

Dear Editor,

We sincerely thank you and the two reviewers for the careful and constructive assessment of our revised manuscript. We greatly appreciate the reviewers' recognition of the substantial improvements made in the previous revision and their remaining suggestions concerning terminology, statistical robustness, and province-level spatial attribution.

We have considered all comments carefully and revised the manuscript accordingly. In particular, we harmonized the terminology used for media-derived impact reporting throughout the manuscript; added direct uncertainty diagnostics for the drought-reporting sensitivity coefficient (β); introduced a complementary zero-inclusive occurrence model that retains months with no impact reporting; and expanded the limitations concerning province-level attribution where article-level provenance can no longer be reconstructed. Below, each reviewer comment is reproduced, followed by our response. Where useful, we also reproduce the corresponding text from the clean revised manuscript in blue.

We are grateful for the opportunity to further improve the manuscript and hope that the revised version satisfactorily addresses the remaining concerns.

Sincerely,

Ahmed Kenawy

On behalf of all coauthors

Reviewer 1:

Point-by-point response

– I appreciate the distinction between reporting intensity as a measure of social salience, vs unique drought impacts, as elaborated in lines 114-120 and line 136. Please apply this consistently throughout the paper. I particularly appreciate the clarity of the description on line 136, “sector-specific monthly reporting frequencies.” The authors consistently use this language in the new sections but still need to implement it throughout the paper. I suggest searching for each use of “impact” or “sensitivity” and ensuring they are all consistent with this distinction. For example: Lines 27 and 544, I believe “media-reported impact frequency” should be “frequency of media reports on impacts” or something similar, as you specifically state (lines 119-120) that you are NOT quantifying unique drought impacts. Same comment for line 30. Starting on line 198, in the header of section 2.5 you refer to a drought-impact sensitivity coefficient, but in line 199, you refer to a drought-reporting sensitivity coefficient.

We thank the reviewer for emphasizing this important distinction. We performed a manuscript-wide terminology check, with particular attention to every use of “impact” and “sensitivity.” We now use “frequency of media reports on impacts,” “impact reporting,” “reporting response,” and “drought-reporting sensitivity coefficient” when referring to the media-derived response variable or to statistics calculated from it. References to “impacts” without the reporting qualifier are retained only where the text refers to the physical drought consequences themselves or to the impact categories described in the literature. To prevent ambiguity while avoiding excessive repetition, we also added an explicit terminology convention at the beginning of the Results section.

L248-251: “Throughout this section, the terms hydrological and agricultural impacts are used as shorthand for the corresponding media-derived impact-reporting series. Accordingly, reported sensitivities, spatial patterns, and dominant timescales refer to variation in the frequency and intensity of reporting on these impacts, rather than direct changes in physical impact magnitude or the number of unique impact events.”

Other examples include:

In addition, Sect. 2.5 (L179) is now consistently titled “Drought-reporting sensitivity coefficient.”

L25–30: “Sensitivity analyses reveal pronounced nonlinear responses, with the frequency of media reports on impacts increasing disproportionately during severe drought conditions with a peak at 12–36 months. [...] Our results indicate that the frequency of media reports on impacts is associated with accumulated moisture deficits and temperature-dependent processes across timescales [...]”

L180-183: “To quantify how media reporting varies with drought magnitude, we calculated a drought-reporting sensitivity coefficient (β) from a log-log relationship between drought magnitude and monthly media-reporting counts.”

- Minor comment: Lines 30-33: I believe “across timescales” should go either before “accumulated” or after “processes”, as it should apply to both precipitation and temperature.

Agreed. We have revised the sentence so that “across timescales” clearly applies to both accumulated moisture deficits and temperature-dependent processes.

L28–30: “Our results indicate that the frequency of media reports on impacts is associated with accumulated moisture deficits and temperature-dependent processes across timescales and provide a framework for improved impact-based drought monitoring in a warming climate.”

- Minor comment: Line 111: Need a period at the end.

Corrected (L110). The missing period has been added in the revised manuscript.

- Minor comment: Line 593: missing the “P” in “Persistent”.

Corrected. The word now reads “Persistent” in the revised manuscript.

Reviewer 2:

Point-by-point response

– I would like to thank the authors for the careful and thorough revision of the manuscript as well as for their detailed responses to my original comments. The revised version is considerably clearer, particularly with respect to the construction and interpretation of the media-derived impacts. I also appreciate the clarification of the drought-reporting sensitivity coefficient. The new supplementary figures showing the number of observations supporting the severity-specific regressions are also useful and directly respond to some of my previous questions. Overall, I find that the manuscript has improved substantially and that most of my previous conceptual concerns have been satisfactorily addressed. I nevertheless have two remaining questions concerning the robustness of the sensitivity analysis and the reliability of the province-level spatial attribution.

