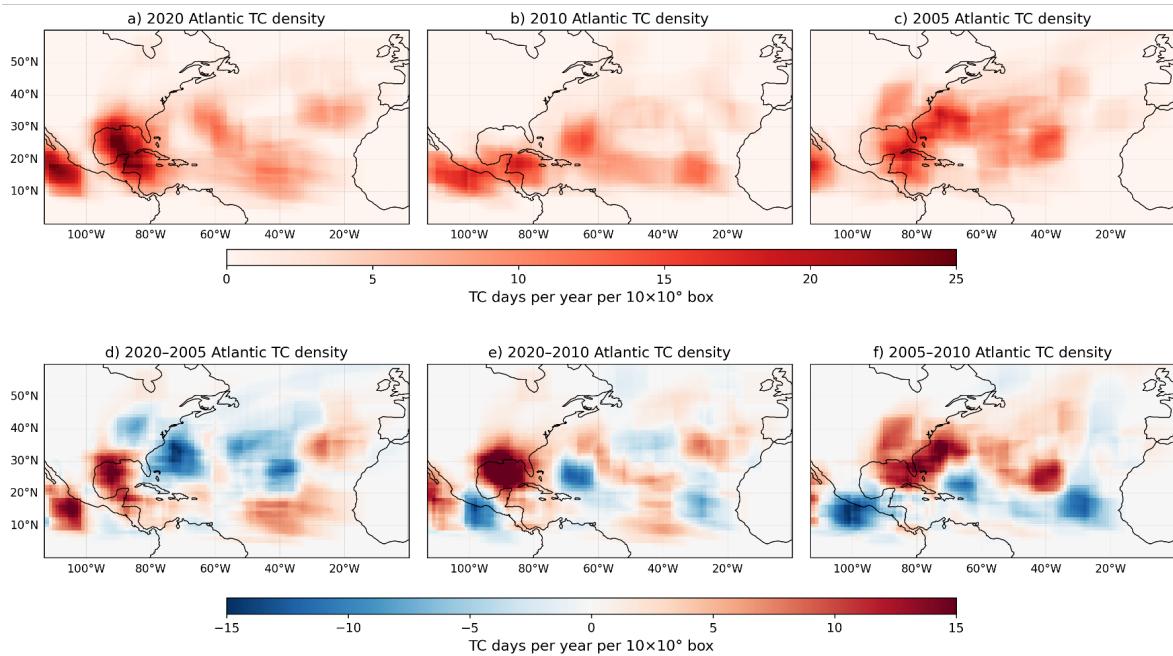
FIG. B1. Density maps of Atlantic tropical cyclone (TC) locations for observed TCs during the (a) 2020, (b) 2010, and (c) 2005 seasons. Density is defined as the number of TC days per year within each model grid cell, aggregated into $10° × 10°$ spatial bins. Panel (d) shows the difference in TC density between the 2020 and 2005 seasons, panel (e) shows the difference between the 2020 and 2010 seasons, and panel (f) shows the difference between the 2005 and 2010 seasons.

We also compare storm intensities among the active seasons of 2020, 2010, and 2005 (Fig. B2). Although no statistically significant difference is found in maximum wind speed between storms in 2020 and those in 2005 or 2010, the 2020 season exhibits a higher density of weaker, pre-hurricane storms with maximum wind speeds between 18 and 30 m s^{-1}. Hurricanes have a minimum wind speed threshold of 33 m s^{-1}. This pattern suggests that the models used in this study, both physics-based and deep learning-based, may struggle to represent or may underdetect these lower-intensity storms, which were more common in 2020 than in the other active seasons. Further investigation of this behavior is warranted.
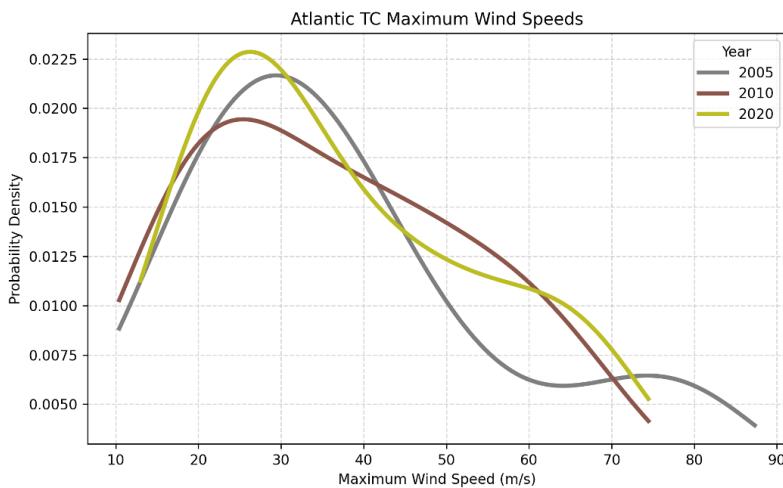
FIG. B2. Kernel Density Estimate (KDE) of the maximum Atlantic TC wind speeds in m s^{-1} for storms in 2020 (olive), 2010 (brown), and 2005 (gray).

