Unveiling causal interactions in complex systems

Stavros K. Stavroglou[a,1] ⓘ**, Athanasios A. Pantelous[b,1]** ⓘ**, H. Eugene Stanley[c,d,1]** ⓘ**, and Konstantin M. Zuev[e]**

[a]Department of Mathematical Sciences, University of Liverpool, Liverpool, L69 7ZL, United Kingdom; [b]Department of Econometrics and Business Statistics, Monash University, Clayton, VIC 3800, Australia; [c]Center for Polymer Studies, Boston University, Boston, MA 02215; [d]Department of Physics, Boston University, Boston, MA 02215; and [e]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125

Throughout time, operational laws and concepts from complex systems have been employed to quantitatively model important aspects and interactions in nature and society. Nevertheless, it remains enigmatic and challenging, yet inspiring, to predict the actual interdependencies that comprise the structure of such systems, particularly when the causal interactions observed in real-world phenomena might be persistently hidden. In this article, we propose a robust methodology for detecting the latent and elusive structure of dynamic complex systems. Our treatment utilizes short-term predictions from information embedded in reconstructed state space. In this regard, using a broad class of real-world applications from ecology, neurology, and finance, we explore and are able to demonstrate our method's power and accuracy to reconstruct the fundamental structure of these complex systems, and simultaneously highlight their most fundamental operations.

complex systems | causality | ecosystem | brain | CDS markets

F or centuries, philosophy illuminated the course of humanity's greatest endeavors. Science gave philosophy a methodological way of empirically testing theories and concepts that helped philosophers to become almost completely disentangled from superstitions, seeking nature's mechanisms for the first principles of phenomena (1). As an example, Thales of Miletus was able to predict the next big harvest and reserve olive presses in advance by observing the long-term impact of the weather on olive trees. Thales' predictions were accurate (2), as he was able to demonstrate profoundly that elaboration on the causes of things leads to a higher understanding of nature's mechanisms (3). This long-standing desire to understand the first principles of phenomena provides the strongest motivation for the present study.

Natural laws govern planetary to particle motions indisputably. However, when it comes to ecosystems, brain functions, and stock markets, we strive to derive first principles, causal relationships, and driving factors. This lack of clear understanding is the scourge of decision and policy makers, who will eventually follow ad hoc rules or best practices (4). Unavoidably, without a clear interpretation of the systems' elements and functions, fatal errors lie in wait (5). Nowadays, fortunately, the recent advances in data availability and computational power have created a fertile soil in which to develop fastidious tools for the deeper understanding of such unfathomable systems.

In this work, we develop a robust methodology (see *Methods* for the details) for detecting the hidden structure in dynamic complex systems. In practice, identifying the most important components of a dynamic complex system and its causal interactions provides an important step toward optimizing the performance and ensuring the stability in its operations (6–9). Our aim is to effectively demonstrate and scholastically test the method's power and accuracy to reconstruct the fundamental structure of complex systems, also highlighting the most essential operations and components. In particular, for one-step-ahead predictions on time series with a priori interdependencies known, our method demonstrates a remarkable accuracy of 90% over 100,000 simulations. Furthermore, to clearly reveal the multidisciplinary nature of our treatment and its robustness, we delve into three highly complex systems from ecology, neurology, and finance, which often have a large component of noise in

data. The present paper expands our understanding of dynamic complex systems.

Applying our method in three distinct areas of research where we already have an a priori knowledge of the crucial components and operations, we reconstruct the most fundamental structure and convincingly evaluate the effectiveness this methodology provides. In this direction, first for a desert ecosystem, we capture both the meaningful invasion and subsequent assimilation dynamics of the invader plant species, *Erodium cicutarium*, as well as the effects of drought as charted from precipitation and temperature. Second, for a brain activity experiment, we explore and are able to detect an expected (from literature) more intense activity in the frontal region of the control (compared to alcoholic) brain, a negative regime in the alcoholic brain between frontal and parietal regions associated with motor functions, as well as higher concentration of activity in the visual cortex of the control brain. Finally, for a set of banking credit default swaps (CDSs), we capture the driving force of Nordic banks, which is confirmed by the International Monetary Fund; the competitive role of German banks given their balance sheets; as well as the central role of certain banks during the 2007 to 2008 crisis.

