Trends in Extreme Precipitation and Streamflow

Caren Marzban[1,2]*
Dennis Lettenmaier[3]
Laura Bowling[3]

[1] Center for Analysis and Prediction of Storms
University of Oklahoma, Norman, OK 73019

[2] Department of Statistics, University of Washington, Seattle, WA 98195
[3] Department of Civil Engineering, University of Washington, Seattle, WA 98195

Abstract

Studies of precipitation trend suggest that both the amount and frequency of precipitation across the U.S. have been increasing over the last century. Although a similar result has been found for streamflow across most of its distribution, extreme streamflows appear to follow a down trend. Here, the question of whether there is a contradiction between trends in precipitation and streamflow will be addressed. It will be argued that the question can be addressed when only "causally connected" precipitation and streamflow events are examined for trends. It is shown that in situations where statistically significant results are found there is no contradiction. (Be more specific??).

1 Introduction

Bla bla (Dennis??).

The main aim of this study is to examine the *relationship* between the *trend* in precipitation and that of streamflow. The outline of the paper is as follows: Section 2 offers a summary of previous work on trends in precipitation and streamflow. Section 3 describes the data and outlines the method of analysis. Section 4 presents the results, and is followed by a conclusion and discussion section.

---

*Corresponding author: http://www.nhn.ou.edu/~marzban

2 Background

It must be acknoweldeged that most physical time series are such that a linear (or even monotonic) trend analysis can simply overlook interesting nonlinear behavior. Precipitation and streamflow are no exception. They both display significant nonlinear behavior on different time scales which can shed light on the underlying physical processes. However, a linear trend analysis is often the starting point for more sophisticated analysis, for it too can shed light on the underlying processes. Most existing literature on precipitation and streamflow trends assume a monotonic trend, and we too shall do the same.

The trends for precipitation and streamflow, separately, have been thoroughly examined in the literature. The results depend on a host of issues, like the region of the country, season, the percentile of the distribution, etc. Many different measures have been utilized as well (proportion of the country affected by some type of precipitation, number of days with precipitation, amount, intensity, etc.). Here, we look only at the daily amount.

Figure 1 is a graphical presentation of the results of Lins and Slack (1999), specifically their Table 1. The x-axis labels the percentiles of the distribution. The solid curve is the percentage of significant trends. The remaining curves represent the percentage of significant downtrends (red) and significant uptrends (green). From top to bottom the length of the data increases from 30 years to 80 years in 10-yr increments. The many features of these graphs are fully discussed in the original paper; here, we simply point out that the convergence (and even intersection) of the uptrend and downtrend curves at the highest percentiles. For example, according to the 30yr record, the percentage of stations with a downtrend is always higher than that of the stations with an uptrend, except for the higher percentiles of the distribution, where that pattern reverses. For longer record lengths, this reversal does not always occur, but the difference between the two curves does always diminish at the higher percentiles. Of course, without error-bars on the curves it is impossible to ascertain whether or not the reversal of the trends (between the upper and lower percentiles) is statistically significant. However, the various plots in Figure 1 provide some evidence for the possibility that the upper percentiles display a different pattern of trends than the middle and lower percentiles.

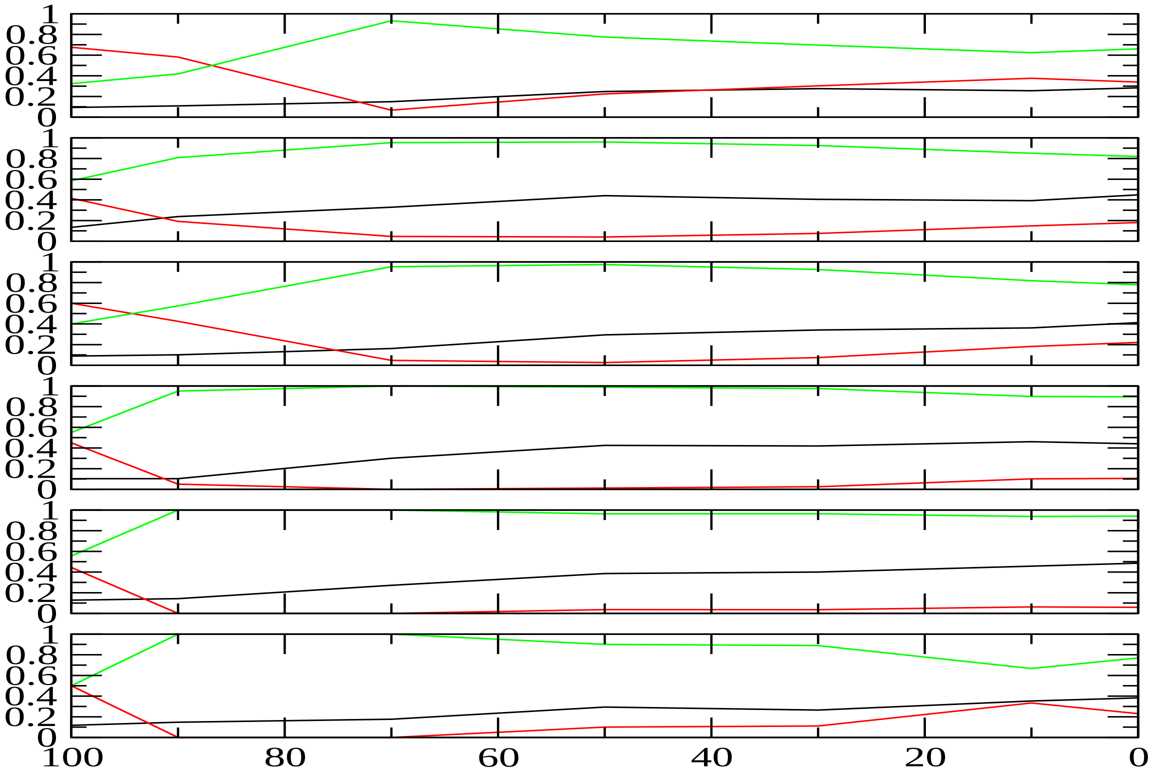

Fig. 1: Streamflow trends at different percentiles of the distribution according to Lins and Slack (1999). See text for details.

Lins and Slack also performed a similar analysis for different regions of the U.S. Their results can be coarsely summarized as in Table 1, where "mixed" reflects a pattern of flows that does not clearly disambiguate between downtrends and uptrends.

Annual Max.		Median		Annual Min.	
West	East	West	East	West	East
down	mixed	down	up	down	up

Table 1: A coarse summary of the results of Lins and Slack (1999). Unnecessary??

The precipitation results are stated in Karl and Knight (1998), as well as in numerous other studies (refs??). The former serves for a point of comparison with the current work. They find a general pattern of increasing trends in the U.S. over the last century. The exception is a downtrend in Winter flows in several regions of the country. However, the Winter results are not statistically significant. Figure 2 is a restatement of their annual and nationwide result. The trend is expressed as percent of mean precipitation per century, and the x-axis represents the different percentiles of the distribution in 5% increments. A coarse summary of the results for different regions and seasons is presented in Table 2.

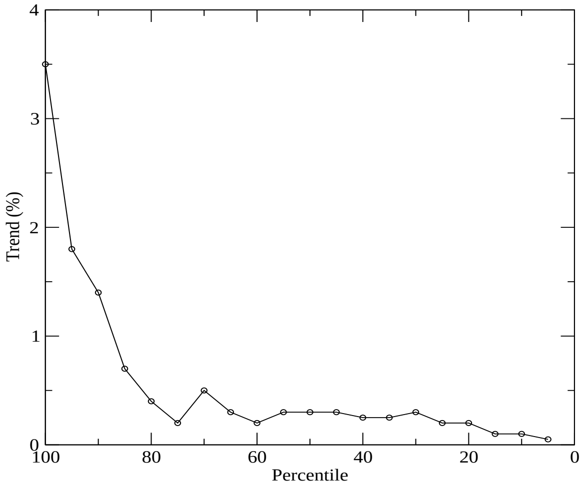

Fig. 2: Precipitation trend over the entire U.S. at different percentiles of the distribution (Karl and Knight, 1998).

	95^{th} percentile		50^{th} percentile	
	West	East	West	East
Winter	mixed	mixed	down	mixed
Spring	up	up	up	up
Summer	up	up	up	up
Autumn	mixed	up	up	up

Table 2: A coarse summary of the results of Karl and Knight (1999). Results for the lowest percentiles are not shown. Unnecessary??

A comparison of Tables 2 and 3 is not straight forward since the former is not seasonally partitioned. However, it would appear that the Spring and Summer uptrends in precipitation, across the U.S. and in both the upper and middle percentiles, are not matched by an analogous uptrend in streamflow. The exception is in the Eastern U.S., where a general pattern of precipitation uptrends is reflected by a similar pattern of streamflow uptrends. As mentioned in the Introduction section, it has been noted that the lack of a correlation between precipitation and streamflow trends is inconsistent with the hydrologic cycle. On the other hand, Lins and Slack suggest that there is no inconsistency, and that the increase in precipitation may simply be insufficient to cause flooding.

