Assessing causality from multivariate time series

P. F. Verdes
Institute for Environmental Physics, University of Heidelberg, Im Neuenheimer Feld 229, D-69120 Heidelberg, Germany
(Received 16 February 2005; revised manuscript received 23 June 2005; published 25 August 2005)

In this work we propose a general nonparametric test of causality for weakly dependent time series. More precisely, we study the problem of attribution, i.e., the proper comparison of the relative influence that two or more external dynamics trigger on a given system of interest. We illustrate the possible applications of the proposed methodology in very different fields like physiology and climate science.

DOI: 10.1103/PhysRevE.72.026222 PACS number(s): 05.45.Tp, 05.45.Xt, 89.70.+c, 87.80.Tq

I. INTRODUCTION

Detecting interdependencies and causal relationships is one of the most relevant problems among those addressed by time series analysis. As more and better data become available along with technological progress, this issue is receiving increasing attention in the literature [1–14]. Applications are ubiquitous in almost all fields of science, like physics, economics, and ecology, to name a few, or physiology—most notably in brain studies. In this last case, for example, information about the interaction among recorded channels of an electroencephalogram can aid clinical practice by pinning out the region of the brain that is acting as a recruiting focus in epilepsy.

Commonly used tools for the estimation of dependencies are linear cross-correlation, cross-spectra [16], and mutual information [17]. However, these measures share the property of being symmetric and, as such, are not suited for assessing causality within relationships. To study the directional aspect of interactions, many approaches have been employed. One of them consists in examining whether the prediction of one series can be improved by incorporating information from the other [1–3]. This was originally proposed by Wiener [18] and later formalized by Granger [19] in the context of linear regression models of stochastic processes. Specifically, if the variance of the prediction error of a given time series at the present time is reduced by the inclusion of past measurements from a second time series in a linear regression model, then the latter is said to have a causal influence on the former. Of course, the roles can be inverted to address the question of causal influence in the opposite direction. Since Granger causality was formulated for linear models, its direct application to nonlinear systems may not be appropriate. Granger's idea has thus been extended to the nonlinear case either (i) by restricting its application to local linear models in reduced neighborhoods and then averaging the resulting statistical quantity over the entire dataset [13], or (ii) by considering the error reduction triggered by added variables in global nonlinear models [1]. Despite the success of these strategies, it must be noticed that model-based methods (linear or not) may suffer from the shortcomings of a model misspecification. In order to avoid this problem, we have used in this work a general *nonparametric* test of causality for stationary and weakly dependent time series based on Information Theory [9,11,12]. More precisely, the proposed measure builds upon conditional entropy, which, in contrast to the standard form of mutual information [20,21] or delayed-related measures, allows us to distinguish actually transported from shared information. Furthermore, since correlation integral-based entropies need minimal assumptions about the underlying dynamics of the systems and the nature of their coupling, their application does not assume the existence of any deterministic process [14,15]. Specifically, our purpose is to compare the relative influence that two or more external dynamics exert on a given system of interest. To this end, we start by considering a construction by Schreiber [14] and derive a new information-theoretic attribution technique, as described in the next section.

II. THEORETICAL ASPECTS

We first need to briefly recall the concept of *transfer entropy* (see Ref. [14] for details). Assume that we observe three physical processes X, Y, and Z, visiting states x_i, y_i, and z_i with probabilities $p(x_i)$, $p(y_i)$, and $p(z_i)$, respectively. We think of X, Y, and Z as interacting subsystems of a whole, x_i, y_i, and z_i being, in principle, multivariate descriptors of their respective internal states [22]. We are interested in comparing the relative influence that Y and Z have on X. The transfer entropy from Y to X is defined by

$$T_{Y\to X} = \sum_i p(x_i,y_i,z_i)\log \frac{p(x_i|y_i,z_i)}{p(x_i|z_i)}, \tag{1}$$

where the sum extends over all states i visited by the composite system. Its most appealing feature is that—unlike mutual information or conditional entropy—it is essentially asymmetric in X and Y, detecting the directional transport of information. By virtue of the conditioning to the state of Z under the logarithm, it ignores all static correlations between X and Y that could be possibly triggered by their common interaction with Z. (The generalization to the case of more subsystems is straightforward.) For continuous systems, in time series applications, the transfer entropy is obtained by coarse graining the state space at resolution ε, thus making $T_{Y\to X}$ dependent on the size of the partitioning. The transition probabilities involved in (1) read as $p(x_i|y_i,z_i) = p(x_i,y_i,z_i)/p(y_i,z_i)$, and the latter can be approximated by the visitation frequency

$$p(x_i,y_i,z_i)=\frac{1}{N_{pairs}}n(\Delta x_{ij}<\varepsilon,\Delta y_{ij}<\varepsilon,\Delta z_{ij}<\varepsilon),$$

where $\Delta x_{ij}=|x_i-x_j|$ (analogous definitions hold for Δy_{ij} and Δz_{ij}), N_{pairs} denotes the total number of ij pairs available, and $n(\cdot)$ is the number of couples satisfying the corresponding distance constraints.

Defined in this way, transfer entropy is best suited for the study of homogeneous spatially extended systems, like identical maps coupled on a lattice. In contrast, as noted in [14], when the processes are of a different nature, the driven system will most probably respond with a different sensitivity to each forcing agent, and the transfer entropy (1) will depend on the overall dynamic range of the signals. To overcome this difficulty, and inspired by the work of Pi and Peterson [23], we propose to redefine probabilities by distorting cubic neighborhoods in state space:

$$p^*(x_i,y_i,z_i)=\frac{1}{N_{pairs}}n(\Delta x_{ij}<\varepsilon,\Delta y_{ij}<\delta_Y,\Delta z_{ij}<\delta_Z).$$

Now the question naturally arises as how to choose appropriate values for δ_Y and δ_Z. To answer it, e.g., in the case of Y, we consider the conditional probability

$$p^*(x_i|y_i)=\frac{n(\Delta x_{ij}<\varepsilon,\Delta y_{ij}<\delta_Y)}{n(\Delta y_{ij}<\delta_Y)}.$$