Results

Overseeing Ecosystem Interdependencies. Ecosystems are characterized by recurring perturbations, swinging among multiple equilibria and chaotic disturbances. A small change in the native

> ### Significance
>
> Patterns in nature and society are described as complex systems due to their complicated and highly interconnected properties. Capturing the ebb and flow of their structures sheds a light on our better understanding of nature's rules and social connectedness. In this context, a methodology is proposed that unveils the most important operations and components of complex systems. The method's power is effectively demonstrated by reconstructing the essential structure of a desert ecosystem, discovering distinguishing features of the alcoholic brain, and locating key assets in the CDS market. The proposed framework serves as an exceptional tool for decision and policy makers, and its demonstrated effectiveness establishes its excellent potential for capturing hidden interactions in a much broader area of applications.

pool of species can have unpredictable impacts on the long-term balance of a given ecosystem (10). Environmental sentinels are concerned with species invasions and the impact of the weather on erratic regions such as desert ecosystems.

Thus, we employ our methodology in a dataset from a Chihuahuan desert scrubland site established in 1977 near Portal, Arizona (11), which contains four types of measurements: weather variables, quantities of various rodent species, several plant species, and some ant species, a detailed list of which can be seen in *SI Appendix*, Table S2. Our primary purpose is to retrieve, on the one hand, the causal interdependencies centered around the invader species, *Erodium cicutarium* (12, 13), and on the other hand to track the traceable impact of the weather on the ecosystem. To infer the type of influence among each species, an adequate backtesting procedure is employed, enabling us to assess the type of interdependence. Finally, we use the maximum spanning tree (14) algorithm to eliminate weak interdependencies, and thus to keep the most influential ones. Thus, the strongest links will also be the most meaningful in the ecosystem.

Diagnosing Disorders from Brain Activity. The brain, as a system of synaptic activity, is affected by most if not all mental disorders. For example, people suffering from alcoholism tend to exhibit adverse effects in their social life due to the neurotoxic effects on the brain, especially the frontal region. Sometimes, it even leads to persistent functional changes in brain neural circuits (15, 16). Principals of large-scale treatment programs can benefit from tools that are able to identify factors that differentiate afflicted subjects from control ones.

Inspired by the apparent impact of alcohol on the brain, we use a dataset made available publicly by Henri Begleiter of the Neurodynamics Laboratory of the State University of New York Health Center in Brooklyn (17). We use electroencephalography (EEG) measurements from 10 alcoholic and 10 control subjects. The dataset contains recordings from 64 electrodes placed on the subjects' scalps, which were sampled at 256 Hz (3.9-ms epoch) for 1 s. For our analysis, we consider each subject's exposure to a single stimulus of object pictures chosen from a curated picture set (18). The electrode positions were located at standard sites (Standard Electrode Position Nomenclature according to the American Electroencephalographic Association). The data collection process is described in detail in ref. 19. Additionally, summary details for the electrodes corresponding to specific brain regions are provided in *SI Appendix*, Table S3. Our purpose is to reconstruct the vital causal structure of the alcoholic brain compared to the control one. To that end, we perform backtesting to infer at each time step the type of causality for each pair of electrodes.

Monitoring Derivatives' Systemic Risk. Ever since the inauguration of derivative financial products, such as options and CDSs, the selection and subsequent management of portfolios has become increasingly challenging. Furthermore, all market participants are intrinsically linked, and a small decision by one can have far-reaching consequences for the market. Therefore, fund managers need to constantly investigate the ever-increasing volume of data, to optimize decision making and mitigate systemic risk.

Banks, with the incentive of hedging risk with respect to their lending operations, as well as freeing up-regulatory capital, have been the prevalent actors in the CDS market. By March 1998, the global CDS market was estimated at about $300 billion, with JP Morgan alone accounting for about $50 billion of this (20). Starting from early 2008, the global financial crisis has been quite intertwined with the role of banking CDSs. Nordic and German banks have been key components of the global financial network from 2008 onward. This motivates us to investigate further the interdependencies of banking CDSs and test whether our method can identify the de facto key players during global financial crisis and postcrisis periods. We use a dataset of daily CDS spreads

from the banking sector with a 5-y maturity (*SI Appendix*, Table S4) spanning from December 14, 2007, to May 13, 2019. The time series were retrieved from Thomson Datastream. To assess the nature of causality, we perform backtesting to infer at each time step how each stock influences each other stock daily.