APPENDIX C

Large-scale environmental conditions in ACE2

Using the large ensemble sizes generated by the ACE2 model, we examine several large-scale environmental conditions that influence TC development, comparing the active 2005 TC season with the inactive 1982 season. For each season, we analyze the ensemble-mean June–November averages derived from monthly mean model output across a randomly selected 50 ensemble members. Due to data storage limitations, monthly mean data are used instead of higher temporal frequency data. Based on the available ACE2 output variables, we select several fields that represent

key components of the TC development framework described in Section 2. The following ensemble-mean, seasonal-mean variables are compared between the active and inactive TC seasons:

- **Surface precipitation:** an indicator of convection and vertical ascent, which are critical for seed development arising from mesoscale tropical convective systems with intense precipitation rates. Surface precipitation serves as an estimate of $(-\omega)$, which is an input variable to the Seed Propensity Index (SPI), a proxy for TC precursors or seeds (Hsieh et al. 2020).

- **Specific total water (vapor plus condensates) at sigma-pressure level 5 (approximately 600 hPa):** an indicator of mid-tropospheric humidity, which is essential for sustaining TC development. The Ventilation Index (Λ) suggests that a moist mid-troposphere is necessary to prevent the entrainment of dry air into the storm, which can deplete its energy (Tang and Emanuel 2012).

- **Vertical wind shear between sigma-pressure levels 3 (approximately 250 hPa) and 6 (approximately 800 hPa):** a measure of atmospheric vertical wind shear, an explicit component of the Ventilation Index. Strong shear can induce storm asymmetries, tilting the vortex and generating mesoscale eddies that transport low-entropy (colder, drier) air into the storm's midlevels, thereby disrupting the thermodynamic structure required for intensification (Tang and Emanuel 2012). Lower shear is therefore more favorable for storm development.

The results comparing the large-scale, seasonally averaged environmental conditions between the randomly selected 50 ensembles of the active 2005 season and the inactive 1982 season are shown in Figs. C1a–c. Overall, the conditions during 2005 were statistically significantly more favorable for TC development than those during 1982. Surface precipitation was notably higher across the Atlantic main development region (MDR), except in the Gulf of Mexico, indicating stronger vertical ascent and convective activity during 2005. In addition, the mid-troposphere was generally moister across the 2005 ensembles, with higher specific total water content throughout much of the MDR, again with the exception of the Gulf of Mexico. Vertical wind shear was also lower across the 50 ensemble members of the 2005 TC season compared to those of 1982. Together, these results highlight two key points. First, they demonstrate the ability of the DL ACE2 model to distinguish between active and inactive TC seasons for physically grounded reasons. Second,

45

they suggest that, when forced with observed SSTs, the contrasting seasonal mean environmental conditions in 2005 and 1982 can largely explain their differing levels of TC activity.
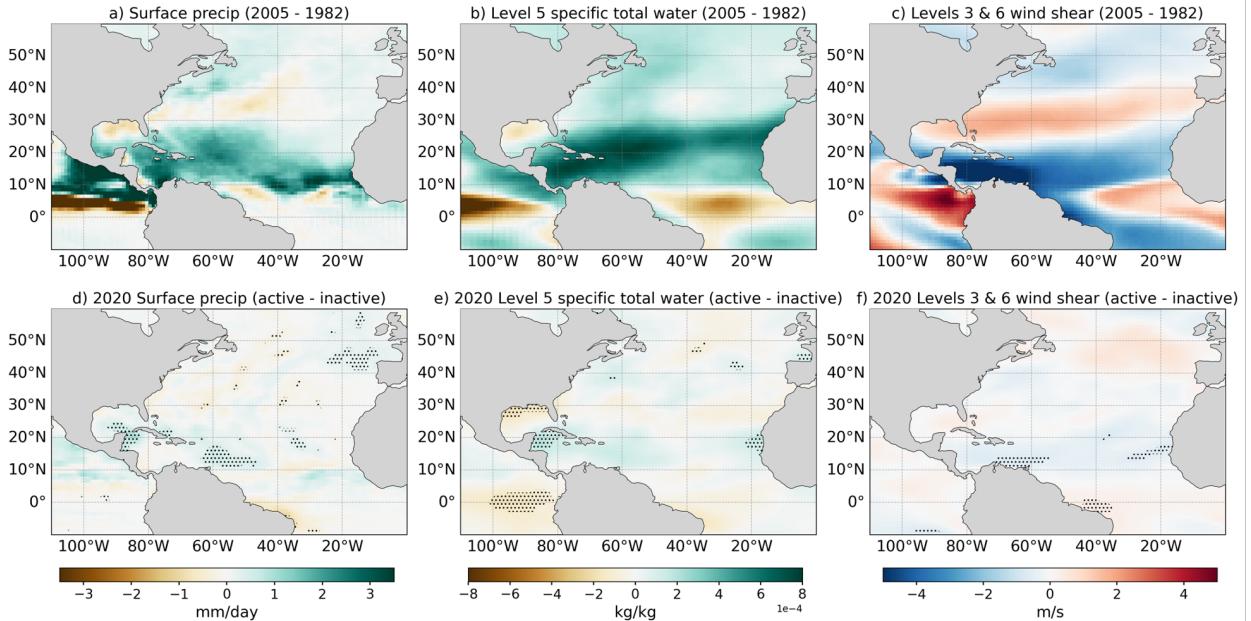
FIG. C1. (a–c) Differences in ensemble mean environmental conditions between the randomly selected 50 ACE2 model–generated ensembles of the active 2005 season and the randomly selected 50 ensembles of the inactive 1982 season. Panel (a) shows surface precipitation, (b) shows specific total water (vapor plus condensates) at sigma-pressure level 5 (approximately 600 hPa), and (c) shows vertical wind shear between sigma-pressure levels 3 (approximately 250 hPa) and 6 (approximately 800 hPa). Panels (d–f) show the same environmental variables but for the 2020 season, comparing the 37 ensembles that produced at least 20 Atlantic tropical cyclones to the 37 ensembles that produced 8 or fewer. These ensembles are subsampled from the full set of 1,000 ACE2 model ensembles for 2020. Stippling in panels (d–f) indicates statistically significant differences based on a two-sample t-test ($p < 0.05$). Stippling is omitted from panels (a–c) because most regions show statistically significant results. All values represent June–November seasonal means derived from monthly mean model output. Together, these panels illustrate the environmental differences between active and inactive Atlantic TC seasons.