Groissman et al. (2001), too, ascertain that there is no inconsistency, but their method is indirect. Their argument goes as follows: First, they show that the data provide evidence for an uptrend in precipitation. Second, they argue that the uptrend is accompanied by an increase in snow cover, and as such is not expected to yield an uptrend in streamflow. Finally, they confirm that the data supports both the increase in snow cover and the downtrend in flow.

In the following sections a method will be put forth that allows one to examine the trends in "causally connected" precipitation and streamflow events. In understanding the relationship between the trends in precipitation and streamflow, jointly, only such events should be considered. In other words, precipitation events that for whatever reason do not yield streamflow events should not be included in an analysis of the relationship between precipitation and streamflow trends. Similarly, streamflow events that are not "caused" by some precipitation event must be excluded from a trend analysis.

It has been shown that the spatial correlations present in the data can artificially magnify the significance of trends. As such, it must be acknowledged that the statistical significance of the findings herein is likely positively optimistic.

3 Data and Method

The daily data for 19,438 streamflow stations is obtained from four USGS CDs. The flow data span the years 1889 to 1997.

As for precipitation data, daily data for 1,061 stations is obtained from National Climate Data Center (NCDC) (also at http://www.hydro.washington.edu/Lettenmaier/Data/met_data.html.) The precipitation data span the time period 1877 to 1997.

Although a full range of flow and precipitation amounts is examined, the primary emphasis of the current study is the trends in their extremes. To that end, the framework of Peaks Over Threshold (POT) will be adopted (refs??). As such, an event in a time series whose trend is being examined is defined as a data value that exceeds a threshold. This applies to both precipitation and flow events. In what follows, the threshold itself is varied as to allow an average of one, two, three, etc., events per year. Given that a given precipitation or streamflow occurrence can last several days, this definition of an event leads to a dependency among events. Most tests for trends, however, assume a time series consisting of independent events. As such, at the stage in the analysis when trends are to be computed and tested, the definition of an event is revised to refer to the maximum of a *cluster* of data values exceeding the threshold. A cluster, in turn, is defined as a sequence of adjacent data values exceeding the threshold. It is assumed that the clusters are independent. (Confusing??)

Only flow stations with a drainage area less than ?? are selected. These are believed to be unaffected by human activity. Additionally, for a given time interval (e.g., 30 yr, 60 yr) the stations are required to have no missing data in the first and last year of the interval. The existence of missing data within the interval is not detrimental to the analysis, because the measure of trend (Kendall's tau) is independent of the time spacing of the events.

The selection of the precipitation stations is a bit more involved. First, only those stations that are unaffected by snow (during the Season of interest) are retained. Then, only those which fall within the basin of some pre-selected flow station are selected. Each basin is approximated by an ellipse, and Laura??

It was found that many precipitation stations residing outside of the ellipse still affect the flow at the corresponding flow station. The question then becomes one of determining how far a precipitation station can be from the centroid of a basin before it can be selected for analysis? The answer requires further analysis which will be outlined next.

Central to the goal of this article is the premise that not only precipitation and flow stations can be naturally paired together, but also individual flow and precipitation events at the paired stations can be paired. In fact, both tasks can be done simultaneously, and Figure 3 shows how. The black curve in Figure 3a is the conditionally probability of a flow event at a flow station, given that a precipitation event occurred some days before at a precipitation station. The number of days between the precipitation and flow events, called the lag, is plotted on the x-axis. It can be seen that the probability is low for small lags, peaks at around 12 days, and asymptotically falls off to a constant value (modulo some "noise"). Figure 3b (??) shows the same quantities but for a different pair of flow and precipitation stations. It is then natural to conclude that the pair of stations examined in Fig. 3a are "causally" connected, while the pair in Fig 3b are independent. It is the distinctive feature of the black curve in Fig. 3a that is utilized to pair (or match-up) precipitation with flow stations. The red strip is the two-sigma interval based on the interval $0 \leq lag \geq 250$. When the black curve exits the red region, it is likely that the effect is statistically significant.

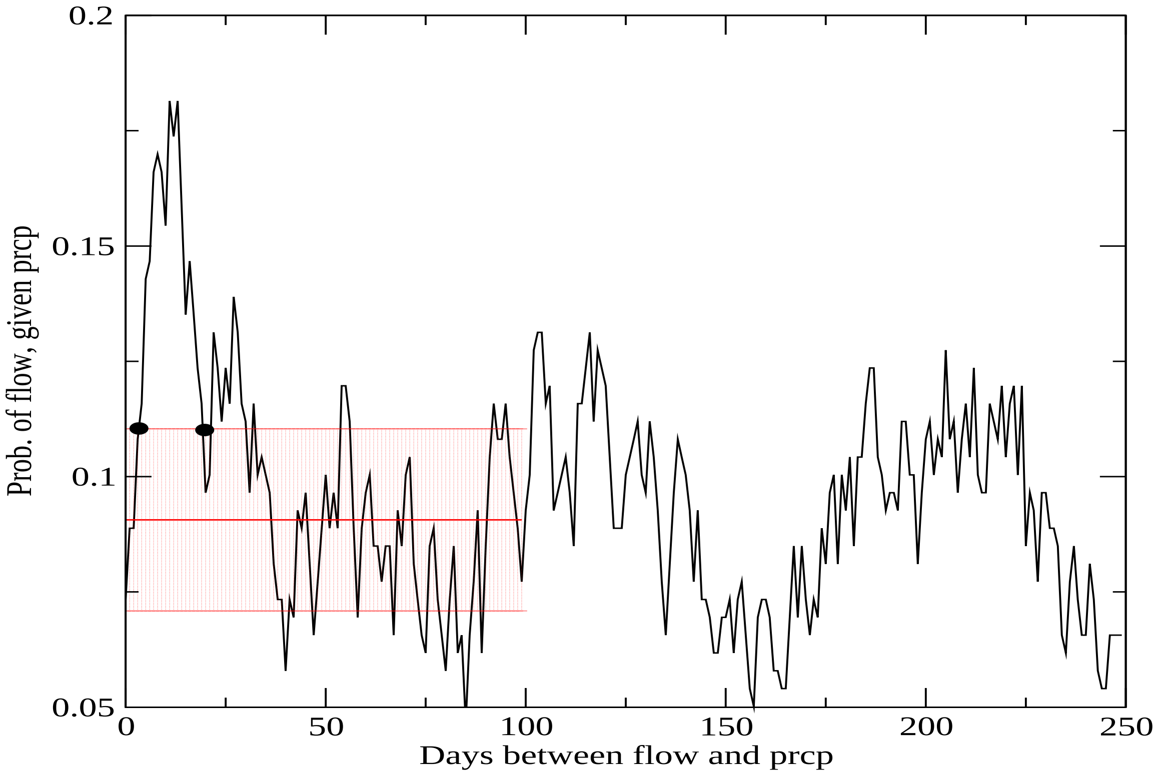

Figure 3. The conditional probability of a flow event, given a precipitation event as a function of the number of days between the two events.

Fig. 3a also suggests a natural rule for pairing individual precipitation and flow events as well. The lag at which the curve peaks (lagmax) is the number of days between precipitation and flow events at which the probability of a flow event is maximum, given a precipitation event. As such, flow and precipitation events that are precisely lag days apart are naturally (or loosely speaking, causally) connected. It is these matched pairs of precipitation and flow events whose trends are the subject of this analysis. Actually, given that many events exceeding a threshold are not independent, the trend analysis is performed on not all pairs that are lag days apart, but only on those belonging to an independent cluster (see above). (Confusing??) [1]

The above scheme for pairing stations and events induces a number of different time series whose trends can be examined. For example, one can examine the trend in the time series of all precipitation events at a station that is matched with some flow station. Conversely, one can estimate the trend of all flow events at a flow station that is matched with some precipitation station. The sample size of each of these

---

[1] In a further attempt to isolate only the "causally connected" events only precipitation events that have a high Antecedent Precipitation Index (refs??) are included in the trend analysis. More??

time series is given in the first two rows of Table 3. However, given that we are interested in the hydrologic cycle, it is then natural to examine the trends in only the precipitation and flow events that are matched. The sample size for these time series are suffixed with "_pair" in Table 4. It is worth emphasizing that the events in these paired time series are subsets of the original time series. As such, they are of smaller size, but they are more significant in addressing the relationship between precipitation and flow trends.