Tracking Invasion Dynamics and Weather Impact in a Desert Ecosystem. During the "preinvasion" period (Fig. 1*A*), *Erodium cicutarium* (the invader) accounted for a very small percentage of the local flora (12, 13). This information is captured with our method given that two species of ants and one species of plants are negatively related to the invader, attesting to an underlying hostility. At the "breakout" of the entrenchment (Fig. 1*B*), the invader's abundance rose to account for 25% of the flora measured (21, 22), probably related to the positive influence from an ant species and the subsequent (Fig. 1*C*) positive causality from some plant and ant species. Still, however, another plant species had a negative causality on the invader, a pattern we can also see in the preinvasion period. Later (Fig. 1*D*), despite some insisting negative influences on the invader's abundance, an ant species is found to be positively associated with the invader. Down the line (Fig. 1*E*), we find again that the invader is involved in a mixed
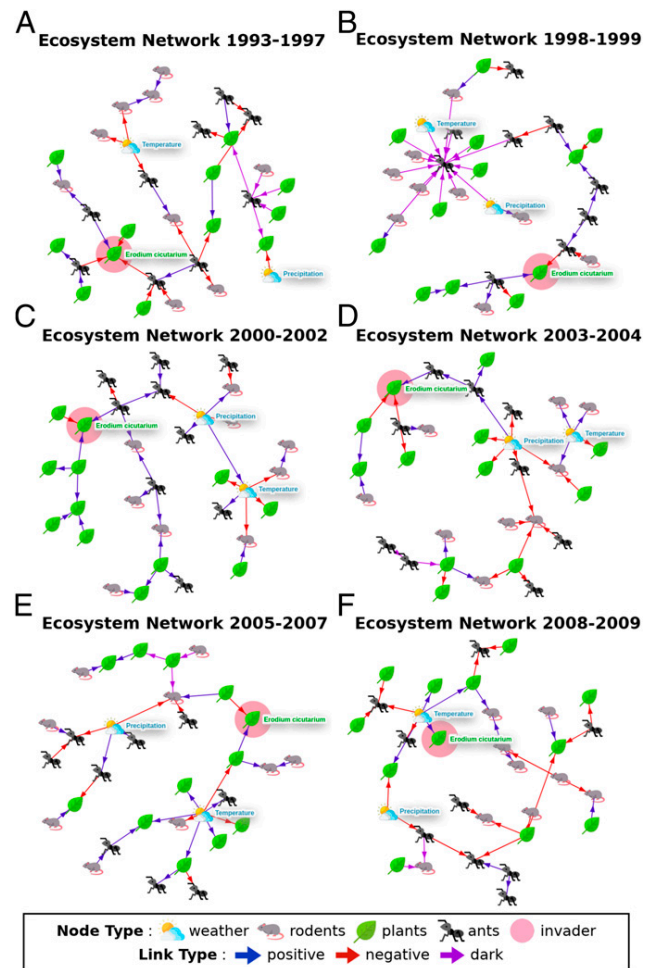
Fig. 1. Cumulative causal networks (using Eq. 11) for six separate periods: (*A*) from 1993 to 1997, before the "aggression" of *Erodium cicutarium*; (*B*–*E*) from 1998 to 2007, invasion periods of interest; (*F*) from 2008 to 2009, postinvasion period. The node icon is representative of the node's type (ants, plants, rodents, weather). The link color denotes type of causality (blue for positive, red for negative, and purple for dark).

Stavroglou et al.

triangle, with a plant species affecting it positively, and another ant species negatively. In the final period (Fig. 1*F*), only temperature affects the invader, suggesting an imminent assimilation with the rest of the ecosystem. The main insight here is that sporadic positive causalities on the invader species during the postinvasion period (Fig. 1 *C*–*F*) probably aided its successful spread in the ecosystem.

As far as the impact of the weather is concerned (Fig. 1*A*), both temperature and precipitation negatively impact two rodent species, one ant species, and one plant species, attesting to the severe drought that occurred in this period (23). Later (Fig. 1*B*), we observe the development of a dark causality regime, again involving temperature and precipitation, with an ant species at its center. Subsequently (Fig. 1*C*), temperature and precipitation play a persistent driving role in the rest of the ecosystem, in both positive and negative ways, with precipitation later claiming more of a negative force (Fig. 1*D*) and reverting to a more balanced role thereafter (Fig. 1*E*). Ultimately (Fig. 1*F*), only temperature maintains a central role in the ecosystem, affecting plant species in a positive way. However, the fact that this period is characterized by a drought is captured by two ant species being affected by a negative causality from temperature (16). For details on the exact species, see *SI Appendix*, Figs. S1–S6.