The exceptionally large ACE2 model 1,000-member ensemble for the 2020 season provides an opportunity to compare large-scale environmental conditions between ensemble members that produce a hyperactive Atlantic TC season and those that yield a quiet one. Specifically, we compare the ensemble mean of the 37 members that simulated at least 20 Atlantic TCs with the

46

ensemble mean of the 37 members that simulated eight or fewer TCs. This approach enables a comparison of the large-scale environments between plausible realizations of the 2020 season that share identical SST boundary forcing but result in markedly different TC outcomes. The results are presented in Figs. C1d–f. Overall, the conditions are more favorable for TC development in the ensemble members that produced hyperactive seasons compared to those that produced inactive ones. Statistically significant increases in surface precipitation are evident in the main development region (MDR), particularly off the northeast coast of Venezuela and in the Gulf of Mexico. Mid-tropospheric moisture is also statistically significantly higher in the active ensemble members, especially near Cuba and more broadly across the MDR. Finally, vertical wind shear is generally lower in the active ensemble members, with statistically significant reductions off the northeast coast of Venezuela and the west coast of Senegal.

When comparing the differences in large-scale environmental conditions between the active 2005 and inactive 1982 hurricane seasons (Fig. C1a-c) to those between the active and inactive ensemble members of the 2020 season (Fig. C1d-f), the contrasts for 2020 are markedly weaker. The 2005 season typically exhibited a seasonally averaged surface precipitation rate of approximately 2–4 kg m^{-2} s^{-1} higher than that of the 1982 season, whereas the most statistically significant differences between active and inactive 2020 ensemble members are only about 1 kg m^{-2} s^{-1}. Similarly, while the ensemble-mean mid-tropospheric moisture in the active 2020 ensembles is at most 1×10^{-4} kg kg^{-1} higher than in the inactive ensembles, the corresponding difference between 2005 and 1982 reaches roughly $5–8\times10^{-4}$ kg kg^{-1}. The reduction in vertical wind shear is also much smaller: about 1 m s^{-1} between the active and inactive 2020 ensembles compared with approximately 3–6 m s^{-1} between 2005 and 1982.

These results indicate that, given the boundary forcing from the observed 2020 SSTs, the ACE2 model produces only a limited range of seasonally averaged atmospheric conditions among different realizations of the 2020 season. This range is far smaller than that which distinguishes an observed hyperactive season, such as 2005, from a quiet season like 1982. Therefore, some of the variability that separates highly active from inactive outcomes in 2020 likely arises from subseasonal atmospheric fluctuations that cannot be captured by seasonally or monthly averaged quantities. These results further underscore the critical role of SSTs in modulating the subseasonal atmospheric environment conducive to TC development. Notably, all 1,000 ACE2 ensembles of

47

the 2020 season were forced with identical SST boundary conditions, whereas the 2005 and 1982 ensembles each used their respective, distinct SST fields.

Influence of 2020 on historical trends

To evaluate the sensitivity of linear trend estimates to individual seasons, we conducted a leave-one-out analysis over 1980–2024 for observed Atlantic TC, hurricane, and major hurricane counts. For each year, we removed that season from the record and recomputed the ordinary least squares linear trend of seasonal storm counts against time (Fig. D1). This approach quantifies the influence of individual years on the estimated long-term trend. The results indicate that the hyperactive 2020 season exerts substantial leverage on the historical record. Excluding 2020 reduces the estimated linear trend by approximately 6 storms per century for TCs, 2 storms per century for hurricanes, and 1 storm per century for major hurricanes. For TCs and major hurricanes, the sign of the trend remains positive even when 2020 is removed, and the corresponding 95% confidence intervals do not cross zero. However, excluding 2020 from the hurricane record is uniquely impactful: it is the only omission that causes the 95% confidence interval of the estimated trend to dip below zero, implying that the statistical evidence for a positive hurricane trend weakens considerably in its absence. These results underscore the disproportionate influence of the anomalously active 2020 season, occurring at the end of the record, on inferred historical trends in Atlantic storm activity.