Finally, the match between precipitation and flow stations is not one-to-one. Many flow stations can be matched to a single precipitation station, and vice versa. Each match between precipitation and flow stations induces a time series at each station whose trends are to be computed. The number of these "links" between stations is given on the last row of Table 4. [2]

nyrs	30 1968-	40 1958-	50 1948-	60 1938-	70 1928-	80 1918-
nprcp	440	440	437	221	107	62
nflow	139	139	139	82	51	29
nprcp_pair	410	421	414	217	105	61
nflow_pair	132	136	136	81	49	28
nlinks	3763	?	?	1564	?	?

Table 4. Number of precipitation and streamflow stations, paired and otherwise, for different record lengths. (Update??)

The geographic distribution of the paired stations is shown in Figure 4. It's worth emphasizing that this map depends on a number of parameters including season, record length, and even NPOT, because the method for pairing precipitation stations with flow stations depends on those parameters. The map in Figure 4 is for Summer, a 30-yr record, and NPOT=2. Discuss more??

---

[2] It is instructive to step through this process of pairing). For example, for the 30-yr time interval and for Summer, i.e., the first column in Table 4, Initially, there are 1061 precipitation stations, 955 of which fall within the basin of flow stations with a drainage area of less than ??. Of these 955 precipitation stations 677 are unaffected by snow during Summer, and of those 619 have no missing data in the 30-year time interval. These 619 precipitation stations are "causally" related to 208 flow stations, but only 139 of these flow stations pass the missing data criterion, which in turn reduces the number of paired precipitation stations to 440. Note that this selection scheme allows for two or more precipitation stations to be paired to a single flow station.

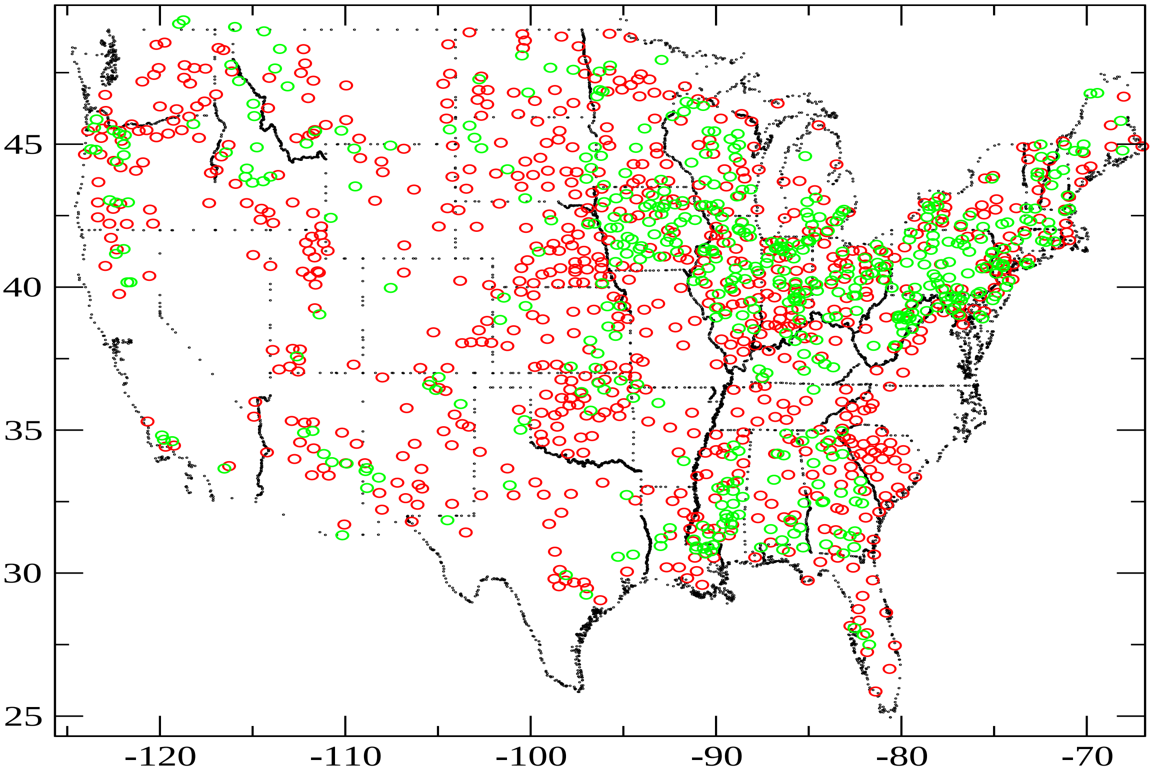

Figure 4. The geographic distribution of the paired precipitation (red) and streamflow (green) stations.

Kendall's tau is utilized to assess the statistical significance of a trend. Then for a given geographic region and season, the percentage of station time series with a statistically significant trend at the 0.05 level is computed. Similarly computed is the percentage of those with a down- and an up-trend. For the flow trends, these are analogous to those computed by Lins and Slack (1998) (in their Table 1), and will allow a comparison of the present results with theirs. The same set of percentages is also computed for the paired time series.

Although, all of these percentages carry a wealth of information, in understanding the relationship between flow and precipitation trends the most relevant quantity is the joint probability of up and down trends. Specifically, we define the p_{--} as the joint probability of obtaining a down-trend in flow and a downtrend in the matched precipitation time series. Similarly, p_{-+} is the probability of a downtrend in flow and an uptrend in the matched precipitation time series; etc. A perfect correlation (association) between flow and precipitation trend would yield a diagonal p matrix.

A measure of association that also takes into account randomness is κ (kappa)

(Wilcox, 1996 p. 358). It is based on the proportion of agreement, $p_{--} + p_{++}$. It is defined so that a perfect positive association yields $\kappa = 1$, and a perfect negative association gives $\kappa = -1$. A chance association yields $\kappa = 0$. The reason this particular measure is chosen is the availability of analytic formulae for computing the standard error of κ (Wilcox, 1996 p. 359).

4 Results

The analysis in this subsection is based only on flow stations that pass the vicinity test and the snow test for being selected. As such, the results may not be identical to those found in other reports.

Before proceeding with the general analysis, it is instructive to consider one specific question that was raised above, namely the significance of the aforementioned reversal of the streamflow uptrends and downtrends between upper and lower percentiles of the distribution. Figure 5 displays these trends for the stations that pass our criteria; the x-axis is the average number of POTs, and the error-bars are standard errors. This figure is for flows during the Summer season and for the entire U.S..

These figures can be compared with those in Fig. 1. First, note that NPOT=1 in Fig. 5 implies that only the highest flows are selected, while maintaining an average of one flow event per year. As such the left-side of these figures explores the highest percentile of the flow distribution. However, the lowest percentiles of the distribution do not correspond to high NPOT values, which simply allow all flow events into the analysis.

For both the 30yr record and a 60yr record the overlap of the error-bars in the central region of the graphs suggests that the difference between the percentage of stations with a significant downtrend and those with a significant uptrend is not statistically significant, except for very low and very high NPOT values. A reversal akin to those noted in Fig. 1 is observed. In other words, the extreme streamflow events (i.e., the left side of the graphs) are more likely to have a downtrend than an uptrend. However, for flows of any strength (right side of the graphs) that pattern is reversed. There exists an anomalous reversal at NPOT=5 in the 60-yr record which may not be statistically significant, given that the error-bars are only standard errors. Other record lengths too noisy??