Revealing Distinct Features in Alcoholic Brain Networks. In Fig. 2, we

are comparing cumulative adjacency matrices of the "average" alcoholic and control subjects, where darker colors correspond to greater accumulated intensity, according to Eq. **11** of our algorithm. Apparently, the frontal region's positive interdependencies of the average alcoholic brain (Fig. 2*A*) are much fainter compared to the average control brain (Fig. 2*B*). This finding is allegedly related to the exhaustion of the frontal lobe due to the neurotoxic effects of alcohol (15, 16).

However, in terms of negative structure, it is evident that the average alcoholic brain has two specific regions in the adjacency matrix (Fig. 2*C*), with much more intense interdependencies than in the average control brain (Fig. 2*D*). These two regions translate to a negative causal regime, between frontal and parietal regions. Frontal region is responsible for the motor functions, while parietal region is responsible for the perception of space as well as navigation. Our results suggest that, in the average alcoholic brain, these two regions cause opposite electrical fluctuations on each other. This is consistent with the known motor impairments as well as sensory handicaps found in an alcoholic (24–26).

Distinctive features are also discovered in the microstructure of dark-type interactions. Most notably, in the average alcoholic brain, the voltage measurement from electrode CZ (rightmost of central
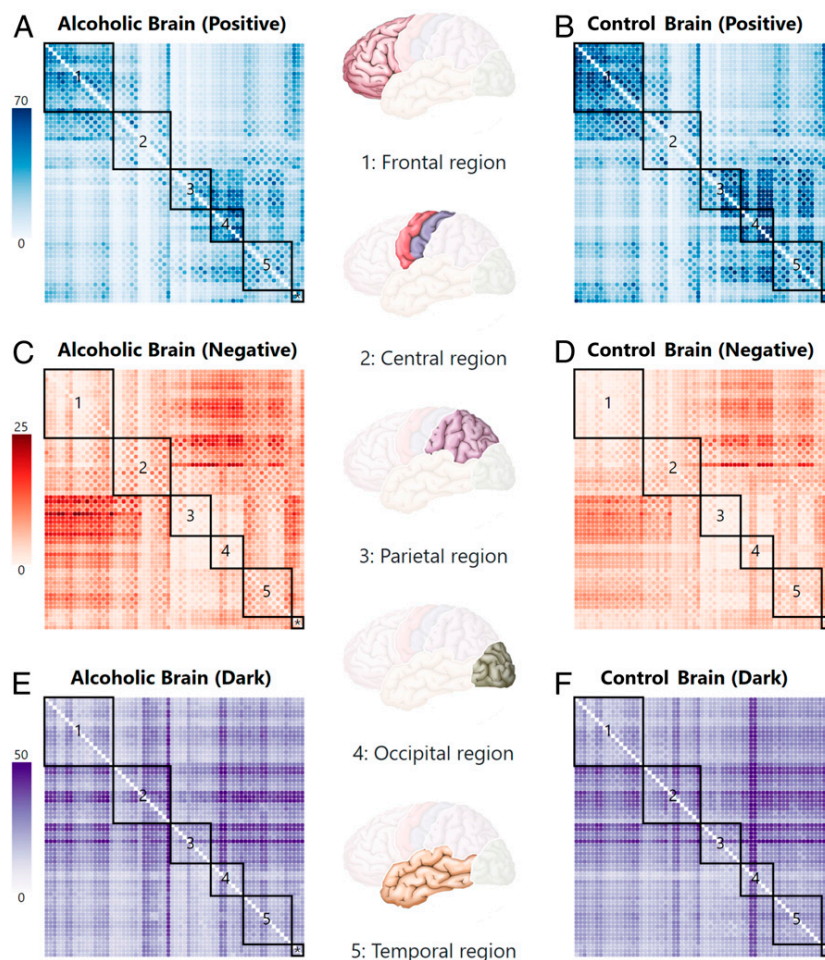
Fig. 2. Cumulative adjacency matrices for the average positive/negative/dark network structures of alcoholic (*A*, *C*, and *E*) and control subjects (*B*, *D*, and *F*) for the whole experiment duration. The darker color denotes higher accumulated link strength. Positive cumulative interdependencies (aggregating with Eq. **11**) range from 0 to 70 for all of the time horizon of the experiment. Similarly, negative cumulative interdependencies range from 0 to 25, and dark cumulative interdependencies range from 0 to 50. Moreover, box 1 corresponds to frontal region, box 2 corresponds to central region, box 3 corresponds to parietal region, box 4 corresponds to occipital region, box 5 corresponds to temporal region, and box * concerns auxiliary electrodes.