48
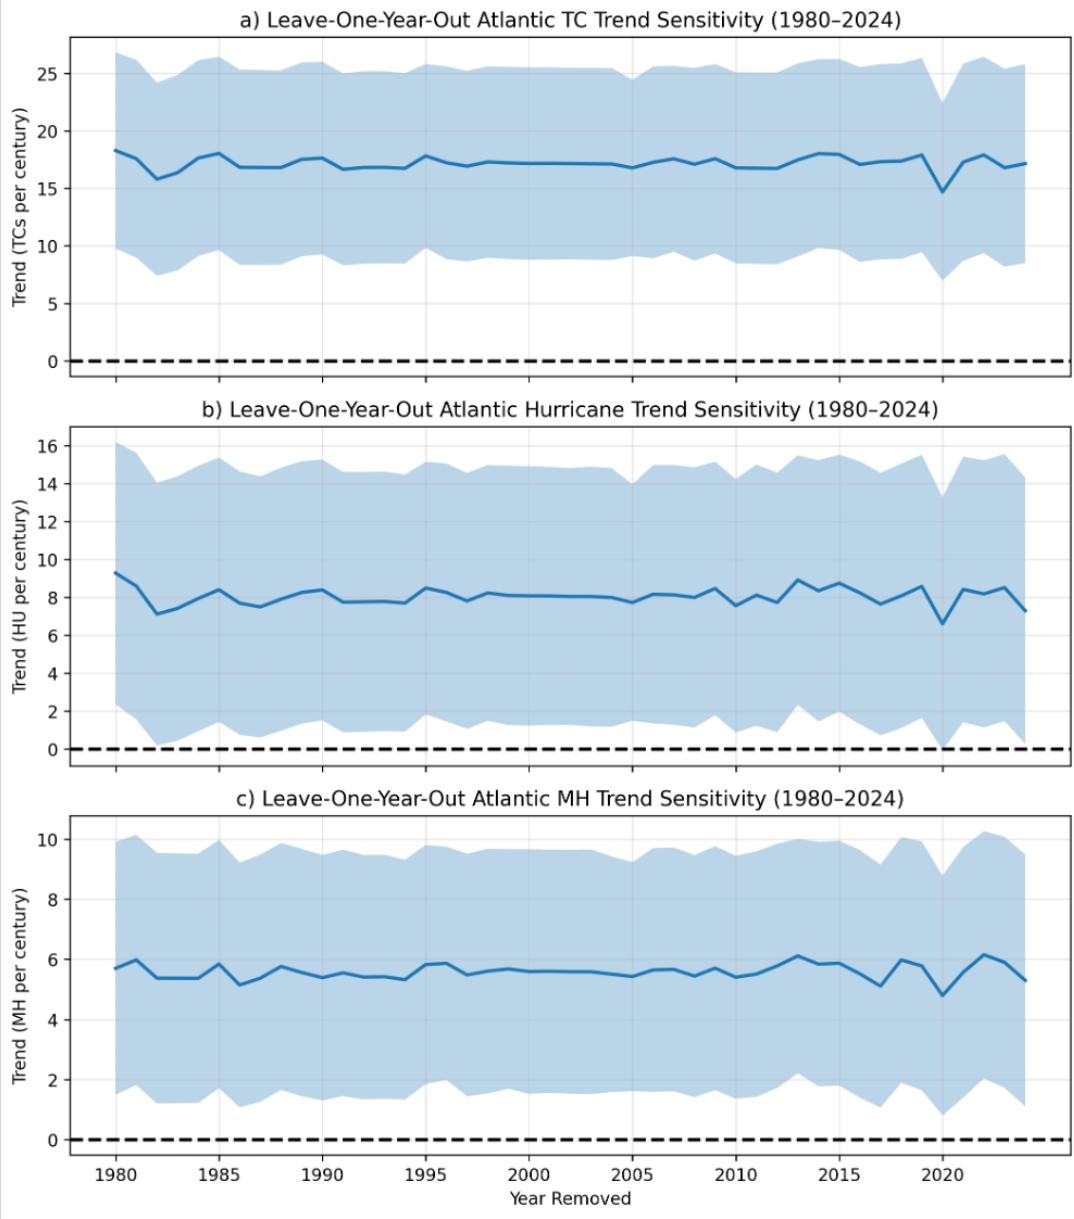
FIG. D1. Leave-one-year-out sensitivity of linear trends in observed Atlantic storm activity over 1980–2024 for (a) TCs, (b) hurricanes, and (c) major hurricanes. For each year, the indicated trend (storms per century) is computed after removing that year from the record and refitting an ordinary least squares linear regression of seasonal counts against time. Shading denotes the analytic 95% confidence interval of the estimated slope based on the Student-t distribution with $n-2$ degrees of freedom. The horizontal dashed line marks zero trend. This analysis quantifies the influence of individual seasons on the computed multidecadal historical trend.