If for any fixed ε (for instance $\varepsilon=\sigma_x$) we think of $p^*(x_i|y_i)$ as a function of δ_Y only, then a possible choice is given by

$$\delta_Y=\arg\max p^*(x_i|y_i).$$

The rationale behind this maximization follows. How does the conditional probability $p^*(x_i|y_i)$ behave as a function of δ_Y? For $\delta_Y\to\infty$ the conditioning has no effect and hence $p^*(x_i|y_i)=p(x_i)$. As $\delta_Y\to 0$, $p^*(x_i|y_i)$ either remains flat (if X is independent of Y) or grows monotonically until saturation to a certain value smaller than 1. For vanishing δ_Y, the computation of $p^*(x_i|y_i)$ becomes spoiled by the finite sample size and deflects down to zero [24]. Thus, the interesting features of the information transport from Y to X are to be found neither at big nor small values of δ_Y, but at some appropriate scale in between. It is this relevant scale that will be picked up by the maximization process, something previously forbidden by the rigidity of cubic neighborhoods. This procedure makes now irrelevant the assignment of a particular overall scale to Y. Along the same lines, we choose $\delta_Z=\arg\max p^*(x_i|z_i)$.

Finally, we add a normalization factor to build an intensive measure and define the *information transfer* from Y to X as

$$\mathrm{IT}_{Y\to X}=\frac{1}{N}\sum_i p^*(x_i,y_i,z_i)\log\frac{p^*(x_i|y_i,z_i)}{p^*(x_i|z_i)}, \tag{2}$$

where N is the length of the available database. Now, an important issue is how to assess the significance of the information transfer $\mathrm{IT}_{Y\to X}$ detected by (2) at each scale ε. To this end, after the composite state vectors (x_i,y_i,z_i) have been built, we generate randomized controls by time shuffling $\{y_i\}$ while leaving the setting otherwise identical [25]. However, we observe that this scrambling will whiten the power spectrum of Y. (For simplicity, let us briefly assume that all subsystems are unidimensional.) Its properties will change completely: the random shuffle will be wildly irregular even if the Y data represent a very smooth signal. Thus, seeking a sound basis for comparison, our first step will be to time shuffle the complete original embedding vectors (x_i,y_i,z_i), yielding all spectra initially white.

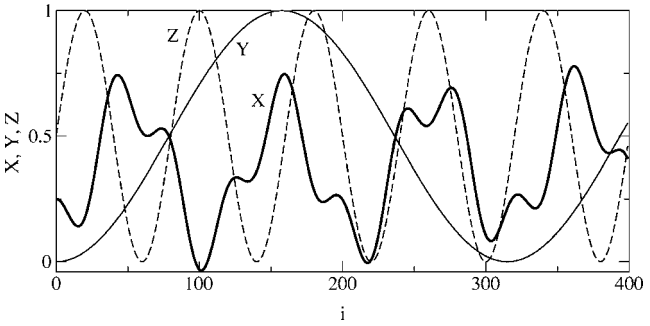

FIG. 1. Worked artificial example: the component time series in the noise-free case, before the initial shuffle described in the main text.

III. A WORKED EXAMPLE

To illustrate the advantages and limitations of this methodology, we first consider a simple example involving unidimensional subsystems. We generate a set of samples (x_i,y_i,z_i), $i=1,\ldots N$, according to

$$x_i=y_i(18y_i^2-27y_i+10)/2+z_i(1-z_i),$$

$$y_i=[1-\cos(2\pi i/315)]/2,$$

$$z_i=[1+\sin(2\pi i/80)]/2, \tag{3}$$

depicted in Fig. 1. As expected by construction, X exhibits both large- and small-scale variability—the Y contribution to X was chosen bigger than that of Z. The existing causal conditionings among these variables seem elusive to the naked eye, probably hidden by their highly nonlinear character. For example, when Z increases, X may equally respond by increasing or decreasing, and the same holds true when Z goes down; for Y we encounter a similar situation. This is readily reflected by vanishing standard cross-correlation coefficients: $R_{XY}=-0.09$ and $R_{XZ}=-0.08$. We have also added Gaussian noise to the time series of all subsystems with an increasing standard deviation ratio $\sigma_{noise}/\sigma_C=0$ to 1, where the symbol C denotes the X, Y, and Z components.

As described in the previous section, we start by shuffling in time the states of the composite system (x_i,y_i,z_i), so that all spectra are initially white. We compute $\mathrm{IT}_{Y\to X}$ and further scramble Y to assess the significance of this result; the same procedure is also applied to Z. The information transfers $\mathrm{IT}_{Y\to X}$ and $\mathrm{IT}_{Z\to X}$ obtained in this way for $N=100$ and 1000 are plotted with thick lines in Fig. 2. Consider the noise-free case first, depicted in the left panels. Not surprisingly, detec-

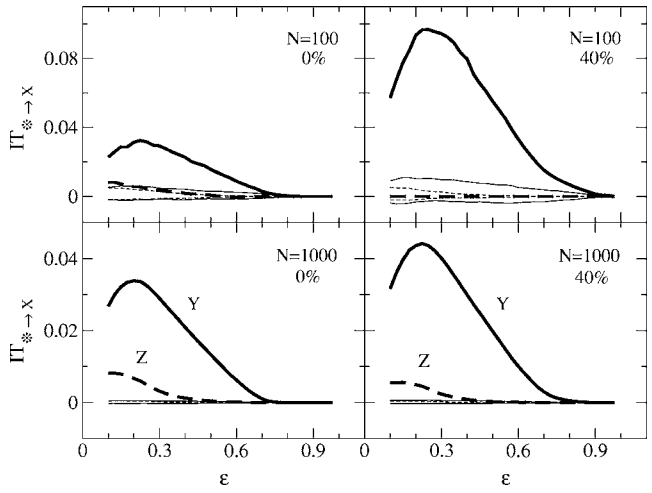

FIG. 2. $IT_{Y\to X}$ and $IT_{Z\to X}$ (thick lines) as a function of length scale ε for the artificial example. Full lines are everywhere associated to Y, and dashed to Z. Thin lines embrace two standard deviations of the distribution of IT over a control population of size 1000. Upper and lower panels correspond to database sizes $N=100$ and $N=1000$, respectively. Increasing noise levels from left to right.