Stavroglou et al.

region) is affected consistently by all other electrodes (Fig. 2*E* and *SI Appendix*, Fig. S11). This pattern is absenting from the average control brain, which exhibits stronger causality on electrodes PO7 and PO8 (Fig. 2*F* and *SI Appendix*, Fig. S12), which are associated with properties related to visual memory (occipital region). Interestingly enough, the occipital region is involved in the processing of pictures, the region of interest in this experiment. Our analysis suggests a higher influence of occipital region from all brain regions in the control brain, a fact already reflected in the brain research literature (19, 27, 28). For details on the exact electrodes, see *SI Appendix*, Figs. S7–S12.

Detecting Persistent Causal Relationships and Influential Assets in the CDS Market. The most straightforward way to rank CDSs' contribution to systemic portfolio risk is via influence exerted and

influence received. Effectively, we can become aware of which are the CDSs that influence others, while at the same time receiving less influence. In Fig. 3, we present a bubble plot where the *x* axis corresponds to cumulative out-strength centrality and *y* axis corresponds to cumulative in-strength centrality (both centralities calculated from the pattern causality networks aggregated at each time step).

We observe that, in terms of positive interdependencies (Fig. 3*A*), the layout of the CDS causality structure seems to be arranged in a homogeneous manner, suggesting that, when considering both exerted and received influence, the majority of CDSs seem to exhibit a balance between the two. Notably, the most influential CDSs are Svenska Handelsbanken, Nordea Bank AB, and Skandinaviska Ensk Banken (Fig. 3*B*; see also *SI Appendix*, Table S5 for the top
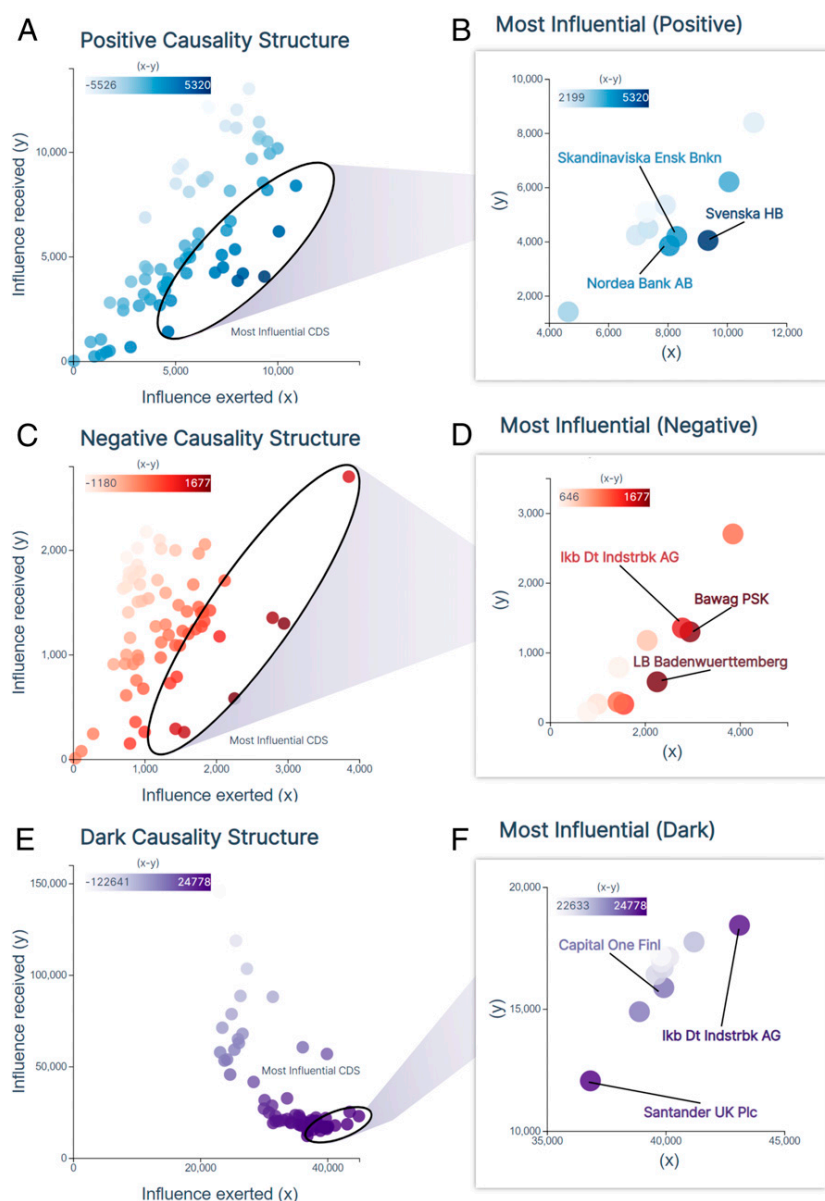