References

Bi, K., L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, 2023: Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, **619 (7970)**, 533–538, https://doi.org/10.1038/s41586-023-06185-3.

Bonev, B., T. Kurth, C. Hundt, J. Pathak, M. Baust, K. Kashinath, and A. Anandkumar, 2023: Spherical Fourier neural operators: Learning stable dynamics on the sphere. *Proceedings of the 40th International Conference on Machine Learning*, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., PMLR, Proceedings of Machine Learning Research, Vol. 202, 2806–2823, URL https://proceedings.mlr.press/v202/bonev23a.html.

Bretherton, C. S., J. R. McCaa, and H. Grenier, 2004: A new parameterization for shallow cumulus convection and its application to marine subtropical cloud-topped boundary layers. part i: Description and 1d results. *Monthly Weather Review*, **132 (4)**, 864–882, https://doi.org/10.1175/1520-0493(2004)132⟨0864:ANPFSC⟩2.0.CO;2.

Chan, D., G. A. Vecchi, W. Yang, and P. Huybers, 2021: Improved simulation of 19th- and 20th-century north atlantic hurricane frequency after correcting historical sea surface temperatures. *Science Advances*, **7 (26)**, eabg6931, https://doi.org/10.1126/sciadv.abg6931.

Chen, J.-H., and S.-J. Lin, 2011: The remarkable predictability of interannual variability of atlantic hurricanes during the past decade. *Geophysical Research Letters*, **38 (11)**, L11 804, https://doi.org/10.1029/2011GL047629.

Chen, L., X. Zhong, F. Zhang, Y. Cheng, Y. Xu, Y. Qi, and H. Li, 2023: Fuxi: A cascade machine learning forecasting system for 15-day global weather forecast. *npj Climate and Atmospheric Science*, **6 (1)**, 190, https://doi.org/10.1038/s41612-023-00512-1.

Chien, M.-T., E. A. Barnes, and E. D. Maloney, 2025: Modulation of tropical cyclogenesis on subseasonal-to-interannual timescales in the deep-learning climate emulator ace2. *Machine Learning: Earth*, **1 (1)**, 015 008, https://doi.org/10.1088/3049-4753/adfd61.

Delworth, T. L., and Coauthors, 2012: Simulated climate and climate change in the GFDL CM2.5 high-resolution coupled climate model. *Journal of Climate*, **25 (8)**, 2755–2781, https://doi.org/10.1175/JCLI-D-11-00316.1.

Deser, C., 2020: Certain uncertainty: The role of internal climate variability in projections of regional climate change and risk management. *Earth's Future*, **8 (12)**, e2020EF001854, https://doi.org/10.1029/2020EF001854.

Deser, C., A. Phillips, V. Bourdette, and H. Teng, 2012: Uncertainty in climate change projections: the role of internal variability. *Climate Dynamics*, **38 (3-4)**, 527–546, https://doi.org/10.1007/s00382-010-0977-x.

Deser, C., and A. S. Phillips, 2023: A range of outcomes: the combined effects of internal variability and anthropogenic forcing on regional climate trends over europe. *Nonlinear Processes in Geophysics*, **30 (1)**, 63–84, https://doi.org/10.5194/npg-30-63-2023.

Diamond, M. S., 2023: Detection of large-scale cloud microphysical changes within a major shipping corridor after implementation of the international maritime organization 2020 fuel sulfur regulations. *Atmospheric Chemistry and Physics*, **23 (14)**, 8259–8269, https://doi.org/10.5194/acp-23-8259-2023.

Eusebi, R., W. Yang, G. A. Vecchi, and S. Fueglistaler, 2025: Statistical modeling of north atlantic hurricane frequency and the impact and role of patterned warming. *Journal of Climate*, **38 (19)**, 5391–5410, https://doi.org/10.1175/JCLI-D-24-0647.1.

Gelaro, R., and Coauthors, 2017: The modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *Journal of Climate*, **30 (14)**, 5419–5454, https://doi.org/10.1175/JCLI-D-16-0758.1.

Gete, P., A. Tsouderou, and S. M. Wachter, 2025: Climate risk in mortgage markets: Evidence from hurricanes harvey and irma. *Real Estate Economics*, **52 (3)**, 660–686, https://doi.org/10.1111/1540-6229.12477.

Gray, W. M., 1984: Atlantic seasonal hurricane frequency. part ii: Forecasting its variability. *Monthly Weather Review*, **112 (9)**, 1669–1683.

Gupta, P., A. Jangid, and R. Kumar, 2023: Covid-19–associated 2020 lockdown: Atmospheric black carbon reduction and human health impacts. *Environmental Geochemistry and Health*, **45 (6)**, 3507–3520, https://doi.org/10.1007/s10653-022-01430-6.