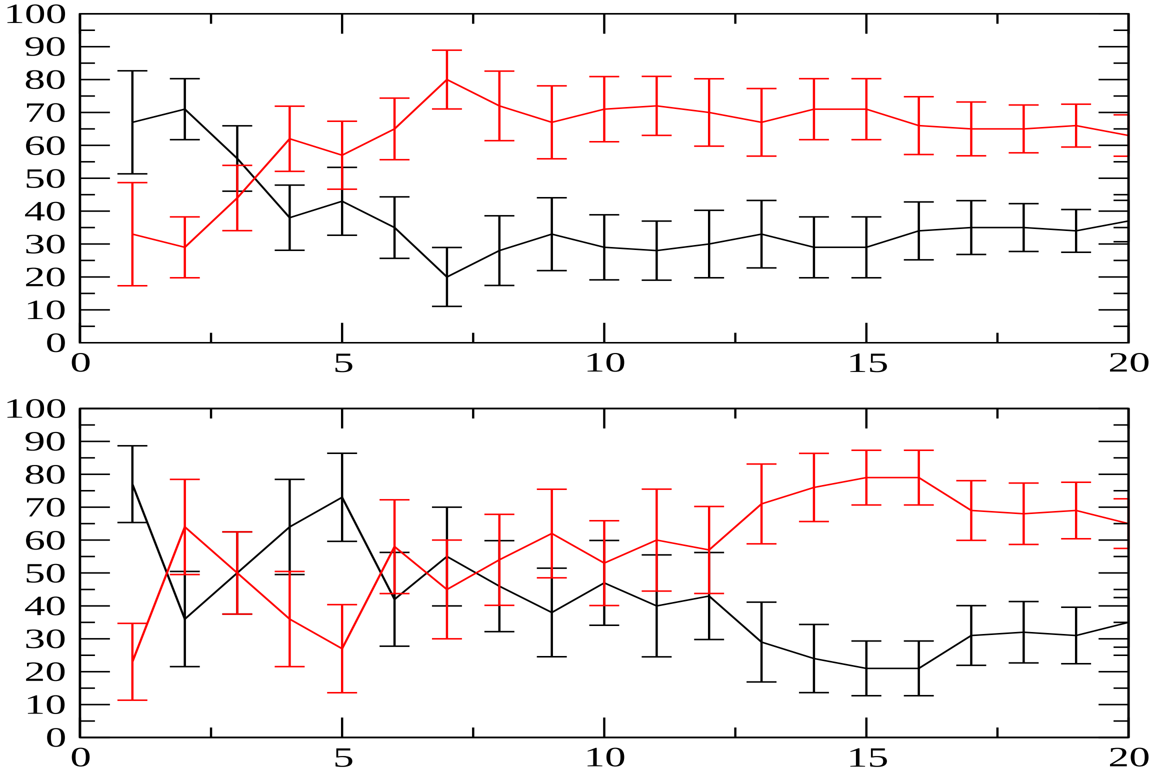

Figure 5. Percentage of stations, nationwide, with a significant downtrend (black) and uptrend (red), during Summer. The upper (lower) figure is based on a 30yr (60yr) record.

Figure 6 shows analogous precipitation (top) and streamflow (bottom) results for the different seasons and regions of the nation based on the 30-yr and 60-yr records. The examined Seasons and regions are Spring and Summer, and East and West, respectively. Winter and Autumn results are mostly nonsignificant because the number of stations unaffected by snow is too small. Similarly, a finer partition of the U.S. leads to excessively small number of stations. These graphs are not ideally suited for examining trends in streamflow or precipitation, separately, because the number of stations is drastically reduced upon enforcing the vicinity criterion. The only reason these figures are included here is to make contact with the work of Lins and Slack (1999) and to show that our results regarding flow are consistent with theirs. Recall that the primary aim of the current work is to examine the *relationship* between flow and precipitation trends, not the trends in each, separately.

Since the current treatment deals with 2 seasons (Spring and Summer), 2 geographic regions (East and West), and 2 time intervals (30yr and 60yr), there are numerous comparisons that can be made. Furthermore, the analysis is performed separately for precipitation events, streamflow events, and only causally connected precipitation and streamflow events.

We begin by examining the 30yr record (Figure 6); the top block of four figures pertain to precipitation events, while the bottom block refers to flow events. Within each block, the left (right) column refers to Spring (Summer), and the top (bottom) row is for the East (West). The black (red) curve is the percentage of stations with significant downtrends (uptrends), and the error-bars are standard errors based on Pratt's test (199??).

It can be seen from the precipitation block that the difference between the two curves is mostly nonsignificant. The exception is the West in Summer, where there is a dominance of downtrends across all NPOT values. The flow block displays more statistically significant features, especially at higher NPOT values. Stations with uptrends outnumber those with downtrends. The exception is the East in Spring time, where the reverse is true. Note that an overlap of the error-bars of one curve with those of the other curve does not imply that there is no trend. It simply implies that there is insufficient evidence for rejecting that hypothesis (i.e., that the data cannot tell us anything about the direction of the trend).

The comparison of the precipitation block (Fig. 6, top) and the flow block (bottom) reveals the aforementioned inconsistencies with the Hydrologic cycle. In particular, In the West, at Summer time, a dominance of downtrends in precipitation is difficult to reconcile with a dominance of uptrends in streamflow. The other regions and seasons display similar inconsistencies, albeit not to a statistically significant degree.

Figure 7 shows analogous results but for causally connected precipitation and flow events. In other words, precipitation and flow events that are not causally connected have been excluded from the analysis. As such, each block offers little useful information regarding the respective events, separately. It is the comparison of the blocks which is important in this figure. It is interesting that a wide separation between the curves in the precipitation block is matched by a narrow separation between the curves in the flow block, and vice versa. In fact, this feature resolves the aforementioned inconsistency; when there is a statistically significant dominance of precipitation stations with uptrends, the causally connected flow events display no statistically significant trend (up nor down). Conversely, when there is a clear downtrend in flow (East, Spring), there is no clear trend in the matched precipitation events.

One can be cavalier and examine these figures beyond the error-bars, i.e., ignoring them. In that case, the curves in the flow block suggest a general pattern of flow uptrends, with the exception of the East at Spring time. This pattern is matched by the precipitation curves (top). As such, an uptrend in precipitation does lead to an uptrend in streamflow. Therefore, again, there is no inconsistency. However, Spring time precipitation and flow events in the East do pose an anomaly, in that even ignoring the error-bars does not yield a consistent pattern of precipitation and flow trends especially for higher NPOT values. Note that the curious (and undesirable) crossing of the two curve for NPOT values between 3 and 5 in Figure 6a disappears in Figure 7a, i.e., when precipitation and flow events are causally paired.

The 60yr record provides a mostly similar conclusion. For precipitation, the Summer figures (Figure 8,a,c) are similar to those based on the 30yr record (Figure 6a,c). The Spring figures differ, but not to a statistically significant degree. As for streamflow, the same is true, except that the spring time flows in the East suggest an uptrend based on the 60yr record (Figure 8,c) in contrast to Figure 6,c.

There are at least two explanations for this change (from downtrends in the 30yr record, to uptrends in the 60yr record). Given that the error-bars are standard errors, and the proximity of the two curves in Figure 7c, one may argue that the difference between the two curves is not statistically significant at a reasonably high level, e.g. 90% or 95%. However, the method employed for computing the error-bars (due to Pratt, 19??) provides only a standard error and no confidence intervals. So, the validity of this explanation cannot be tested.

On the other hand, it is entirely possible that the time series for streamflow behaves differently over the two time intervals. This would imply that the trend is nonlinear or nonstationary. Although a nonlinear trend analysis is beyond the scope of the current project, such a trend is not a detriment to the goal of the current work, for here we are interested in the relationship between precipitation and flow linear trends.

Proceeding with the trends of the causally related events in the 60yr record, Figures 9a-d show the relevant curves. Unlike the curves for the unrelated events (Figs 8a-d), these figures speak of an uptrend in the causally related precipitation events for both seasons and regions. The corresponding flow events follow a similar uptrend when there exists a statistically significant difference between the up- and downtrends. The only place where this is not true is in the East at Summer time, where the uptrends in precipitation are accompanied by a downtrend in the corresponding flows, at least in the mid-range NPOT values (Fig. 9f). In short, there is a consistent pattern of trends between precipitation and the corresponding flow events, with a small anomaly in Summertime flows in the East. The consistency is dramatically demonstrated by the complete reversal of the trends in Figures 8e and 9e.

Finally, it is possible to quantify the connection between precipitation and flow events using a measure of association. Figures 10 show the values of κ and its standard error for all NPOT values. Recall that a significant association between precipitation and flow implies that the error-bar will not intersect the $\kappa = 0$ line. According to the 30-yr record (Fig. 10, top), the West in general displays no significant association between precipitation and flow trends. The East, however, shows a significant association between the trends, but only for lower NPOT values (i.e., more extreme events). The 60yr record (Figs. 10e-h) is similar, but the error bars are too large for a meaningful interpretation.

Plot p++,p+-, etc. for at least one region and season??

5 Conclusion and Discussion

Bla bla.