We sincerely thank the reviewer for this positive assessment of the revised manuscript and for identifying the remaining issues so clearly. We have addressed both points through additional uncertainty and zero-inclusive analyses, as well as a more explicit treatment of the limitations of province-level attribution.

1. Uncertainty and robustness of the drought-reporting sensitivity coefficient

– I appreciate the addition of Supplementary Figs. S1-S3, which demonstrate that the severe and extreme drought regressions are not restricted to isolated timescales with exceptionally small sample sizes. [...] However, I am still uncertain about the statistical precision of the estimated sensitivity coefficients themselves. Reporting the number of observations is useful, but it does not directly indicate how well constrained the regression slope β is. [...] I would therefore encourage the authors to provide some direct indication of uncertainty in β . This could take the form of confidence intervals, standard errors, bootstrap intervals, or variability across provinces around the national mean curves.

We agree that sample size alone does not establish the precision of β , particularly when regressions are fitted within restricted drought-severity ranges. We therefore added direct uncertainty diagnostics for every province–index–timescale regression. Specifically, we calculated the standard error and 95% confidence interval of β and summarized uncertainty across provinces using (i) the interprovincial distribution of β , (ii) the percentage of regressions whose 95% confidence intervals exclude zero, and (iii) the median 95% confidence-interval width. These diagnostics are presented in the new Supplementary Fig. S4 and are discussed explicitly in the Results and Discussion. The revised interpretation now emphasizes that the broad timescale-dependent response is more robust at the aggregated

national and sectoral scales than at the level of individual province–timescale coefficients, especially under severe and extreme drought.

L212–216: “To evaluate the statistical precision of the severity-specific sensitivity estimates, we complemented the sample-size assessment with direct uncertainty diagnostics for β . For each province–index–timescale regression, we calculated the standard error and 95% confidence interval of β . Across provinces, uncertainty was summarized using the interprovincial distribution of β , the percentage of regressions whose 95% confidence intervals excluded zero, and the median 95% confidence-interval width. These diagnostics distinguish the amount of sample support from the precision with which the regression slope is estimated.”

L310–314: “Direct uncertainty diagnostics show substantial interprovincial variability in β , a generally modest proportion of province-level regressions whose 95% confidence intervals exclude zero, and broader confidence intervals under severe and extreme drought conditions (Supplementary Fig. S4). These results indicate that the broad timescale-dependent response is more robust at the aggregated national scale than for individual province–timescale coefficients, which should therefore be interpreted cautiously.”

– A related point concerns the treatment of months with zero impact reporting. [...] the analysis does not capture the transition from no reporting to some reporting. [...] I do not necessarily think the main analysis needs to be replaced, but some supplementary assessment including zero-report months would strengthen the interpretation. A count model or a simple two-stage approach distinguishing the probability of any reporting from reporting intensity conditional on reporting could be informative. At minimum, I suggest discussing more explicitly how conditioning on positive reporting may affect comparisons among drought-severity classes, sectors, and provinces.

We agree and have added the requested complementary zero-inclusive assessment. We retained the original log–log β analysis because it quantifies the elasticity of reporting intensity conditional on positive reporting, but we now complement it with a zero-inclusive first-stage logistic occurrence model for the probability of any reporting, while retaining the original log–log model as the second component describing reporting intensity conditional on positive reporting. For each province, drought index, and accumulation timescale,

reporting occurrence is coded as a binary response (reporting present/absent) and related to positive drought magnitude using logistic regression. We report the coefficient γ , its 95% confidence interval, and the corresponding odds ratio. The new Supplementary Fig. S5 summarizes the probability of any reporting, the occurrence-model odds ratios, and the spatial robustness of the occurrence response. We also expanded the Discussion to state explicitly that conditioning β on positive reporting may omit part of the drought response operating through the transition from zero to positive reporting.