Harris, L. M., S.-J. Lin, and C. Tu, 2016: High-resolution climate simulations using GFDL HiRAM with a stretched global grid. *Journal of Climate*, **29 (11)**, 4293–4314, https://doi.org/10.1175/JCLI-D-15-0389.1.

Hersbach, H., and Coauthors, 2020: The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, **146 (730)**, 1999–2049, https://doi.org/10.1002/qj.3803.

Hsieh, T.-L., G. A. Vecchi, W. Yang, I. M. Held, and S. T. Garner, 2020: Large-scale control on the frequency of tropical cyclones and seeds: a consistent relationship across a hierarchy of global atmospheric models. *Climate Dynamics*, **55 (11)**, 3177–3196, https://doi.org/10.1007/s00382-020-05446-5.

Hsieh, T.-L., W. Yang, G. A. Vecchi, and M. Zhao, 2022: Model spread in the tropical cyclone frequency and seed propensity index across global warming and enso-like perturbations. *Geophysical Research Letters*, **49 (7)**, e2021GL097 157, https://doi.org/10.1029/2021GL097157.

Hsieh, T.-L., B. Zhang, W. Yang, G. A. Vecchi, M. Zhao, B. J. Soden, and C. Wang, 2023: The influence of large-scale radiation anomalies on tropical cyclone frequency. *Journal of Climate*, **36 (16)**, 5431–5441, https://doi.org/10.1175/JCLI-D-22-0449.1.

Huang, B., C. Liu, V. Banzon, E. Freeman, G. Graham, B. Hankins, T. Smith, and H.-M. Zhang, 2021: Improvements of the daily optimum interpolation sea surface temperature (doisst) version 2.1. *Journal of Climate*, **34 (8)**, 2923–2939, https://doi.org/10.1175/JCLI-D-20-0166.1.

Huang, B., and Coauthors, 2017: Extended reconstructed sea surface temperature, version 5 (ersstv5): Upgrades, validations, and intercomparisons. *Journal of Climate*, **30 (20)**, 8179–8205, https://doi.org/10.1175/JCLI-D-16-0836.1.

Ikehata, K., and M. Satoh, 2021: Climatology of tropical cyclone seed frequency and survival rate in tropical cyclones. *Geophysical Research Letters*, **48 (18)**, e2021GL093 626, https://doi.org/10.1029/2021GL093626.

Jordan, G., and M. Henry, 2024: IMO2020 regulations accelerate global warming by up to 3 years in UKESM1. *Earth's Future*, **12 (8)**, e2024EF005 011, https://doi.org/10.1029/2024EF005011.

Klotzbach, P. J., S. G. Bowen, R. Pielke, and M. Bell, 2018: Continental u.s. hurricane landfall frequency and associated damage: Observations and future risks. *Bulletin of the American Meteorological Society*, **99 (7)**, 1359–1376, https://doi.org/10.1175/BAMS-D-17-0184.1.

Klotzbach, P. J., and W. M. Gray, 2003: Forecasting september atlantic basin tropical cyclone activity. *Weather and Forecasting*, **18 (6)**, 1109–1128, https://doi.org/10.1175/1520-0434(2003) 018⟨1109:FSABTC⟩2.0.CO;2.

Klotzbach, P. J., and Coauthors, 2022: A hyperactive end to the atlantic hurricane season october–november 2020. *Bulletin of the American Meteorological Society*, **103 (1)**, S1–S28, https://doi.org/10.1175/BAMS-D-20-0312.1.

Knutson, T., and Coauthors, 2020: Tropical cyclones and climate change assessment: Part II: Projected response to anthropogenic warming. **101 (3)**, E303–E322, https://doi.org/10.1175/ BAMS-D-18-0194.1.

Kochkov, D., and Coauthors, 2024: Neural general circulation models for weather and climate. *Nature*, **632 (8027)**, 1060–1066, https://doi.org/10.1038/s41586-024-07744-y.

Kortum, G., G. A. Vecchi, T.-L. Hsieh, and W. Yang, 2024: Influence of weather and climate on multidecadal trends in atlantic hurricane genesis and tracks. *Journal of Climate*, **37 (5)**, 1501–1522, https://doi.org/10.1175/JCLI-D-23-0088.1.

Lam, R., and Coauthors, 2023: Learning skillful medium-range global weather forecasting. *Science*, **382 (6677)**, 1416–1421, https://doi.org/10.1126/science.adi2336.

Landsea, C. W., G. A. Vecchi, L. Bengtsson, and T. R. Knutson, 2010: Impact of duration thresholds on atlantic tropical cyclone counts. **23 (10)**, 2508–2519, https://doi.org/10.1175/2009JCLI3034. 1.

Lang, S., and Coauthors, 2024: AIFS: ECMWF's data-driven forecasting system. URL https://arxiv.org/abs/2406.01465, arXiv preprint, https://doi.org/10.48550/arXiv.2406.01465, 2406. 01465.

