

Reply to Revisions – Smith et al. - ESD

Referee #1

This manuscript addressed an interesting topic by investigating whether measurement noises can impact the inference of resilience using remote sensing data. They used a simulation approach to investigate how signal-to-noise ratio (SNR) influences the calculations of two indicators of resilience, namely lag-1 autocorrelation (AR1) and variance. Their results have implications for assessing the possible impact of measurement errors in observational data. Overall I found this study well designed and conducted. I have a few comments that may help improve the paper.

Thank you for your time with the manuscript and the helpful comments! We will address each comment individually below.

(1) the study generated simulated time series by combining a background time series with a number of random noises. I was wondering if this characterized the realistic errors introduced by changing instruments. I was not an expert in remote sensing, but I thought in some cases the change of instrument might induce sudden increase/decrease in the time series (rather than a small noise term). Such changes can have major impacts on the calculation of resilience indicators, keeping in mind that such indicators were used to detect 'sudden changes' in the time series, whether they were due to measurement errors or underlying processes?

Thank you for this comment. Sharp jumps are indeed a potential issue in time series obtained from combining signals from different sensors that are active over different time periods. There are many different studies on how and when to cross-calibrate long-term data to handle this. For example, in the VODCA data set that we attempt to model, biases between sensor absolute values are rectified with a CDF matching approach.

Given that problems arising from mean-shifts can be handled comparably easily, in our synthetic set up we address problems that go beyond a mean-shift. We assume that any composited multi-satellite data has had such shifts removed; our analysis rather focuses on what is *not* well treated in such common data compositing/calibration schemes – namely changes in higher-order statistics such as variance or autocorrelation.

(2) AR1 and variance are two important indicators of resilience, or early warning signals (EWS) for catastrophic changes, but there are more. Moreover, researchers had been developing composite EWS by combining different metrics. Given that measurement errors may influence AR1 and variance differently or in opposite directions, I was wondering if a composite EWS would be more robust to measurement errors.

You are right that some studies have proposed combinations of the Variance and AR1 coefficient, as well as other indicators such as Skewness or Kurtosis. In the present case, we focused on the Variance and Lag-1 autocorrelation as the underlying theory provides clear equations relating them to the recovery rate λ and – in this sense – to the resilience of the underlying system: $\langle x^2 \rangle = \sigma^2 / (2\lambda)$ and $\alpha(1) = \exp(-\lambda\Delta t)$. We believe that in order to obtain a detailed understanding of the effects of changing sensors on these indicators, it is best to keep them separated and not combine them, at least in the context of our study. Formulating a combined resilience indicator would make it more difficult to attribute the

individual effects of the changing sensor mix to changes in system parameters (e.g., system memory – autocorrelation, system variability – variance).

(3) the authors discussed about the difference between the average of variance from individual time series and AR1 of the aggregate time series, particularly their different behaviors in the presence of measurement errors. Similarly, the reference Feng et al. (2021) found different temporal trends of these two metrics. However, these two metrics represent different properties (i.e., local- vs. larger-scale resilience) and they did not necessarily exhibit different patterns, even if there was no measurement error. The problem is, the local-scale variance did not add up to give the larger-scale variance, but modulated by the synchrony between local grids. I attached a theoretical paper illustrating this:

Wang, S. & Loreau, M. Ecosystem stability in space: α , β and γ variability. Ecol. Lett. 17, 891–901 (2014).

Thank you for providing that reference – this is a very interesting piece of work! This is something that we did not really consider in our work: how does the variance of individual time series match up to the aggregate variance of many time series in the absence of problems due to combining data from different sensors. Our focus was rather on how resilience analysis has typically been presented (e.g., with regional mean time series), and what problems there might be with that approach given data with dynamic uncertainties. It is not unexpected, however, that the sum of many individual variances does not add up to a single mean-series variance – these two cases would be exposed to quite different noise levels (e.g., random fluctuations are suppressed in a multi-series mean). While we find this a very interesting problem for further exploration, we feel that multi-scale and spatially-conditioned changes in resilience are outside of the scope of this current work. We have, however, added in the revised discussion that scale should be considered for when examining differences between the variance/AR1 of aggregated time series and aggregated variance/AR1 time series.