Fig. 3. Bubble plot of CDSs in terms of exerted (*x*) and received (*y*) cumulative influence (out- and in-strength centrality, respectively). Strength centralities (values in the axes) were calculated using as weights the cumulative weights from the aggregate adjacency matrices for the whole-time period. The color scale is according to (*x* − *y*), thus giving to more influential CDS a darker shade. This process is done separately for positive (*A*), negative (*C*), and dark (*E*) causality. We also focus on the top 10 most influential CDSs according to (*x*) − (*y*) explained above for each category (*B*) for positive, (*D*) for negative, and (*F*) for dark causality. The focus allows us to peek into the top CDSs in terms of every type of causality.

Stavroglou et al.

10). This result suggests that the specific Nordic banks' CDSs had the highest same-direction predictive capacity on the rest of the CDSs in our dataset. This result might associate with the fact that the Nordic banks were experiencing significantly higher loan-to-deposit ratios than all other banks, leaving them quite exposed to systemic risk, thus making their CDS spreads quite the driving market force (29). Moreover, Svenska Handelsbanken has been a center of attention in terms of its innovative banking model (30).

A similar structure is evident with regards to negative interdependencies (Fig. 3C), although a bit more dispersed, implying a sharper difference between influence exerted and received. The most influential here are Landesbank Badenwuerttemberg, Bawag PSK, and Ikb Deutschet Industriebank AG (Fig. 3D; see also *SI Appendix*, Table S5 for the top 10). Notably, these are German banks, and, in the period under study, they were found to hold substantially large amounts of sovereign bonds in their balance sheets (31), effectively making them the biggest players in the sovereign derivatives market.

Ultimately, contemplating Fig. 3E, we deduce that the causality structure of the dark interdependencies is different compared to the ones observed in the positive and negative interdependencies. CDSs receiving much influence, exert much less, while CDSs exerting much influence, receive less, compared to the previous two cases (positive and negative). In this case, the most influential CDSs are those of Santander UK PLC, Ikb Deutschet Industriebank AG, and Capital One Financial (Fig. 3F; see also *SI Appendix*, Table S5 for the top 10). At first sight, these banks seem unrelated; however, they were found to be at the very center of the 2007 to 2008 crisis (32, 33). For details on the exact CDS interdependencies, see *SI Appendix*, Figs. S13–S15 and Table S5.

Discussion

In this work, we introduce a framework for the detection of latent and elusive structures in causal networks. Our method is based on short-term predictions drawn from information embedded in a reconstructed state space. The prudent algorithmic design reveals time series causalities in three distinct types, i.e., positive (same-direction), negative (opposite-direction), and dark (mixed-direction) predictive relationships. This targeted partition allows the unique identification of persistent causal structures and dominant influences that would otherwise be lost in the noise of disparate causalities (if we did not discern the three types of interactions). Applying this method to a set of time series measurements from a given complex system allows us to perceive deeply rooted causalities for each of the three types separately. We demonstrate our method's power to discern the most fundamental components, i.e., the "backbone" that drives a system's evolution in three different disciplines.