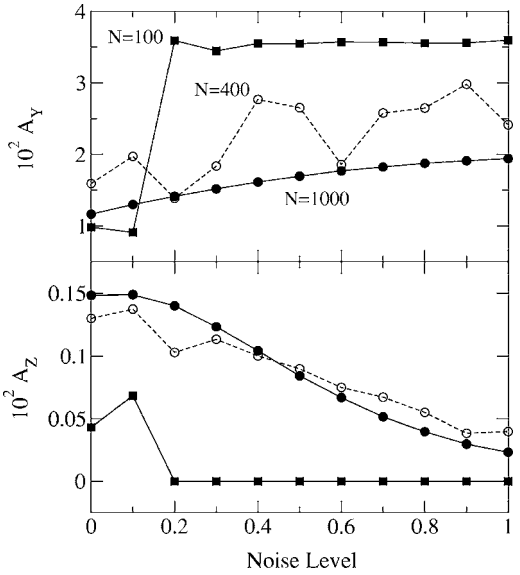

FIG. 3. Area in excess of the two-sigma reference level as a function of noise intensity σ_{noise}/σ_C for $N=100$ (full squares), $N=400$ (open circles, dashed line), and $N=1000$ (dots, full line).

tions become sharper when more observational data are available. In particular, we notice that with less uncertainty IT consistently vanishes for all ε over the randomized control populations. For small datasets and neighborhood sizes our measure seems to be biased toward overestimation. Should we be interested in a precise absolute computation of IT, we would need to implement a correction for this effect. However, since such an offset is irrelevant to our only purpose of comparison, we will ignore it in the following. We further notice the predominance of Y over Z at all scales, and also how $IT_{Y\to X}$ is significant for scales where $IT_{Z\to X}$ is not ($0.4<\varepsilon<0.7$, $N=1000$), which is interpreted as Y being responsible for a larger-scale variability of X than Z.

An examination of the noisy cases plotted in the right panels of Fig. 2 reveals a surprising increase in $IT_{Y\to X}$, in particular, for $N=100$—but also for $N=1000$. This result contradicts, in principle, the natural expectation that the presence of contaminating noise would attenuate the detecting power of the proposed statistics. To investigate this effect closer, we construct an informal measure A_Y (A_Z) as the area under $IT_{Y\to X}$ ($IT_{Z\to X}$) that lies above the corresponding 2σ control level plotted with thin lines in Fig. 2. We show in Fig. 3 the behavior of A_Y and A_Z as a function of noise intensity for $N=100$, 400, and 1000. We observe that only A_Z indeed follows our intuition, i.e., growing noise levels increasingly hide the existing dependency of X on Z. Looking back at Eq. (2), we see that a faster decreasing trend in $p^*(x_i|z_i)$ as compared to $p^*(x_i|y_i,z_i)$ is responsible for the observed increase in $IT_{Y\to X}$. Stated in other words, as the less influential input Z goes buried first into noise (and this effect is more noticeable with fewer data), the importance of Y is strengthened by the relative character of our measure. As we can see in Fig. 3 for $N=100$, a noise-level $\sigma_{noise}/\sigma_C=0.2$ is strong enough to spoil the detection of any information transfer from Z to X, and this transition yields a dramatic increase in A_Y.

A final remark: We have checked that these results remain unaffected by arbitrary rescalings of Y and Z, and that only the corresponding transport scales δ_Y and δ_Z change accordingly.

IV. APPLICATIONS

A. Cardiorespiratory interaction

We now consider Dataset B of the Santa Fe Time Series Competition [26]. This is a multivariate dataset recorded from a patient in the Sleep Laboratory of the Beth Israel Hospital in Boston, MA. The recorded magnitudes are heart rate (H), chest volume or respiration force (R) and blood oxygen concentration (O), and are plotted in Fig. 4.

Under normal, physiological conditions, the heart rate is modulated by respiration through a process known as Respiratory Sinus Arrhythmia (RSA). It is the natural cycle of arrhythmia that occurs through the influence of breathing on the flow of sympathetic and vagus impulses to the sinoatrial node of the heart. More precisely, its rhythm is primarily under the control of the vagus nerve, which inhibits the heart rate and the force of contraction. When we inhale, vagus nerve activity is impeded and the heart rate begins to increase. When we exhale, this pattern is reversed. This quasiperiodic modulation of heart rate by respiration is most notable in young, healthy subjects [27] and decreases with age. The degree of fluctuation in the heart rate is also controlled significantly by regular impulses from the baroreceptors (sensors of partial pressure of oxygen and carbon dioxide) in the aorta and carotid arteries.

However, Dataset B corresponds to a patient suffering from sleep apnea, which is a breathing disorder characterized by brief interruptions of breathing during sleep. There are two types of sleep apnea: central and obstructive. Central sleep apnea, which is less common, occurs when the brain

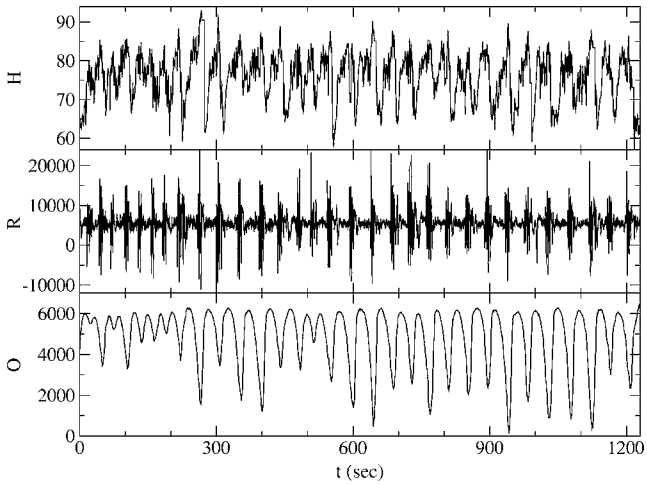

FIG. 4. Simultaneously measured time series of instantaneous heart rate (upper panel), respiration force (middle), and blood oxygen concentration (lower panel) of a sleep apnea patient. The data are given in uncalibrated A/D bits with a sampling frequency of 2 Hz.

fails to send the appropriate signals to the breathing muscles to initiate respirations. Obstructive sleep apnea is far more common and occurs when air cannot flow into or out of the person's nose or mouth although efforts to breathe continue. During apneic events, the person is unable to breathe in oxygen and to exhale carbon dioxide, resulting in low levels of oxygen and increased levels of carbon dioxide in the blood, which alert the brain to resume breathing and cause an arousal. With each arousal, a signal is sent from the brain to the upper airway muscles to open the airway; breathing is resumed, often with a loud snort or gasp [28].