(4) while the manuscript was overall well written, I had to say that I was confused by the different metrics involved in the figures, which seemed to be quite related but differ in important ways. For instance, the authors calculate resilience indicators using several approaches, e.g., deriving the numbers for an individual time series, first aggregating the time series and then calculating AR1 and variance, or first calculating AR1 and variance and then averaging them. They also calculate correlation between AR1 and variance at different levels of complexity. I would suggest to add a table to clearly define all key metrics in the figures, with explanations what a positive/negative or higher/lower value mean.

Thank you for this suggestion – we endeavored to keep our terminology as clear and explicit as possible, but it remains difficult to parse in some places! We have added a comprehensive table to help clarify our terms and interpretations.

Specific comments:

L52: What does 'synthetic series' mean? I think it is simply a simulated time series.

Yes, we refer here to simulated time series. To keep our language specific, we refer to each individual series as 'synthetic', and when we produce multiple realizations each is referred to as a single 'simulation'. We have tried to make this clearer in this Revision.

L79: How was this 'aggregating' implemented?

In each case of aggregation, we use either a time-explicit (e.g., daily) mean, or a time-explicit maximum value. To simulate VOD, we aggregate multiple instruments with a mean. To simulate AVHRR, we take the multi-week maximum, to match the method used in the original data set. For mixing multiple multi-instrument series, we use a time-explicit mean (e.g., averaging 100 synthetic series into a single averaged series). We have clarified this in our revision.

Figure 2: Not sure that I understood these figures correctly. Did the 'median corrcoeff median signals' on the right represent the median of the 'corrcoeff of median signal' on the left? Why they were so difficult, even by sign?

This is admittedly a rather dense figure. We think that the suggestion of the referee to add a table has greatly helped in describing the differences in the metrics and making this figure more accessible. On the left, AR1/Variance are first calculated, and the median of 1000 iterations is displayed. On the right, we take the median signal from 1000 simulations, then calculate AR1/Variance on that. We perform that calculation 100 times, and display the median of that result on the right panels. The difference between the columns comes from when the data was aggregated – aggregating many time series and then calculating AR1/Variance produces a significantly different result from calculating many AR1/Variance series and then aggregating those. We have updated the plot titles to be clearer here.

L130: "the correlation between AR1 and variance is generally positive for individual synthetic series" – any result supporting this argument?

This refers to the labels of the left hand column (Fig 2) and the histogram shown in Figure 4 – for higher SNRs, the median correlation coefficient between individual-series AR1/Variance is positive, while it is negative when those AR1/Variance series are instead aggregated. We have updated the reference on this statement to be clearer.

Figure 3: I must miss something. How to determine the SNR in the real data? And did it happen to exhibit SNR = 0.1 and 2 in the empirical data?

Determining the SNR of the real data is indeed challenging – there are many factors that influence sensor noise, and most change dramatically in space and time even for individual sensors. Cloud cover, atmospheric water content, and time-of-day of satellite overpass are a few key influences on local-scale sensor noise. Since we cannot quantify sensor noise effectively in space and time, we instead aimed to construct a synthetic experiment which mimicked the global-scale patterns seen in VOD (e.g., blue lines) using only noise-level changes (red lines). SNR is then a secondary factor – this allows us to vary not only the relative noise between sensors (e.g. through time) but also the amplitude of that noise relative to the underlying signal. The values of 0.1 and 2 are thus not related to the SNR of the real data, but rather constraints we placed on the synthetic data to illustrate the potential influence of not only changes in noise through time, but also how the amplitude of

that noise plays a role. Note that we deliberately chose very low values for the SNRs; in reality, a sensor with SNR equal to 0.1 will of course not be considered useful, and even SNR = 2 is still far lower than what is reported for operational sensors. MODIS aims for SNRs of individual bands between ~70-1000 (<https://modis.gsfc.nasa.gov/about/specifications.php>); Landsat 8 OLI aims for SNRs of at least 100 (<https://landsat.gsfc.nasa.gov/satellites/landsat-8/spacecraft-instruments/operational-land-imager/oli-requirements/>); Sentinel-2 aims for SNRs around ~100 ([https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-2-
msi/resolutions/radiometric](https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-2-
msi/resolutions/radiometric)). Our values of SNR from 0.1-2 are rather related to our synthetic experiment than trying to match real-world data.

L242: what does it mean by 'aggregated'? You had explained that first aggregating time series and then calculating AR1 and variance can remove the influence of changes in satellite instruments to some extent. So did you mean 'first calculating AR1 and variance and then taking the averaging of these metrics'?