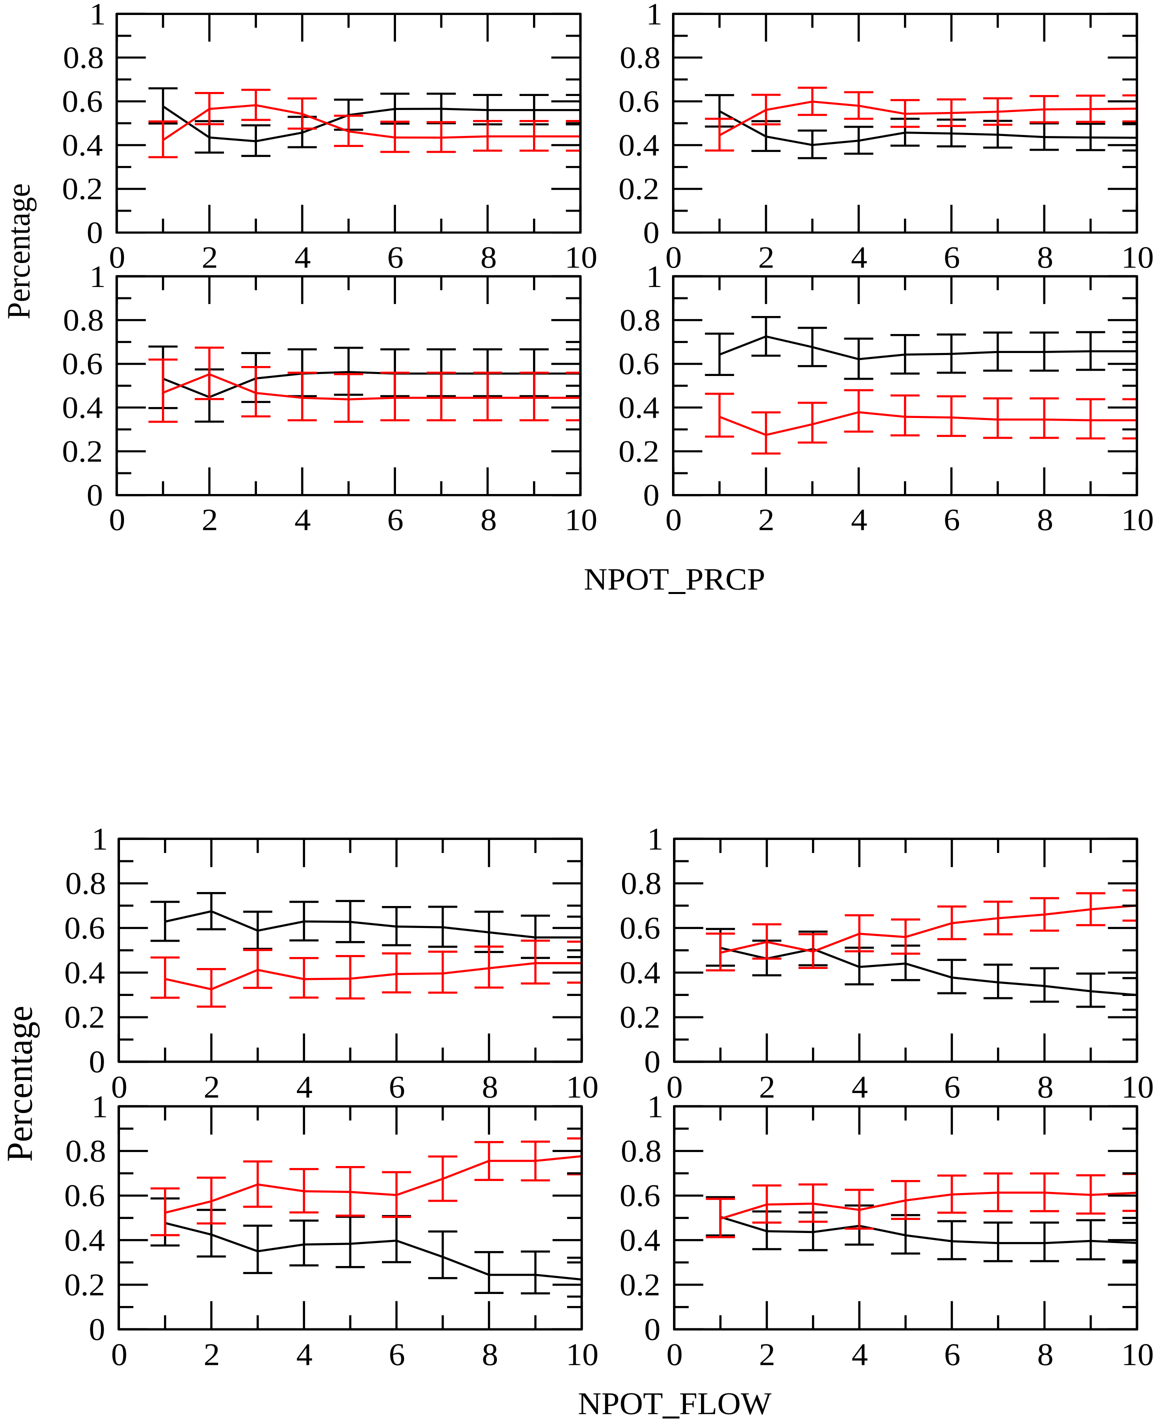

Figure 6. Percentage of precipitation (a-d) stations with downtrends (black) and uptrends (red). Percentage of streamflow (e-h) stations with downtrends (black) and uptrends (red). Record length = 30yr.

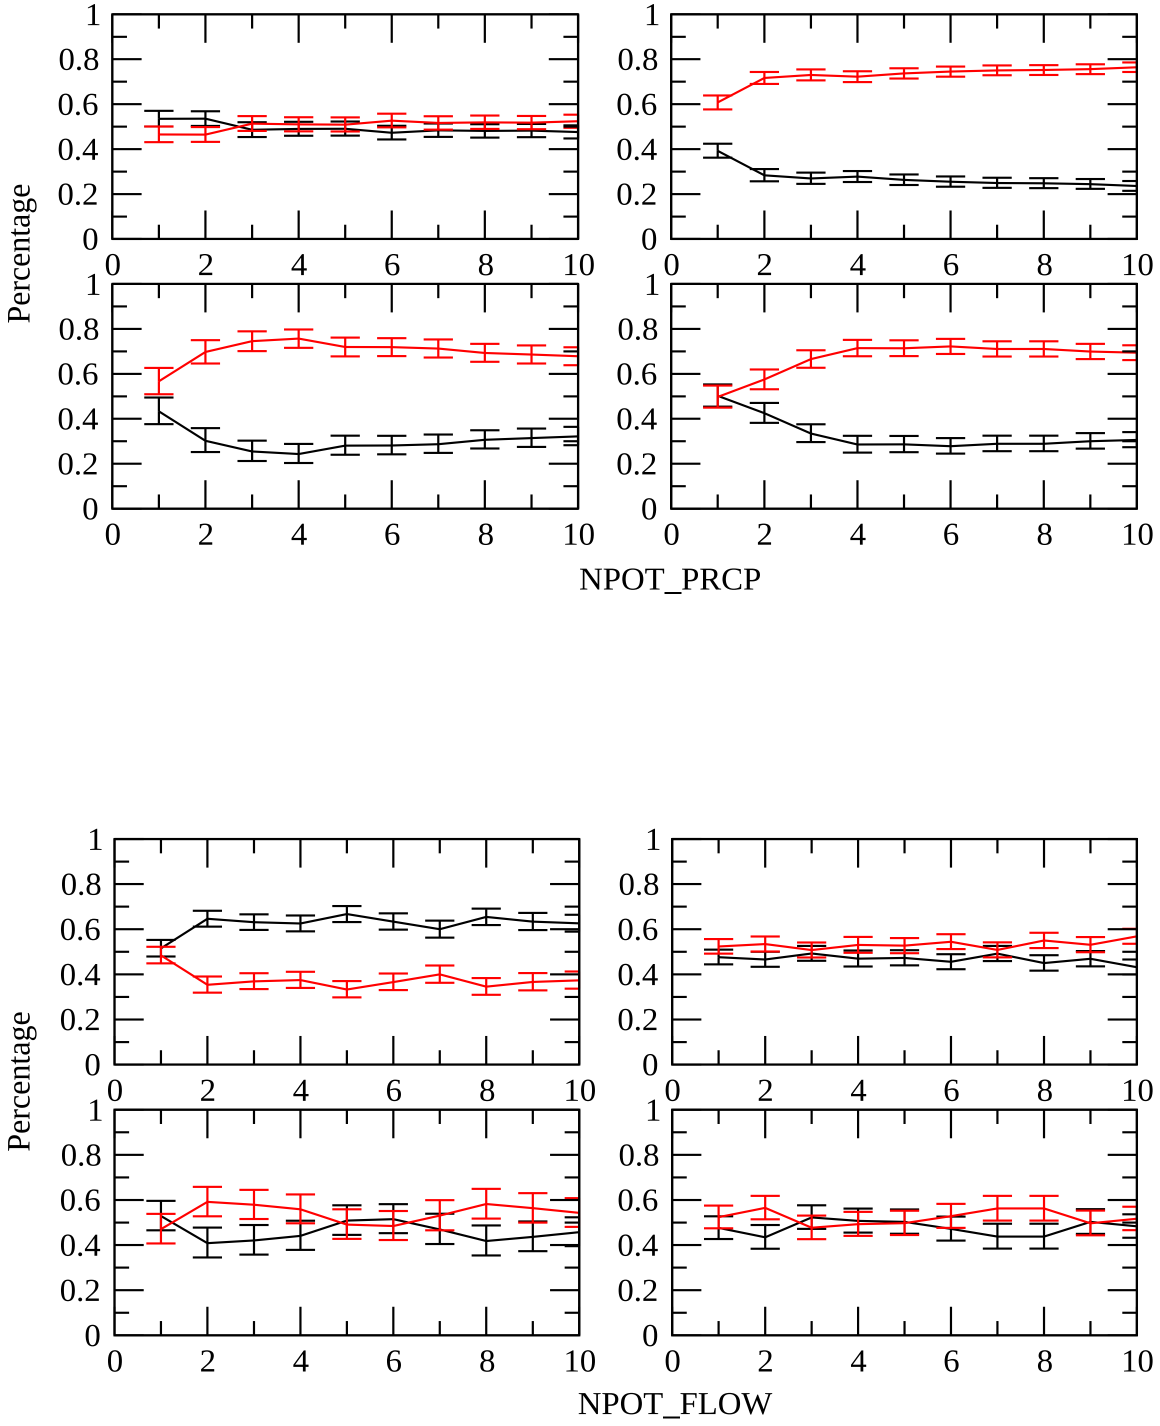

Figure 7. Percentage of precipitation (a-d) stations with downtrends (black) and uptrends (red), and percentage of streamflow (e-h) stations with downtrends (black) and uptrends (red), for paired events. Record length = 30yr.