As described above, sleep apnea affects the normal process of RSA, disturbing the usual patterns of interaction and feedback among the heart rate, respiration, and blood oxygen concentration. As a result, the control of the heart rate by respiration becomes unclear. It may well be blocked, in accordance with the change in dynamics that is characteristic of the so-called "dynamical diseases." Furthermore, in some studies it has been claimed that a weak coupling in the reversed direction could be observed [14,29,30]. However, these investigations were based on a bivariate analysis of only the heart and respiration time series, disregarding feedback interplays with blood oxygen concentration. Here we raise the following question: How much information is transferred from R to H and vice versa, *excluding* possibly shared dynamics triggered by their common interaction with O?

To answer this question, we begin by scrambling in time the states of the composite system (H_i, R_i, O_i), so that all spectra are initially white. We compute $\mathrm{IT}_{R\to H}$ and shuffle R to assess the significance of this result; the same procedure is also applied to O. The information transfers $\mathrm{IT}_{R\to H}$ and $\mathrm{IT}_{O\to H}$ obtained in this way are plotted with thick lines in Fig. 5. As we can see in this figure, the influence of R on H lies within the 2σ limit of the control population. This means that R yields no significant information on H in excess of that already furnished by O. In contrast, we find a significant flow of information from O to H beyond their common interaction with R.

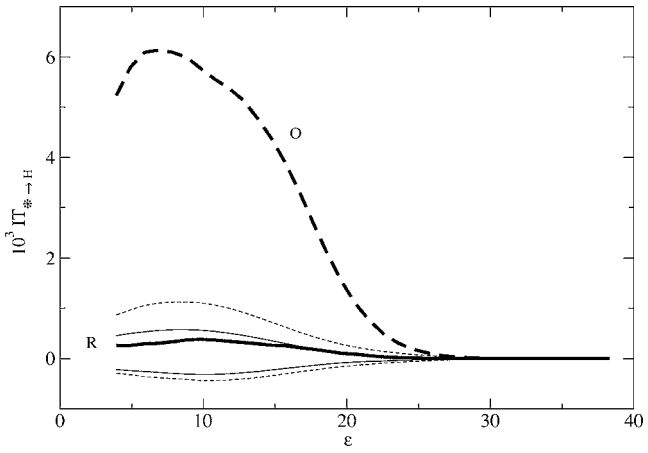

FIG. 5. Information transfer to the heart rate time series coming from respiration (full line) and blood oxygen concentration (dashed line). Thin lines indicate two standard deviations of the results over the control population.

In Fig. 6 we compare the influence that heart rate and blood oxygen concentration exert on the respiration force. Much like the previous case, the flow from H to R falls within the values expected for the control population and therefore cannot be considered significant. Instead, we observe that blood oxygen concentration is sensibly affecting the dynamics of respiration, a result that is in nice agreement with the known mechanism of control by baroreceptors described above.

Finally, in Fig. 7 we compare for completeness the influence that heart rate and respiration exert on blood oxygen concentration. Not surprisingly, we find that both dynamics significantly affect the level of blood oxygen. Together with our previous results, they constitute a reflection of complex two-way interactions between H and O, and R and O. Overall, respiration seems to be a more decisive factor in determining blood oxygen levels than is the heart rate.

Now, as Fig. 4 shows, the raw series cannot be assumed stationary, a property that is, however, essential for the validity of the present analysis. One of the two standard solu-

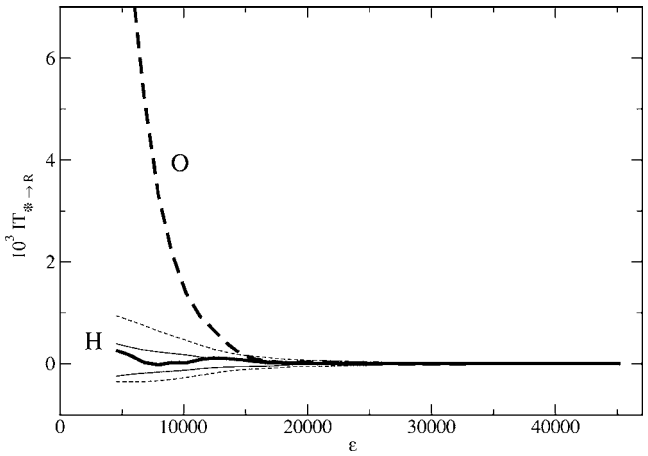

FIG. 6. Information transfer to the respiration time series coming from heart rate (full line) and blood oxygen concentration (dashed line).

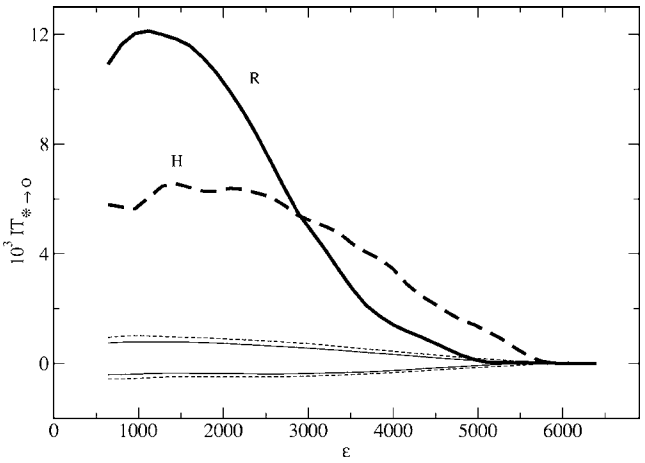

FIG. 7. Information transfer to the blood oxygen concentration time series coming from heart rate (dashed line) and respiration (full line).