Lazante, J. R., 1996: Resistant, robust and non-parametric techniques for the analysis of climate data: Theory and examples, including applications to historical radiosonde sta-

tion data. **16 (11)**, 1197–1226, https://doi.org/10.1002/(SICI)1097-0088(199611)16:11⟨1197::AID-JOC89⟩3.0.CO;2-L.

Levin, E. L., G. A. Vecchi, and W. Yang, 2025: Influence of sea surface temperature patterns and mean warming on past and future atlantic hurricane activity. *Journal of Climate*.

Mahesh, A., and Coauthors, 2024a: Huge ensembles part i: Design of ensemble weather forecasts using spherical fourier neural operators. *arXiv preprint arXiv:2408.03100*, https://doi.org/10.48550/arXiv.2408.03100.

Mahesh, A., and Coauthors, 2024b: Huge ensembles part ii: Properties of a huge ensemble of hindcasts generated with spherical fourier neural operators. *arXiv preprint arXiv:2408.01581*, https://doi.org/10.48550/arXiv.2408.01581.

Menemenlis, S., G. A. Vecchi, W. Yang, S. Fueglistaler, and S. P. Raghuraman, 2025: Consequential differences in satellite-era sea surface temperature trends across datasets. *Nature Climate Change*, **15 (8)**, 897–903, https://doi.org/10.1038/s41558-025-02362-6.

Molina, R., and I. Rudik, 2024: w32548. The social value of hurricane forecasts. URL https://www.nber.org/papers/w32548, https://doi.org/10.3386/w32548.

Moon, J., D. Kim, A. A. Wing, S. J. Camargo, G. N. Emlaw, J. C. Starr, and D.-H. Cha, 2025: Tropical cyclone seed disturbances in ERA5. *Journal of Climate*, **38 (18)**, 4625–4639, https://doi.org/10.1175/JCLI-D-24-0291.1.

Murakami, H., 2022: Substantial global influence of anthropogenic aerosols on tropical cyclones over the past 40 years. *Science Advances*, **8 (19)**, eabn9493, https://doi.org/10.1126/sciadv.abn9493.

Murakami, H., 2024: Effect of regional anthropogenic aerosols on tropical cyclone frequency of occurrence. *Geophysical Research Letters*, **51 (21)**, e2024GL110 443, https://doi.org/10.1029/2024GL110443.

Murakami, H., T. L. Delworth, N. C. Johnson, F. Lu, C. E. McHugh, and L. Jia, 2025: Seasonal forecasts of tropical cyclones using GFDL SPEAR and HiFLOR-s. *Journal of Climate*, **38 (9)**, 3601–3620, https://doi.org/10.1175/JCLI-D-24-0356.1.

Ng, C. H. J., and G. A. Vecchi, 2020: Large-scale environmental controls on the seasonal statistics of rapidly intensifying north atlantic tropical cyclones. *Climate Dynamics*, **54 (9-10)**, 3907–3925, https://doi.org/10.1007/s00382-020-05207-4.

Pielke, R. A., J. Gratz, C. W. Landsea, D. Collins, M. A. Saunders, and R. Musulin, 2008: Normalized hurricane damage in the united states: 1900–2005. **9 (1)**, 29–42, https://doi.org/10.1061/(ASCE)1527-6988(2008)9:1(29).

Roberts, P. S., and K. Wernstedt, 2016: Using climate forecasts across a state's emergency management network. *Natural Hazards Review*, **17 (3)**, 05016 002, https://doi.org/10.1061/(ASCE)NH.1527-6996.0000222.

Schneider, D. P., C. Deser, J. Fasullo, and K. E. Trenberth, 2013: Climate data guide spurs discovery and understanding. *Eos, Transactions American Geophysical Union*, **94 (13)**, 121–122, https://doi.org/10.1002/2013EO130001.

Sobel, A. H., A. A. Wing, S. J. Camargo, C. M. Patricola, G. A. Vecchi, C. Lee, and M. K. Tippett, 2021: Tropical cyclone frequency. *Earth's Future*, **9 (12)**, e2021EF002 275, https://doi.org/10.1029/2021EF002275.

Tang, B., and K. Emanuel, 2010: Midlevel ventilation's constraint on tropical cyclone intensity. *Journal of the Atmospheric Sciences*, **67 (6)**, 1817-1830, https://doi.org/10.1175/2010JAS3318.1.

Tang, B., and K. Emanuel, 2012: A ventilation index for tropical cyclones. *Bulletin of the American Meteorological Society*, **93 (12)**, 1901-1915, https://doi.org/10.1175/BAMSD1100165.1.

Ullrich, P. A., and C. M. Zarzycki, 2017: TempestExtremes: a framework for scale-insensitive pointwise feature tracking on unstructured grids. *Geoscientific Model Development*, **10 (3)**, 1069–1090, https://doi.org/10.5194/gmd-10-1069-2017.

Vecchi, G. A., and B. J. Soden, 2007: Effect of remote sea surface temperature change on tropical cyclone potential intensity. *Nature*, **450 (7172)**, 1066–1070, https://doi.org/10.1038/nature06423.