L194–201: “To complement this conditional log–log analysis and explicitly account for zero-report months, we also modelled the occurrence of drought-impact reporting using all dry months ($\text{SPI} < 0$ or $\text{SPEI} < 0$), including months with zero reports (i.e., no identified impact article). For each province, drought index, and accumulation timescale, reporting occurrence was represented as a binary response ($\text{It} > 0 = 1$; $\text{It} = 0 = 0$) and related to positive drought magnitude ($\text{Dt} = -\text{SPIt}$ or $-\text{SPEIt}$) using logistic regression. The coefficient γ , its 95% confidence interval, and the corresponding odds ratio $\exp(\gamma)$ were used to quantify the change in the probability of any reporting with increasing drought magnitude. This occurrence model complements, rather than replaces, the log–log elasticity coefficient β , which quantifies reporting intensity conditional on positive reporting.”

L506–511: “The zero-inclusive analysis further shows that drought reporting comprises two related but distinct components: the initiation of reporting and the subsequent intensity of reporting once it occurs. Greater drought magnitude is generally associated with higher odds of any reporting, with the most spatially coherent occurrence responses concentrated at short-to-intermediate timescales, particularly for SPEI. Consequently, conditioning the log–log analysis on positive counts may omit part of the reporting response where drought primarily affects whether reporting is initiated, and comparisons among sectors, provinces, and timescales should consider both occurrence and conditional intensity.”

2. Robustness of province-level spatial attribution

– I appreciate the transparent clarification that impact labels and province labels were extracted independently and that the workflow does not establish sentence-level links between individual impacts and individual provinces. [...] Still, because province-level analysis remains an important component of the manuscript, I would still like some indication of how frequently this ambiguity occurs. [...] I therefore suggest quantifying, if possible, the proportion of articles that contain: more than one province; more than one broad impact category; and particularly

both multiple provinces and multiple impact categories. A useful sensitivity check could also be to repeat one representative province-level analysis using only single-province articles.

We agree that these diagnostics would be informative if the necessary article-level provenance were still available. However, the final analytical database retained for this study contains monthly province–sector aggregates rather than the original article-level links between each extracted impact category and each geographic reference. Consequently, the requested proportions of articles containing multiple provinces, multiple impact categories, or both cannot be reconstructed reliably from the retained data, and a single-province-only sensitivity analysis cannot be performed without returning to and reprocessing the original source corpus. We therefore cannot reliably infer or approximate these quantities from the aggregated monthly data. Instead, we have made this limitation explicit in the Discussion and have further constrained the interpretation of province-level coefficients and maps. The revised manuscript now states that fine-scale spatial results represent the geographic and sectoral framing of drought in media reporting rather than precisely geocoded inventories of physical impacts, and that individual-province interpretations require greater caution than aggregated national or broad regional patterns.

L513–521: “A further source of uncertainty in the province-level spatial analysis concerns the independent extraction of impact categories and geographic references. In articles mentioning several provinces and several impact types, the available workflow cannot determine whether every reported impact applies to every referenced province. Because the final analytical dataset retains monthly province–sector counts rather than article-level provenance, the frequency of such combinations cannot be quantified retrospectively, nor can the analysis be restricted to single-province articles. This limitation is particularly relevant to the interpretation of fine-scale spatial contrasts. Accordingly, province-level coefficients and maps should be viewed as patterns of the geographic and sectoral framing of drought in media reporting, rather than as precisely geocoded inventories of physical impacts. Aggregated national and broad regional patterns are less dependent on exact within-article attribution, whereas interpretation of individual provinces warrants greater caution.”

– Minor comment: In Sect. 2.3, the manuscript still refers to the “monthly relative frequency” and to the raw data as “proportional frequencies”, although Sect. 2.1 now defines the variable

as the monthly number/count of drought-related articles. I suggest using “monthly reporting frequency” or “monthly article count” consistently unless an actual proportional normalization has been applied.

Agreed. We checked Sect. 2.1–2.5 and the remainder of the manuscript for this inconsistency. Because the variable used in the current analysis is the monthly number/count of drought-related articles for each province and sector, Sect. 2.3 now consistently refers to monthly counts/reporting counts rather than “relative” or “proportional” frequencies. We also use “reporting frequency” in the interpretive text when referring generically to the temporal intensity of reporting.

L148–156: “To allow a robust comparison of the temporal dynamics of drought indices (SPI and SPEI) and the count of media reports on drought impacts, we standardized the time series of the monthly count of drought-related news across the provinces of Spain for the period 1976–2023. [...] Specifically, it deals with important issues of discrete and zero-inflated time series and provides a flexible yet rigorous framework to convert monthly reporting counts into standardized values.”

We again thank both reviewers for their careful assessment and constructive suggestions. The additional uncertainty analysis, zero-inclusive occurrence model, terminology harmonization, and explicit treatment of province-level attribution limitations have substantially strengthened the transparency and interpretive robustness of the manuscript.