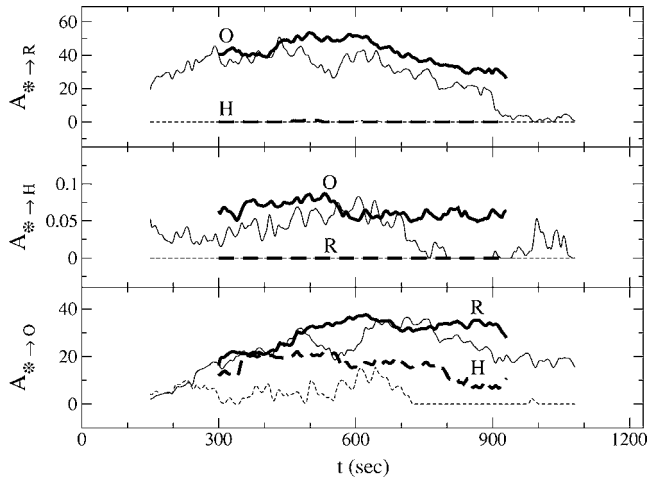

FIG. 8. Temporal evolution of detected causalities in sliding windows, as measured by the area in excess of the 2σ reference level versus the interval middle-point position. Thick lines correspond to a window size $L=N/2$; thin lines of the same type are employed for the case $L=N/4$.

tions to this problem is to define some new variable that can be regarded as sufficiently stationary. In financial time series applications, e.g., this is usually accomplished by considering either the *increments* $\delta x_\tau(t)=x(t+\tau)-x(t)$, *returns* $\delta x_\tau(t)/x(t)$, or *log returns* time series $\ln[x(t+\tau)]-\ln[x(t)]$. However, this procedure is only optimal to cope with the so-called weak nonstationarities—drifting mean and variance. In the general case of a single long time series generated by a dynamics slowly varying in time, the nonstationarity can be approximately mapped to a stationary situation by cutting the time series into (potentially overlapping) intervals, if the change in the dynamics is sufficiently slow. In a strict sense, the resulting segments will still be slightly nonstationary. Hence, time averages will no longer represent averages according to an invariant distribution, which in turn prevents the interpretation of relative frequencies as true probabilities. This loss of the ergodicity condition constitutes the main conceptual drawback of nonstationarity for the whole time series analysis branch. Nevertheless, the accepted view in the literature is that—by dropping their corresponding physical interpretations and considering the measures as having only a *relative* character—useful conclusions can still be drawn by comparison, both with surrogate populations as well as with neighboring windows. It is in this spirit that many concepts derived from the theory of nonlinear time series analysis, like Lyapunov exponents or the correlation dimension, have been applied in sliding windows for the study of EEG or other records of physiological activity [31–34]. In our case the nonstationarity should not necessarily be considered disadvantageous, but can instead be used to identify periods of weaker and stronger coupling between the subsystems. We have thus estimated causalities as a function of time—a procedure that naturally raises a tradeoff between stationarity and the window length L considered. On one hand, the stationarity condition is better fulfilled for decreasing L. On the other hand, statistical significance testing against control populations may yield negative results in this limit due only to an insufficient amount of data. The competition between these two factors was already illustrated for the worked example of the previous section, where we showed their interplay for the system (3) for varying window lengths and noise levels. In the present case, we have looked for converging evidence by probing different window lengths. In order to summarize the new information, we plot in Fig. 8 the area under the IT vs ε curves that lie above the corresponding 2σ control levels for all possible causal combinations among respiration, heart rate, and blood oxygen concentration. We first employed almost-overlapping windows of size $L=N/2=1230$ data points, corresponding to a temporal span of roughly 600 s (see Fig. 4), with consecutive windows differing only in a single data point. These results are depicted with thick lines in all panels of Fig. 8. We have also performed a finer scanning, sliding a half-sized window of length $L=N/4=615$ data points (approx. 300 s), and plot the (more fluctuating) obtained causalities with thin lines. The evolving relative influence that heart rate and blood oxygen concentration exert on the respiration rhythm can be compared in the upper panel. A null information transfer from H to R is confirmed for all epochs of the window sizes considered, whereas the oxygen influence slightly increases during the first half of the record and steadily decreases in the second half. Furthermore, $A_{O\to R}$ eventually vanishes for $L=N/4$ (beyond $t=900$ s). Although it cannot be completely ruled out that this may be a reflection of a small window size collecting an insufficient amount of data, the fact that detection was possible for earlier epochs together with an overall evolution of $A_{O\to R}$ consistent with $L=N/2$ (thick solid line) suggest that the information flow is indeed vanishing for the last epochs of the record. The case corresponding to the heart rate is depicted in the central panel, and it shows a general behavior rather similar to that already encountered for the respiration time series in the upper panel. Finally, in the lower panel we plot the temporal evolution of the estimated influences that R and H exert on O. In spite of the fluctuating character of these curves, both window sizes confirm the contention that respiration is a more decisive factor in determining blood oxygen levels than is heart rate.

In conclusion, the analysis in this section has shown that a bivariate analysis of heart rate and respiration is insufficient, it being essential that their interaction with the baroreceptor system of control be also encompassed. This fact was previously suggested in other studies in the literature [14,30]. In accordance with the intrinsic change in dynamics that is characteristic of the so called dynamical diseases, respiratory sinus arrhythmia seems to be blocked in the pathological case of sleep apnea. Of course, for this conclusion to be valid in general, systematic tests should be performed on a more extensive database, a task that is beyond the scope of this work.

B. Climate change

To illustrate the possibilities of application of this tool in very different areas, in this subsection we make a preliminary study of the problem of climate change. It is well known that the global climate has changed rapidly over recent decades, as reflected by evolving patterns of tropical circulation, monsoon rainfall, and other climatological parameters [35]. In particular, global temperatures have increased by (0.6 ± 0.2) °C since 1860 [36]. A question that has become increasingly crucial is to determine the relative responsibility of several potentially explaining factors: the emission of heat-trapping greenhouse gases, the Sun's increased radiative output, the lack of volcanic activity, the changes in atmospheric ozone, etc. [37]. Among these, the unambiguous characterization of forcing by greenhouse gases (GHG) as opposed to solar total irradiance (STI) is particularly important due to their anthropogenic origin.