Yes, in this case we mean that when many AR1/Variance series are averaged, changes in the noise structure through time are emphasized. Conversely, first averaging many series and then calculating AR1/Variance on the resultant averaged series reduces the impact of changes in sensor noise through time. We have clarified the use of the term "aggregation" throughout the manuscript in our revision.

Referee #2

The impact of data noise on estimating two resilience metrics, variance and lag-1 autocorrelation, with satellite data was assessed in this study. The topic addressed is very importance, because satellite products are widely used to quantify the resilience of terrestrial ecosystems. My major concern is that it is within our expectation that data noise will affect the reliability of the metrics, what's the new finding of this study? I hope two aspects may be investigated in depth: 1. What's the uncertainty of the existing satellite products when used for quantifying resilience? For this purpose, the 'real noise' of the data needs to be quantified. 2. What's the uncertainty of using the products to depict the temporal changes in ecosystem resilience? For this purpose, the temporal changes in the noise are also need to be quantified.

Thank you for your comments and the time spent with the Manuscript. We will address your three main points individually.

1. What's the new finding of this study?

As you note, satellite products are widely used to quantify the stability and resilience of different natural systems. These estimates of stability changes are based on the assumption that changes to higher-order statistics (e.g. variance, autocorrelation) are due to changes in the system under observation, and not in the observation mechanism. However, this assumption is not always true for satellite data – many nominally continuous data sets are in fact made up of a constellation of sensors. You are clearly right that it is within our expectation that data noise will affect the reliability of resilience metrics. As satellite data is used more and more often in these contexts, however, we felt it was important to explore how exactly changes in measurement procedures could propagate into resilience estimates, and what strategies might be used to minimize this issue. This is why we set up a thorough investigation of the detailed effects of combining signals from different sensors, using synthetic time series constructed for a wide range of possible signal-to-noise ratios. We further compared our controlled synthetic experiment to three real-world data sets, showing that while individual time series might be reliable, averaging individual time series of resilience indicators over large regions will tend to enhance the effects of measurement changes and thus reduce the reliability of resilience estimates. We feel that this is an important contribution to best practices when analyzing resilience based on satellite data. We have clarified the above points in our revised manuscript.

2. What's the uncertainty of the existing satellite products when used for quantifying resilience? What's the uncertainty of using the products to depict the temporal changes in ecosystem resilience?

Answering this question is exactly the reason we conducted our investigation. We undertook our study using synthetic data because we cannot fully constrain real-world data – satellite data noise shifts drastically in space and time, and we have no 'perfect' measurement against which to compare as a ground truth. Satellite instruments have reported signal-to-noise ratios and are calibrated against known quantities (e.g., deep space, the Sahara); however, these do not account for all sources of noise in, for example, a vegetation measurement. Other factors – such as atmospheric water content, cloud cover, and satellite viewing angles – will also influence estimated surface parameters. Without a true reference, dis-entangling the time-variable (e.g., seasonal, annual) changes in noise from changes in system stability is not possible with real-world data. Hence, we focused here on a controlled

synthetic system to explore the potential influences of changes in measurement procedures through time.

The lack of true reference also limits our ability to quantify the uncertainty in changes in resilience in real-world applications. We instead attempted to mimic a change in resilience proxies (autocorrelation and variance) using only changes in measurement noise; this could be thought of as a null model for whether or not there was a change in resilience. If changes in autocorrelation and variance are stronger than those implied by the changes in the underlying noise, than that could be interpreted as a more robust resilience change signal. Again, however, in a real-world system we would need a time-explicit estimate of how signal-to-noise ratios are changing to control for the influence of measurement procedure on our resilience estimates.

3. the 'real noise' of the data needs to be quantified ... the temporal changes in the noise also need to be quantified

While we fully agree with your desire for time-explicit noise models for different satellite data, that is not the problem we set out to address with this publication. We instead approached the problem from a synthetic perspective, where we could control the amount and timing of noise, as well as how and when different sensors were mixed. Our work emphasizes the need for better constraints on the noise levels of satellite products and demonstrated the influence of varying signal-to-noise ratios on different ways to estimate resilience. It does not aim to provide a thorough accounting for different remote sensing products. We have revised our introduction to make this clearer.

List of Major Changes

1. Added additional text around the issue of real satellite reliability.
2. Added a Table (Table 1) to better explain the different metrics, aggregation schemes, and their implications in a concise way.
3. Updated the headings of Figure 2 to make them easier to understand.
4. Various text and reference changes for clarity.