

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. I believe these can be addressed without fundamentally redesigning the study, but some additional analysis or clarification would strengthen confidence in the main conclusions of the study.

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. The revised manuscript also appropriately acknowledges that sample support is weaker for some province–timescale combinations, particularly at the longest accumulation periods.

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. This is especially important because the regressions are fitted separately within relatively narrow drought-severity ranges. For example, the severe drought class spans only 1.28–1.65, so even with approximately 10–15 observations, the predictor range may be too restricted to estimate a stable slope in some provinces. 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.

A related point concerns the treatment of months with zero impact reporting. The revised manuscript clearly explains that months with zero reports are excluded from the log–log regression and that β therefore describes the elasticity of reporting intensity conditional on at least one report being present. This clarification is very helpful, but it also means that the analysis does not capture the transition from no reporting to some reporting - what I feel can be potentially important part of drought reporting response.

For example, if mild drought months usually produce zero articles, moderate drought months occasionally produce one article, and severe drought months consistently produce several articles, the 0 to 1 transition may represent an important part of the drought reporting response, but it is absent from the current β estimate. 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 bit more explicitly how conditioning on positive reporting may affect comparisons among drought-severity classes, sectors, and provinces.

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. The revised version of the manuscript now appropriately states that multi-impact, multi-location articles can therefore contribute to more than one province-sector combination and that the resulting maps should be interpreted as spatial patterns of drought-related reporting rather than perfectly geocoded impact events.

I think this framing is much clearer. Still, because province-level analysis remains an important component of the manuscript, I would still like some indication of how frequently this ambiguity occurs.

For example, consider an article reporting agricultural losses in Zaragoza and reservoir water levels depletion in Huesca. If agricultural and hydrological labels and province labels are extracted independently, the resulting representation may associate both impact types with both provinces. Such cases would matter much more if they were common than if almost all relevant articles referred to only one province. 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. If the general spatial structure remains similar, this would provide reassuring evidence that the main provincial patterns are not primarily generated by ambiguous multi-location reporting. Once again, I do not consider the extraction strategy itself invalid as the revised manuscript is now appropriately transparent about what it represents.

Minor comments

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.