On one hand, there is a well-established consensus on warming by increased GHG, supported by extensive simulations with atmosphere-ocean general circulation models (AOGCMs) for the past centuries [38,39]. On the other hand, a number of paleoclimatic reconstructions of global temperatures show a conspicuous correlation to the secular behavior of the Sun's irradiance [40], but the absolute variations in solar intensity are small and the physical mechanisms for the required amplification have not been identified. Attempts in this direction are given, e.g., by recent experiments with the GISS (Goddard Institute for Space Studies) stratospheric general circulation model that show that solar variability affects surface winds, sea-level pressures, and regional surface temperatures [41]. However, the conclusions drawn from these simulations will be continuously revised insofar as GCM models are refined and their performances improved. This constitutes a formidable task, for atmospheric dynamics is extremely complex and not fully understood: among other delicate issues, its models must account for couplings with the land, oceans, and associated nonlinear feedbacks, together with their chemical responses to increased radiation at all wavelengths [42]. It is generally accepted that AOGCMs are at least quantitatively uncertain with regard to the regional patterns of the behavior of climate, the hydrological cycle, or sea ice drift (see Ref. [43] for a review). One of the remaining inconsistencies is that they cannot fully account for the observed difference in the trend between superficial and lower tropospheric temperatures over the last 20 years, even when all known external influences are included [35].

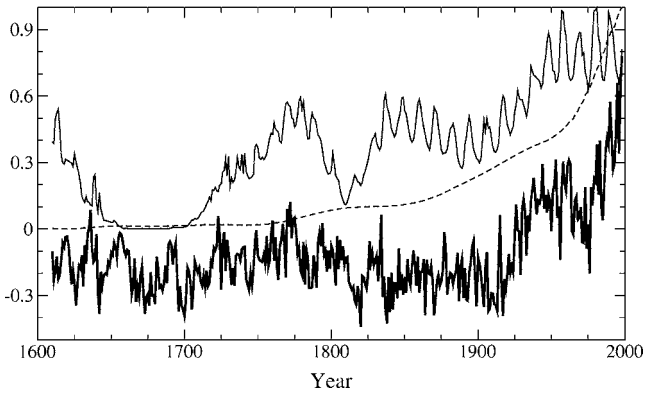

FIG. 9. Available registers for the problem of climate change. Temperature anomalies are represented with a full thick line; STI (full thin line) and GHG (dashed line) have been rescaled to the unit interval only for ease in plotting.

In view of these caveats, a model-free statistical method is attractive. As an advantage, this approach offers the possibility of directly testing hypotheses on the observational data. The standard practice along this line is to employ "optimal detection techniques," which are essentially multiple-regression frameworks that assume that the observations of interest can be explained as a linear combination of exogenous candidate signals plus noise [45]. The procedure consists of estimating the unknown coefficients and testing the null hypothesis that they vanish. More involved techniques developed from information theory have already been applied to some climatic problems [46] but, to the best of our knowledge, no nonparametric approach to this particular problem has been pursued. In this section we apply the proposed tool to measure to which extent the above-mentioned candidate factors contribute to information production in a time series of global temperatures.

To study this problem we consider the history of the last 400 years. We choose to describe the state of this system in the space $(T_i, \mathrm{STI}_i, \mathrm{GHG}_i)$, excluding time-lagged copies of these variables because small dimensions are preferable both for computational reasons and to avoid depopulation of the state space due to the curse-of-dimensionality effect. The data consist of yearly averages of physically based reconstructions of these magnitudes [47] and are depicted in Fig. 9. We here represent GHG by the standard equivalent radiative forcing for CO_2 and other well-mixed trace gases like methane, nitrous oxides, and chlorofluorocarbons. Other potentially forcing mechanisms like tropospheric aerosols were excluded to keep the description as simple as possible. Another reason to exclude volcanic forcing is that pulses of volcanism have a cooling effect, irrespective of the temperature level. Therefore, temperature *variations* should also be included as a new state-space variable to properly account for this interaction.

In Fig. 10 we show $\mathrm{IT}_{\mathrm{CO}_2\to T}$ and $\mathrm{IT}_{\mathrm{STI}\to T}$ as a function of the length scale ε, which in this case has units of temperature anomalies. We observe that (i) there is a predominant influence of CO_2 on T over all scales, and (ii) the information gained about T because of knowing CO_2 becomes significant

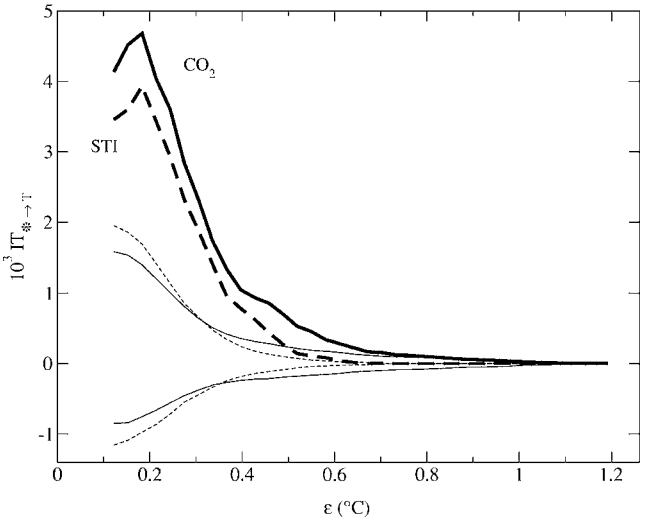

FIG. 10. Information transfer to the Mann *et al.* temperature anomalies time series T coming from CO_2 (full line) and STI (dashed line).

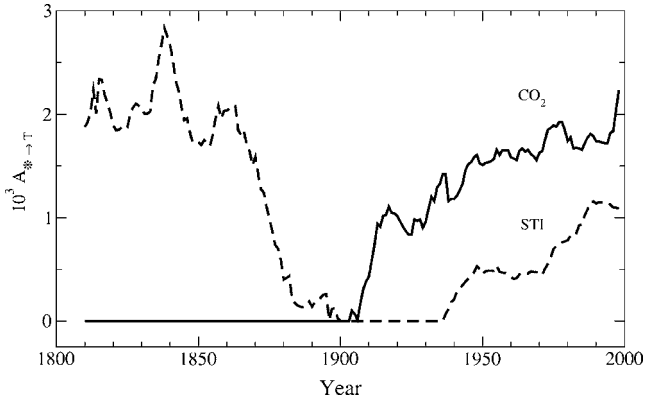

FIG. 11. Climate change problem: temporal evolution of detected causalities in sliding windows of 200 years, as measured by the area in excess of the 2σ reference level versus the interval end-point position.

at slightly bigger scales than STI. In this last case, we notice that the effects of solar irradiance variations on north-hemispheric average temperatures are to be observed most probably around 0.2 °C. This figure agrees with the results presented in [48], where the authors report that their simulations with the GISS GCM for the 20th Century show that solar irradiance changes affect superficial temperatures on Earth by 0.2 °C–0.25 °C.