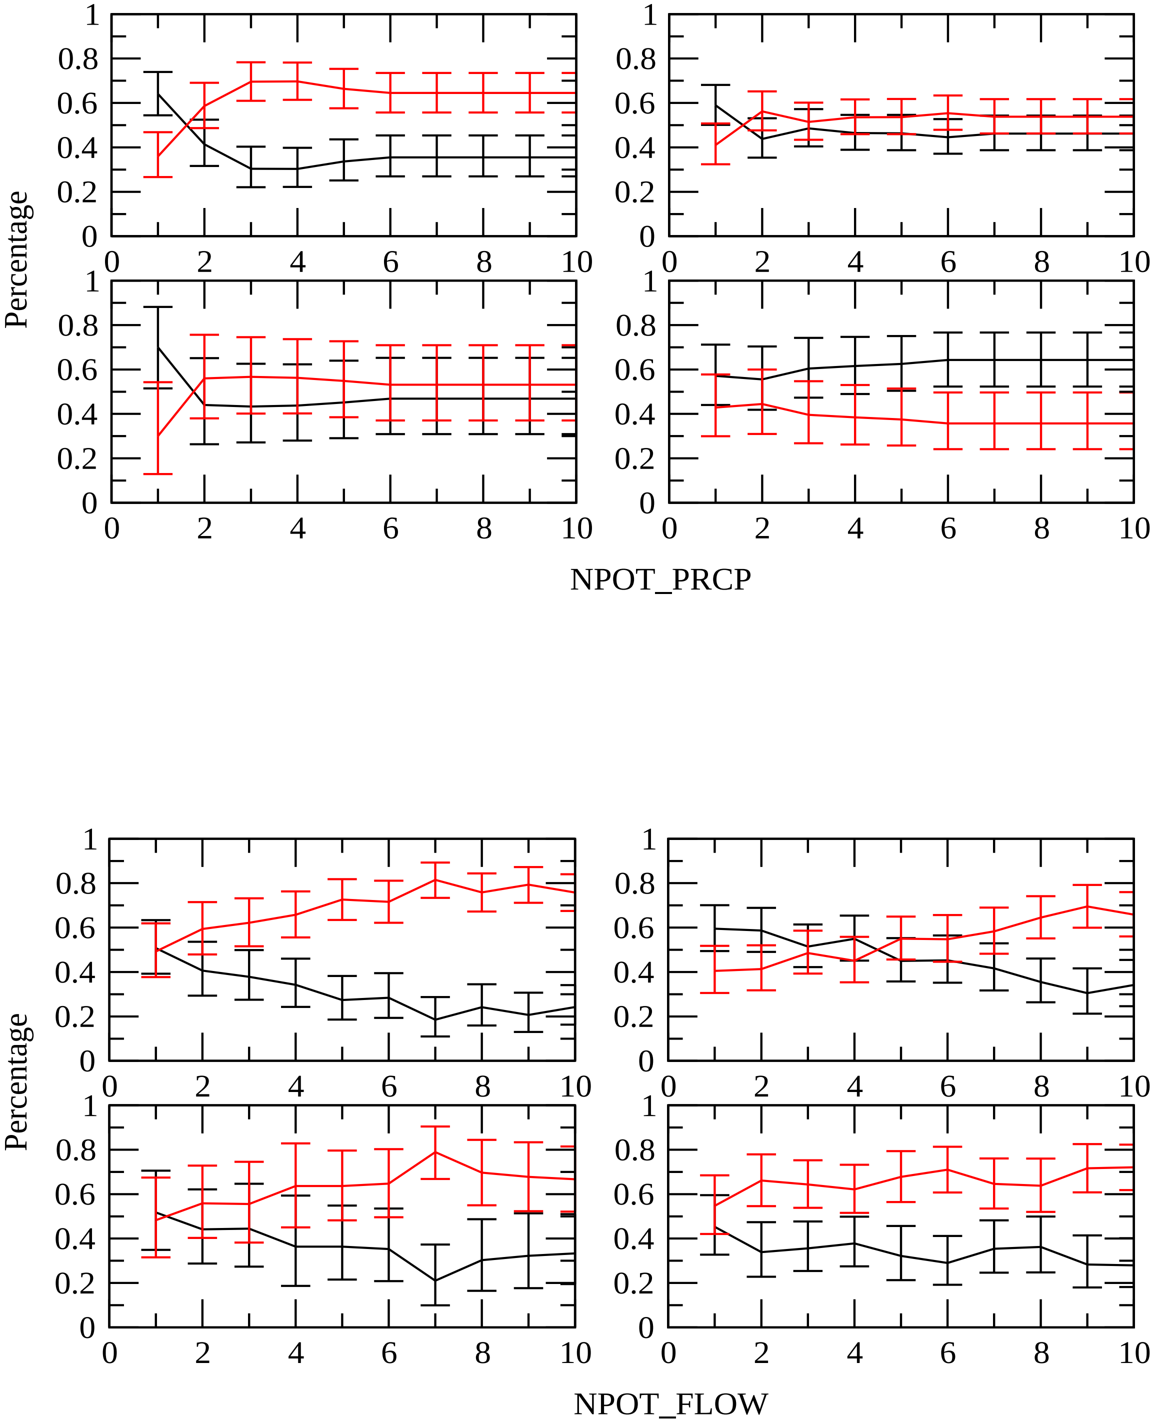

Figure 8. Percentage of precipitation (a-d) stations with downtrends (black) and uptrends (red), and percentage of streamflow (e-h) stations with downtrends (black) and uptrends (red). Record length = 60yr.