Vecchi, G. A., and G. Villarini, 2014: Next season's hurricanes. *Science*, **343 (6171)**, 618–619, https://doi.org/10.1126/science.1247759.

Vecchi, G. A., M. Zhao, H. Wang, G. Villarini, A. Rosati, A. Kumar, I. M. Held, and R. Gudgel, 2011: Statistical–dynamical predictions of seasonal north atlantic hurricane activity. *Monthly Weather Review*, **139 (4)**, 1071–1087, https://doi.org/10.1175/2010MWR3499.1.

Vecchi, G. A., and Coauthors, 2014: On the seasonal forecasting of regional tropical cyclone activity. https://doi.org/10.1175/JCLI-D-14-00158.1.

Vecchi, G. A., and Coauthors, 2019: Tropical cyclone sensitivities to co2 doubling: roles of atmospheric resolution, synoptic variability and background climate changes. *Climate Dynamics*, **53 (9)**, 5999–6033, https://doi.org/10.1007/s00382-019-04913-y.

Villarini, G., G. A. Vecchi, T. R. Knutson, and J. A. Smith, 2011a: Is the recorded increase in short-duration north atlantic tropical storms spurious? *Journal of Geophysical Research: Atmospheres*, **116**, D10 112, https://doi.org/10.1029/2010JD015493.

Villarini, G., G. A. Vecchi, T. R. Knutson, M. Zhao, and J. A. Smith, 2011b: North atlantic tropical storm frequency response to anthropogenic forcing: Projections and sources of uncertainty. *Journal of Climate*, **24 (13)**, 3224–3248, https://doi.org/10.1175/2011JCLI3853.1.

Vishnu, S., W. R. Boos, P. A. Ullrich, and T. A. O'Brien, 2020: Assessing historical variability of south asian monsoon lows and depressions with an optimized tracking algorithm. *Journal of Geophysical Research: Atmospheres*, **125 (15)**, e2020JD032 977, https://doi.org/10.1029/2020JD032977.

Vitart, F., and T. N. Stockdale, 2001: Seasonal forecasting of tropical storms using coupled gcm integrations. *Monthly Weather Review*, **129 (10)**, 2521–2537, https://doi.org/10.1175/1520-0493(2001)129⟨2521:SFOTSU⟩2.0.CO;2.

Watt-Meyer, O., and Coauthors, 2025: ACE2: Accurately learning subseasonal to decadal atmospheric variability and forced responses. *npj Climate and Atmospheric Science*, **8 (1)**, 205, https://doi.org/10.1038/s41612-025-01090-0.

Weinkle, J., C. Landsea, D. Collins, R. Musulin, R. P. Crompton, P. J. Klotzbach, and R. Pielke, 2018: Normalized hurricane damage in the continental united states 1900–2017. *Nature Sustainability*, **1 (12)**, 808–813, https://doi.org/10.1038/s41893-018-0165-2.

Yang, W., T.-L. Hsieh, and G. A. Vecchi, 2021: Hurricane annual cycle controlled by both seeds and genesis probability. *Proceedings of the National Academy of Sciences*, **118 (41)**, e2108397 118, https://doi.org/10.1073/pnas.2108397118.

Yang, W., and Coauthors, 2024: Last millennium hurricane activity linked to endogenous climate variability. *Nature Communications*, **15 (1)**, 816, https://doi.org/10.1038/s41467-024-45112-6.

Young, R., and S. Hsiang, 2024: Mortality caused by tropical cyclones in the united states. *Nature*, **635 (8037)**, 121–128, https://doi.org/10.1038/s41586-024-07945-5.

Zhang, G., M. Rao, J. Yuval, and M. Zhao, 2025a: Advancing seasonal prediction of tropical cyclone activity with a hybrid ai–physics climate model. *Environmental Research Letters*, **20 (9)**, 094 031, https://doi.org/10.1088/1748-9326/adf864.

Zhang, J., Y.-S. Chen, E. Gryspeerdt, T. Yamaguchi, and G. Feingold, 2025b: Radiative forcing from the 2020 shipping fuel regulation is large but hard to detect. *Communications Earth & Environment*, **6 (1)**, 1–11, https://doi.org/10.1038/s43247-024-01911-9.

Zhao, M., and I. M. Held, 2010: An analysis of the effect of global warming on the intensity of atlantic hurricanes using a gcm with statistical refinement. *Journal of Climate*, **23 (23)**, 6382–6395, https://doi.org/10.1175/2010JCLI3837.1.

Zhao, M., I. M. Held, S.-J. Lin, and G. A. Vecchi, 2009: Simulations of global hurricane climatology, interannual variability, and response to global warming using a 50-km resolution gcm. *Journal of Climate*, **22 (24)**, 6653–6678, https://doi.org/10.1175/2009JCLI3049.1.

Zhao, M., I. M. Held, and G. A. Vecchi, 2010: Retrospective forecasts of the hurricane season using a global atmospheric model assuming persistence of sst anomalies. *Monthly Weather Review*, **138 (10)**, 3858–3868, https://doi.org/10.1175/2010MWR3366.1.