We have also performed a sliding window analysis to gain some insight into the possibly changing character of the coupling between these magnitudes. Given that in this case the data availability is more limited than in the physiological application, we have only considered half-sized windows, i.e., $L=N/2=200$. We slid this 200-years window with a minimal step of 1 year, computing as before the area under $IT_{*\to T}(\varepsilon)$ beyond the 2σ significance limit imposed by the control population. We have now (arbitrarily) chosen to assign the obtained causalities to the end point of the corresponding interval. As we can see in Fig. 11, important changes in the relative influence of CO_2 and STI took place along different centuries. More precisely, we observe a dominating solar influence on temperature anomalies in the pre-industrial era that gradually declined toward the end of the 19th Century. With the advent of the 20th Century this pattern reversed to an increasing prevalence of anthropogenic activity as the main explaining factor of temperature variability, despite the also increasing explaining power that STI accused during the last six decades. The most interesting aspect of these results is given by the fact that they were obtained along a purely data-driven path, i.e., independently of any possible AOGCM misspecification error.

Summing up, this preliminary study has shown, using a model-free approach, that STI seems to account for a smaller-scale behavior of global temperatures than GHG. However, it must be noted that we have employed only one particular set of climate reconstructions covering the history of the last 400 years, which represents a very limited time span for paleoclimatic time scales. An extensive study using the many millenial reconstructions of both temperature anomalies [49] and solar irradiance [50] that have been achieved by different authors in recent years will be the subject of future work.

V. CONCLUSIONS

We have presented a new model-free, information-theoretical tool useful to compare the relative influence of two or more external dynamics impinging on a system of interest. Using a synthetic example and two interesting real-world applications belonging to dissimilar fields, we have shown that this methodology allows us to make a rich and scale-resolved analysis—with many potential research applications both within and outside nonlinear dynamics and physics.

The main conclusion of our analysis of the cardiorespiratory interaction is that a bivariate description of the heart rate–respiration feedback is limited and fails to provide a sufficient representation for this problem. In particular, their interaction with the baroreceptor system of control must be also encompassed. In accordance with the intrinsic change in dynamics that is characteristic of dynamical diseases, our results support the contention of a blocking of respiratory sinus arrhythmia by sleep apnea. We emphasize that systematic tests must be performed on a more extensive database in order to corroborate the general validity of this conclusion. This will be the subject of future work.

The concrete application to the attribution problem of climate change revealed that the emission of heat-trapping greenhouse gases—as represented by CO_2—has significantly affected global temperatures on bigger scales than the solar radiative output. Our result is in agreement with the established consensus in the literature but, being itself a data-driven approach, has the advantage of avoiding possible model misspecification errors. However, we would like to point out that our results must be considered preliminary as they were obtained on a restricted database and need to be ascertained using newly available millenial reconstructions [49,50].

Finally, a note of caution concerning nonstationarity. Changing environmental or internal conditions are ubiquitous in real data and constitute a critical issue not only for this work, but for the whole time series analysis branch. That physiological or geophysical data cannot be considered stationary is widely accepted, but few attempts have been made in order to develop statistical methods to appropriately cope with this situation. With reference to our formal derivation, nonstationarity means that we cannot interpret the measured visitation frequencies in phase space as true (single or conditional) probabilities. However, at the expense of abandoning the attractive physical interpretations of the stationary case, useful conclusions can still be drawn. This perspective has also been adopted in previous studies in the literature, and amounts to considering the measures as having essentially a relative character and employing them only in comparison with control populations or neighboring windows. By slightly relaxing the stationarity condition of applicability, the analysis can be actually enriched by identifying periods of weaker and stronger coupling between the indices.

ACKNOWLEDGMENT

This work was supported by the Alexander von Humboldt Foundation.