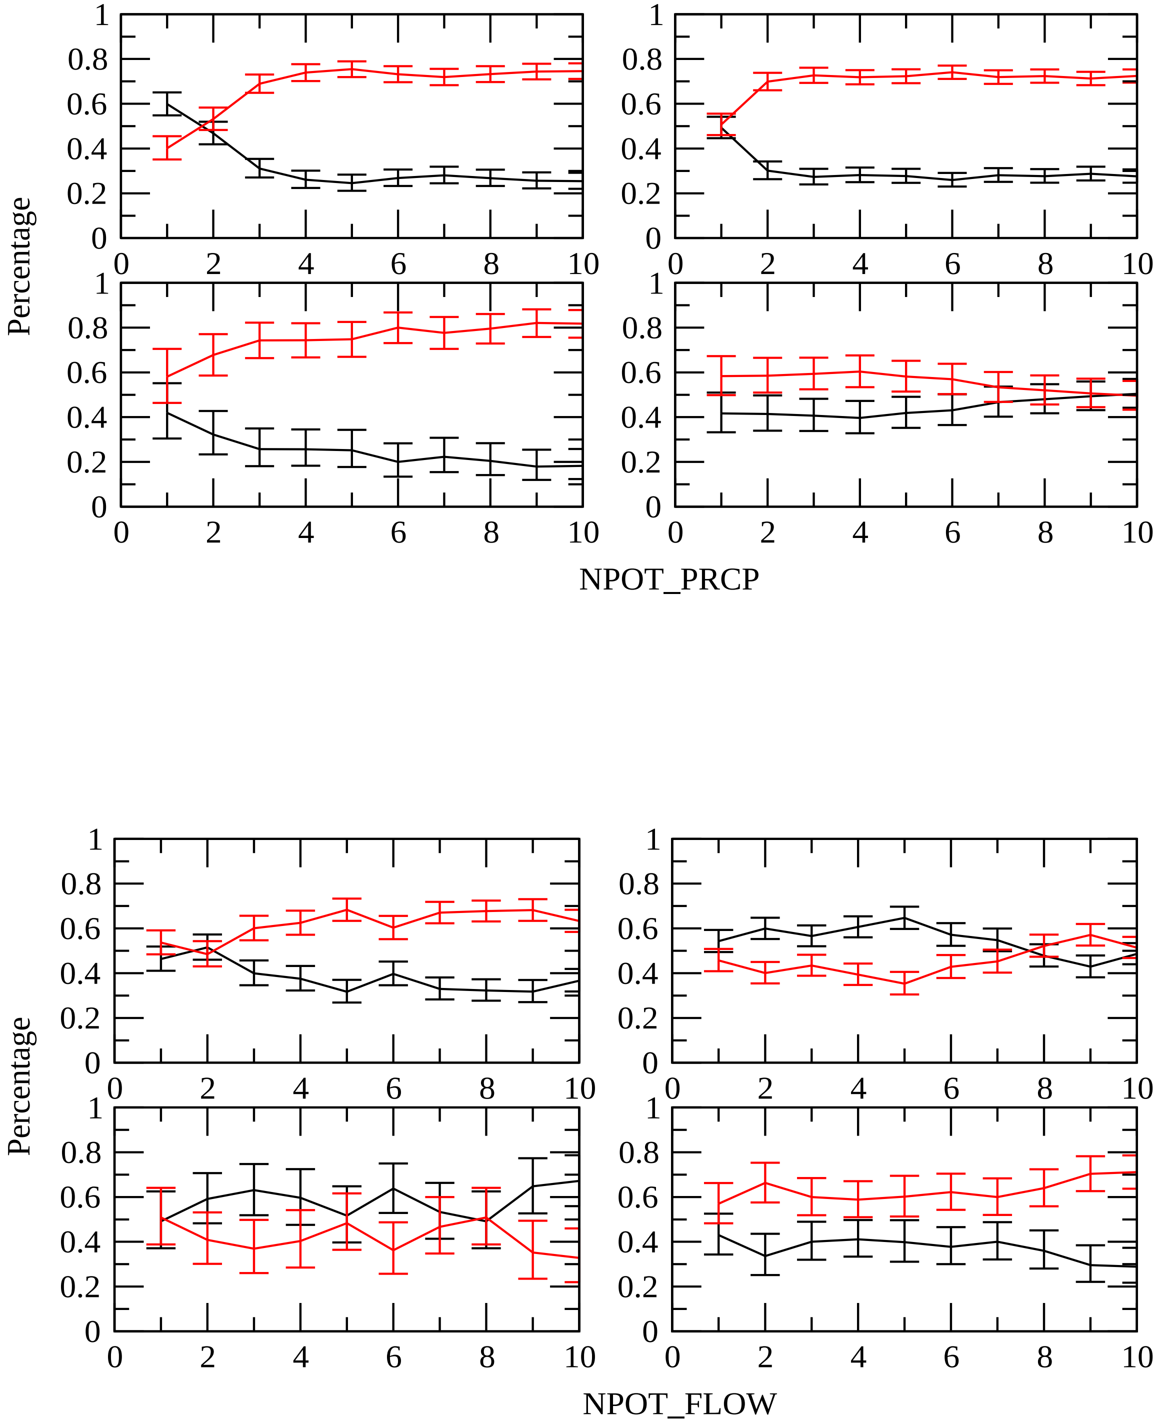

Figure 8. Percentage of precipitation (a-d) stations with downtrends (black) and uptrends (red), and percentage of streamflow (e-h) stations with downtrends (black) and uptrends (red), for paired events. Record length = 60yr.

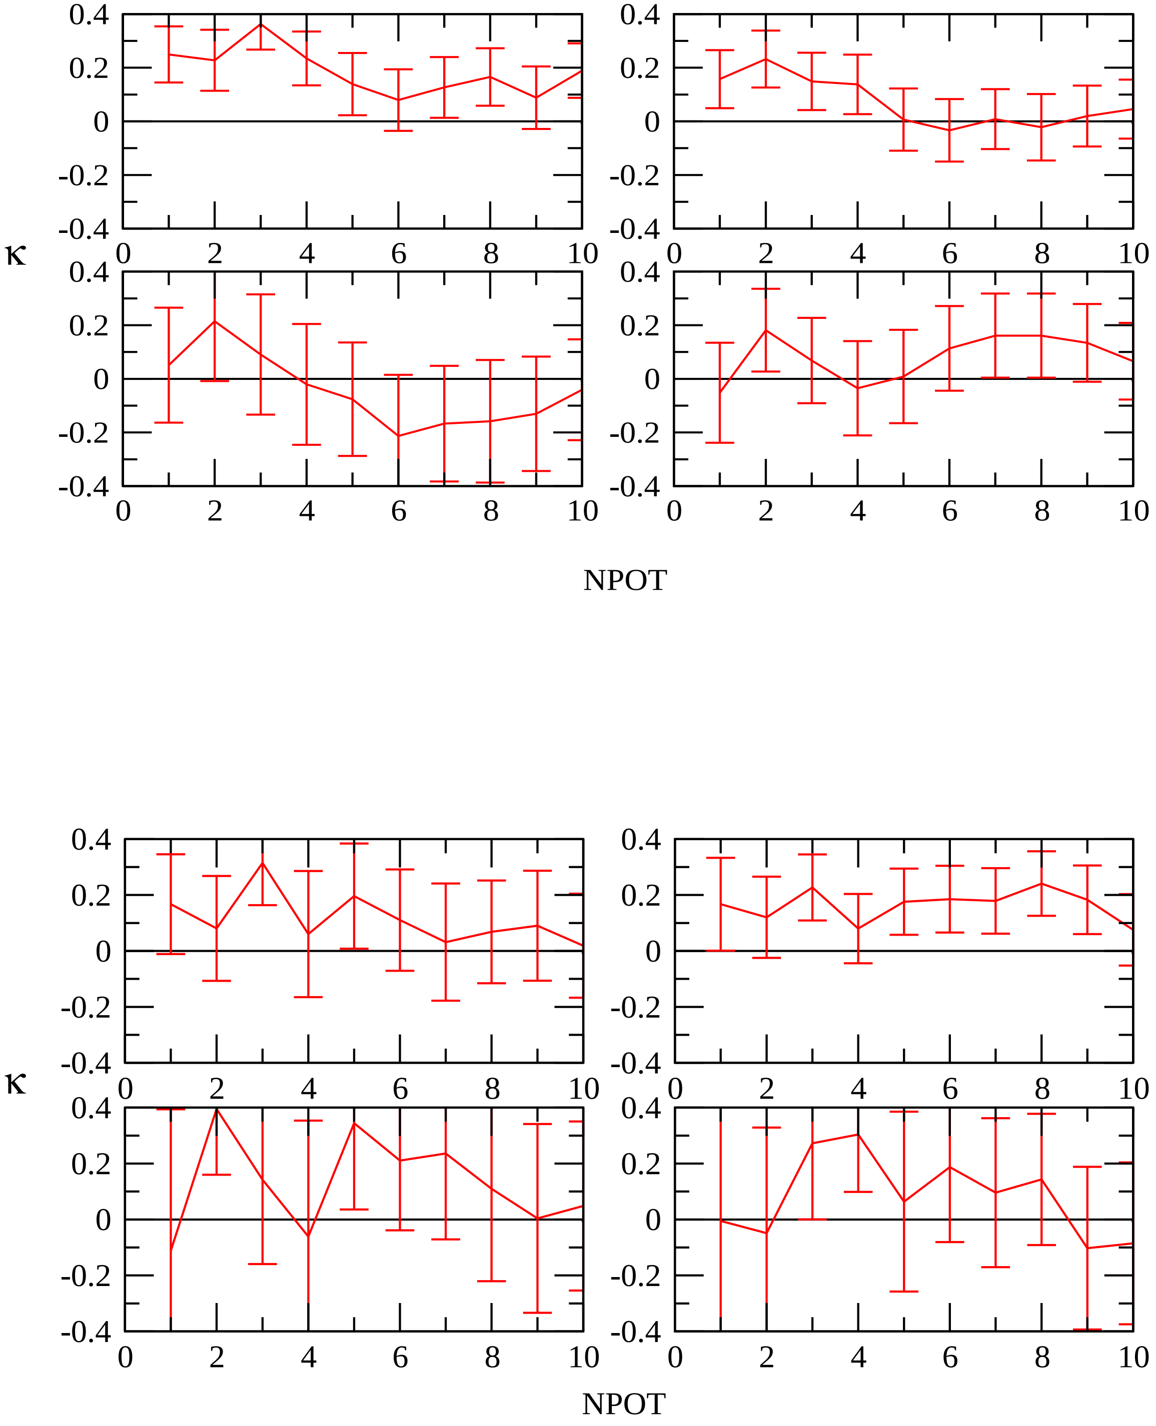

Figure 10. The measure of association (κ) between precipitation and flow trend. The top (bottom) block of figures is based on the 30yr (60yr) record. Left (right) column is for Spring (Summer). Within each block the top (bottom) row is for East (West).