[1] S. J. Schiff, P. So, T. Chang, R. E. Burke, and T. Sauer, Phys. Rev. E **54**, 6708 (1996).
[2] M. Le Van Quyen, J. Martinerie, C. Adam, and F. J. Varela, Physica D **127**, 250 (1999).
[3] K. J. Blinowska, R. Kuś, and M. Kamiński, Phys. Rev. E **70**, 050902(R) (2004).
[4] J. Arnhold, P. Grassberger, K. Lehnertz, and C. E. Elger, Physica D **134**, 419 (1999).
[5] A. Schmitz, Phys. Rev. E **62**, 7508 (2000).
[6] R. Quian Quiroga, J. Arnhold, and P. Grassberger, Phys. Rev. E **61**, 5142 (2000).
[7] L. Pecora and T. L. Carroll, Int. J. Bifurcation Chaos Appl. Sci. Eng. **10**, 875 (2000).
[8] M. Wiesenfeldt, U. Parlitz, and W. Lauterborn, Int. J. Bifurcation Chaos Appl. Sci. Eng. **11**, 2217 (2001).
[9] M. Paluš, V. Komárec, Z. Hrnčíř, and K. Štěrbová, Phys. Rev. E **63**, 046211 (2001).
[10] M. G. Rosenblum and A. S. Pikovsky, Phys. Rev. E **64**, 045202(R) (2001).
[11] M. Paluš and A. Stefanovska, Phys. Rev. E **67**, 055201(R) (2003).
[12] M. Chávez, J. Martinerie, and M. Le Van Quyen, J. Neurosci. Methods **124**, 113 (2003).
[13] Y. Chen, G. Rangarajan, J. Feng, and M. Ding, Phys. Lett. A **324**, 26 (2004).
[14] T. Schreiber, Phys. Rev. Lett. **85**, 461 (2000); A. Kaiser and T. Schreiber, Physica D **166**, 43 (2002).
[15] C. Dicks and J. A. DeGoede, in *Global Analysis of Dynamical Systems*, edited by H. Broer, B. Krauskopf, and G. Vegter (IoP Publishing, London, 2001).
[16] R. K. Otnes and L. Enochson, *Digital Time Series Analysis* (Wiley, New York, 1972).
[17] B. Pompe, J. Stat. Phys. **73**, 587 (1993); J. H. Xu, Z. R. Liu, R. Liu, and Q. F. Yang, Physica D **106**, 363 (1997).
[18] N. Wiener, in *Modern Mathematics for Engineers*, edited by E. F. Beckenbach (McGraw-Hill, New York, 1956).
[19] C. W. J. Granger, Econometrica **37**, 424 (1969).
[20] C. E. Shannon, Bell Syst. Tech. J. **27**, 379 (1948).
[21] I. M. Gelfand and A. M. Yaglom, Am. Math. Soc. Transl. **12**, 199 (1959).
[22] As such, the most probable scenario will be that x_i, y_i and z_i are vectors performing a time-delayed pseudo-phase-space reconstruction of the internal dynamics of X, Y, and Z.
[23] H. Pi and C. Peterson, Neural Comput. **6**, 509 (1994).
[24] To avoid statistical fluctuations in the limit of a small bandwidth, we adopted the prevention of only considering those states satisfying $p(x_i)$, $p(y_i)$, and $p(z_i)$ bigger than $N/10$.
[25] Observing time or computational constraints, the number of bootstrap experiments may, of course, be tailored to each particular case. The control population size was set to 1000 in this work.
[26] D. R. Rigney *et al.*, in *Time Series Prediction: Forecasting the Future and Understanding the Past*, edited by A. S. Weigend and N. A. Gershenfeld (Addison-Wesley, Reading, MA, 1994).
[27] M. G. Rosenblum, L. Cimponeriu, A. Bezerianos, A. Patzak, and R. Mrowka, Phys. Rev. E **65**, 041909 (2002).
[28] Frequent arousals, although necessary for breathing to restart, prevent the patient from getting enough restorative, deep sleep. Early recognition and treatment is important because it may be associated with irregular heartbeat, high blood pressure, heart attack, and stroke.
[29] J. Bhattacharya, E. Pereda, and H. Petsche, IEEE T. Syst. Man Cy. B **33**, 1083 (2002).
[30] N. Ancona, D. Marinazzo, and S. Stramaglia, Phys. Rev. E **70**, 056221 (2004).
[31] D. E. Lerner, Physica D **97**, 563 (1996).
[32] L. D. Iasemidis, J. C. Sackellares, H. P. Zaveri, and W. J. Williams, Brain Topogr **2**, 187 (1990).
[33] K. Lehnertz and C. E. Elger, Phys. Rev. Lett. **80**, 5019 (1998).
[34] J. Martinerie, C. Adam, M. Le Van Quyen, M. Baulac, S. Clémenceau, B. Renault, and F. Varela, Nat. Med. **4**, 1173 (1998).
[35] *IPCC Third Assessment Report—Climate Change 2001: The Scientific Basis*, edited by J. T. Houghton *et al.* (Cambridge University Press, Cambridge, 2001).
[36] P. D. Jones, J. Clim. **7**, 1794 (1994); J. E. Hansen and S. Lebedeff, J. Geophys. Res. **92**, 13345 (1987); K. Ya. Vinnikov, P. Ya. Groisman, and K. M. Lugina, J. Clim. **3**, 662 (1990).
[37] J. L. Lean, J. Beer, and R. S. Bradley, Geophys. Res. Lett. **22**, 3195 (1995); A. Robertson *et al.*, J. Geophys. Res. **106**, 14783 (2001).
[38] D. T. Shindell, R. L. Miller, G. A. Schmidt, and L. Pandolfo, Nature (London) **399**, 452 (1999); T. Crowley, Science **289**, 270 (2000); T. P. Barnett, D. W. Pierce, and R. Schnur, *ibid.*

292, 270 (2001).
[39] P. A. Stott *et al.*, Science **290**, 2133 (2000).
[40] U. Neff *et al.*, Nature (London) **411**, 290 (2001); G. Bond *et al.*, Science **294**, 2130 (2001); D. V. Hoyt and K. H. Schatten, *The Role of the Sun in Climate Change* (Oxford University Press, New York, 1997).
[41] D. T. Shindell *et al.*, Science **284**, 305 (1999).
[42] J. E. Hansen *et al.*, in *Climate Modeling: Past, Present and Future*, edited by D. Randall (Academic Press, San Diego, 2000), p. 127.
[43] D. Rind, Science **296**, 673 (2002).
[44] R. Govindan *et al.*, Phys. Rev. Lett. **89**, 028501 (2002).
[45] G. Hegerl *et al.*, Clim. Dyn. **16**, 737 (2000); S. F. B. Tett *et al.*, Nature (London) **399**, 569 (1999); L. Berliner, R. Levine, and D. Shea, J. Clim. **13**, 3805 (2000); C. Schönwiese, Energy & Environ. **9**, 589 (1998); K. Hasselmann, Clim. Dyn. **13**, 601 (1997).
[46] C. Diks and M. Mudelsee, Phys. Lett. A **275**, 407 (2000).
[47] The data sources for this study were the following: STI from J. L. Lean, Geophys. Res. Lett. **27**, 2425 (2000); equivalent radiative forcing for GHG from D. M. Etheridge *et al.*, J. Geophys. Res. **101**, 4115 (1996); and temperature anomalies from M. E. Mann, R. S. Bradley, and M. K. Hughes, Geophys. Res. Lett. **26**, 759 (1999).
[48] D. T. Shindell *et al.*, Science **294**, 2149 (2001).
[49] K. R. Briffa, Quat. Sci. Rev. **19**, 87 (2000); T. J. Crowley and T. S. Lowery, Ambio **29**, 51 (2000); J. Esper, E. R. Cook, and F. H. Schweingruber, Science **295**, 2250 (2002); P. D. Jones, K. R. Briffa, T. P. Barnett, and S. F. B. Tett, Holocene **8**, 455 (1998); S. McIntyre and R. McKitrick, Energy & Environ. **14**, 751 (2003); J. Luterbacher *et al.*, Science **303**, 1499 (2004); A. Moberg *et al.*, Nature (London) **433**, 613 (2005).
[50] E. Bard, G. Raisbeck, F. Yiou, and J. Jouzel, Earth Planet. Sci. Lett. **150**, 453 (1997); Tellus, Ser. B **52**, 985 (2000); M. Stuiver and T. F. Braziunas, Radiocarbon **35**, 137 (1993); I. G. Usoskin, S. K. Solanki, M. Schussler, K. Mursula, and K. Alanko, Phys. Rev. Lett. **91**, 211101 (2003); S. K. Solanki *et al.*, Nature (London) **431**, 1084 (2004).