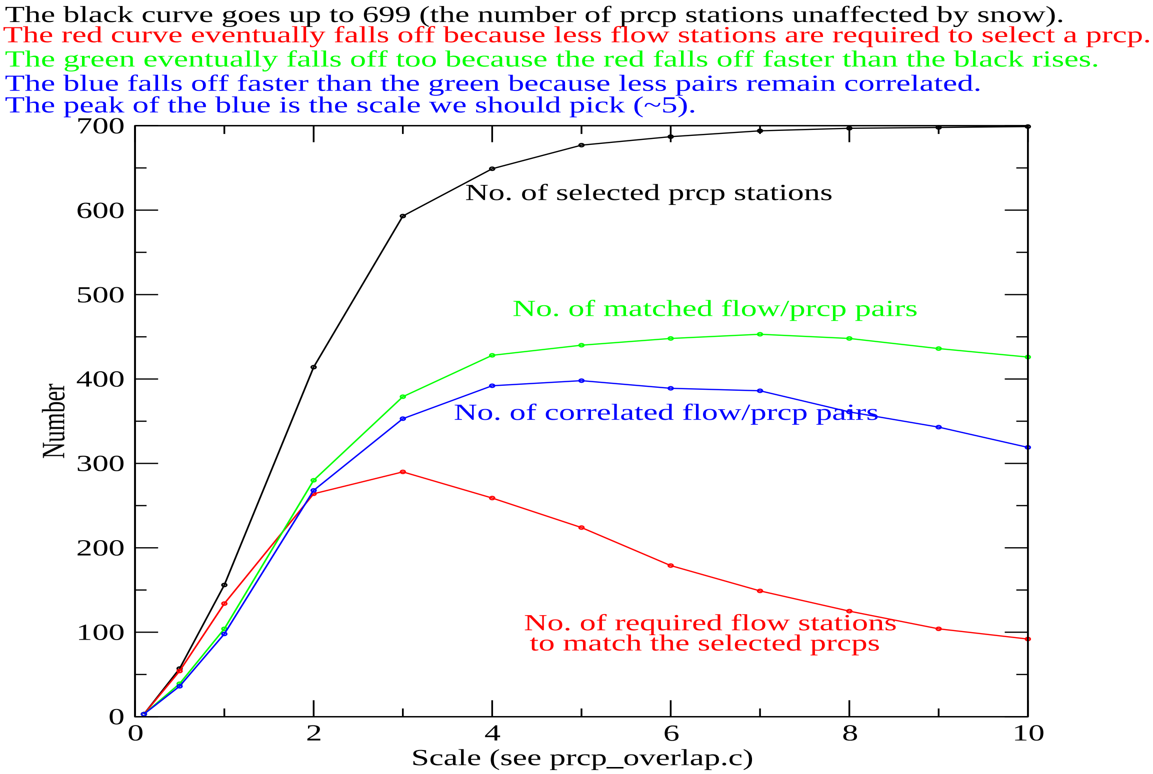
The black curve goes up to 699 (the number of prcp stations unaffected by snow).
The red curve eventually falls off because less flow stations are required to select a prcp.
The green eventually falls off too because the red falls off faster than the black rises.
The blue falls off faster than the green because less pairs remain correlated.
The peak of the blue is the scale we should pick (~5).

Fig. x: The number of precipitation and flow stations before and after pairing. (Unnecesary??).

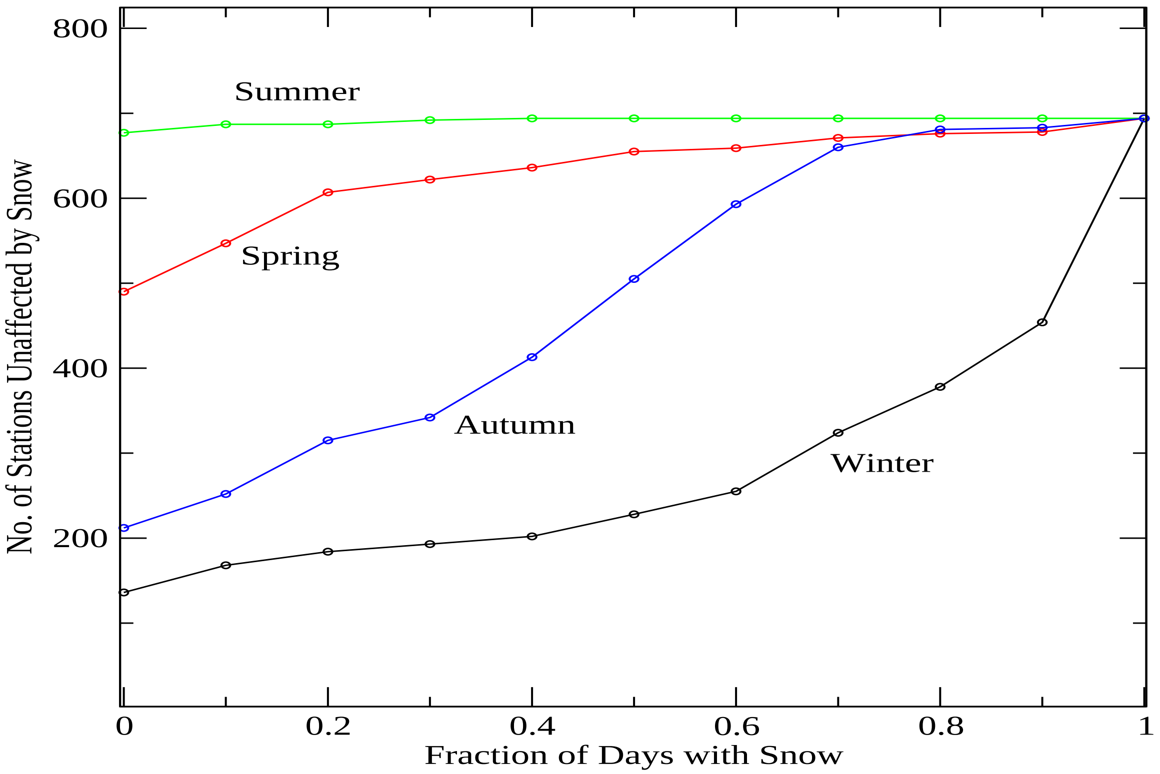

Fig. x: The number of precipitation stations unaffected by snow as a function of the fraction of days with snow. (Unnecesary??).

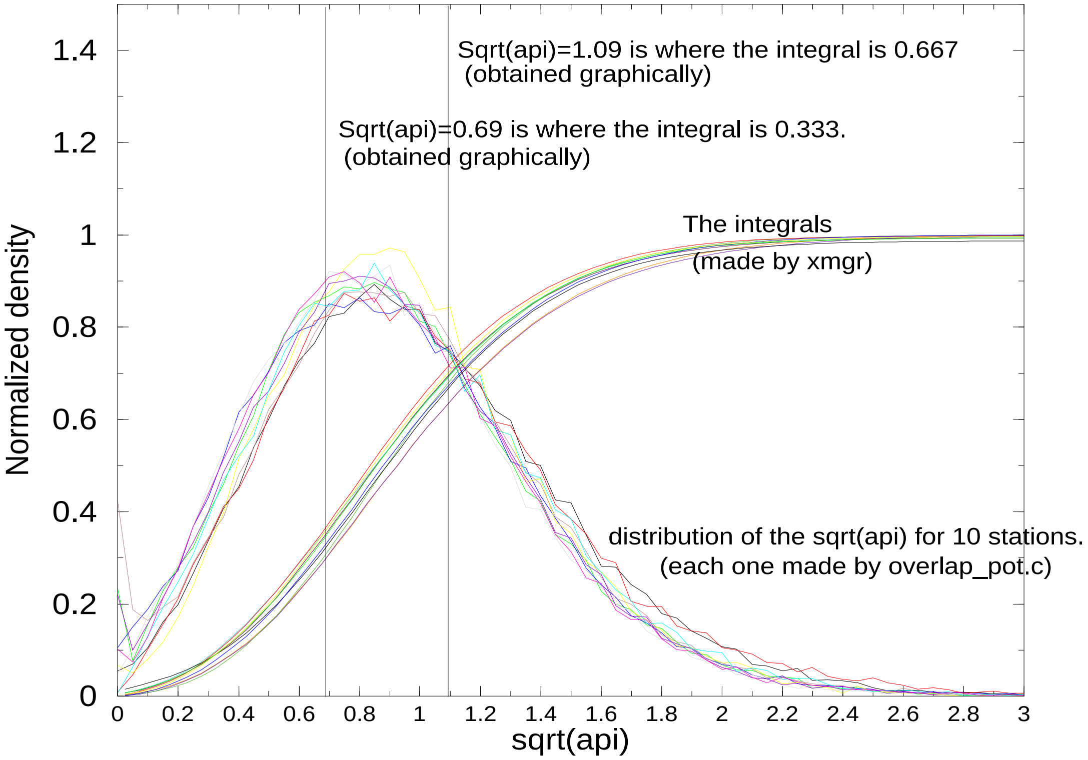

Fig. x. The density and the cumulative distribution of API. Unnecessary??