As a first challenge, we tested our method on a desert ecosystem with imperfect measurements of weather conditions, as well as fauna and flora abundances. From observation (12, 13), it was known that this ecosystem experienced an exotic species invasion as well as two periods of severe drought. Our method was able to quantitatively capture the invasion's dynamics as well as some extra information regarding possible "inside assistance" for the invader species. Moreover, the central role of the weather, both during the droughts and in the other periods, was effectively in tandem with empirical findings (23). Next, we tested our method on a setting from neurology. Well-established literature (24–26) had noted alcoholism's impact on the frontal lobe. Through our method, we found much fainter positive interdependencies in the frontal region of the average alcoholic compared to the average control brain. Furthermore, under the dark causality spectrum, we were able to identify the average control brain's higher activity in the occipital region (visual cortex). Finally, being aware of specific banks' highlighted roles during the last decade (29, 30), we wanted to test our method's capacity to reconstruct the CDS causal network, while capturing

the most impactful components. Indeed, our method was able to identify the high impact of Nordic and German banks on the rest of the banks' CDSs, as well as banks whose role was very central during the 2007 to 2008 financial crisis. In each case, we are able to reveal the most important factors driving the rest of the system under scrutiny.

Finally, the proposed method can capture a range of causal links in a variety of complex systems. However, from these perspectives, we would like to see the application of the suggested methodology beyond the presented examples, and its reach extended to a much broader class of topics.

Methods

We introduce a method that unveils the structure of complex systems through time series data. Thus, taking a pair of time series and testing it for causality, we check whether, how much, and in which direction X causes Y. In this regard, first, we reconstruct the shadow attractors, i.e., their time-delayed representations on at least a two-dimensional space. Finally, we test X's ability to predict Y's values. The better the prediction accuracy is, the stronger the causality from X to Y.

Notational Information. Before the theoretical methodology is developed to reveal causal networks, it is necessary to introduce the following notation:

Variable	Description
$X(t) \in \mathbb{R}$	State variables (time series) of the dynamical system Ω, which operate as a function that maps points from Ω's attractor M to a real-valued scalar. X may correspond to Cartesian coordinates of the actual E-dimensional state space containing M.
$t \in \mathbb{N}$	Denotes time measured in discrete steps $t_1, t_2, t_3, \ldots$ of X's temporal evolution.
$L \in \mathbb{N}$	The time series length, which is also called the library of the time series.
$E \in \mathbb{N}$	The embedding dimension of the attractor.
$\tau \in \mathbb{N}$	The time lag we use to reconstruct a shadow attractor.
$M_X \in \mathbb{R}^E$	The shadow attractor reconstructed using time lags of $X(t)$.
$\underline{x}(t) \in \mathbb{R}^E$	The points (vectors) of M_X corresponding to the state of the system at time t.
$h \in \mathbb{N}$	The prediction horizon h steps ahead of current time t.
$L_1 \in \mathbb{R}$	The Manhattan distance measured as: $d_1(\underline{x}(t_1), \underline{x}(t_2)) = \sum_{i=1}^{E} |\underline{x}(t_1^i) - \underline{x}(t_2^i)|$.
$L_2 \in \mathbb{R}$	The Euclidean distance measured as: $d_2(\underline{x}(t_1), \underline{x}(t_2)) = \sqrt{\sum_{i=1}^{E} (\underline{x}(t_1^i) - \underline{x}(t_2^i))^2}$.
$D_X \in \mathbb{R}^{E-1}$	The distance matrix according to some metric (e.g., L_1 or L_2).
$NN_{\underline{x}(t)} \in \mathbb{R}^E$	The nearest neighbors of $\underline{x}(t)$ according to D_X.
$S_{\underline{x}(t)} \in \mathbb{R}^{E-1}$	The vector of successive percentage changes of $\underline{x}(t)$.
$\hat{S}_{\underline{x}(t+h)} \in \mathbb{R}^{E-1}$	The vector of successive percentage changes of $\underline{x}(t+h)$.
$P_{\underline{x}(t)} \in \mho^E$	The current pattern of $\underline{x}(t)$.
$\hat{P}_{\underline{x}(t+h)} \in \mho^E$	The estimated forecasted pattern of the affected variable. It is extracted as the signature of $\hat{S}_{\underline{x}(t+h)}$.
$PC[P_X, P_Y, t] \in \mho^E$	The pattern causality (PC) matrix, which is a 3D array with dimensions $(3^{E-1}, 3^{E-1}, L)$ that models the influence strength.

A Framework of Causality Assessment. The predictive capability of this approach is assessed by establishing a causal relationship between time series. While, in ref. 34, the influence from X to Y is merely quantified by

comparing patterns from contemporaneous neighborhoods of M_X and M_Y, here we investigate this relationship further using patterns from M_X's current neighborhood to predict M_Y's future patterns (h steps ahead of time t). In other words, the strong predictive power of our treatment is deployed by the algorithm formulated below. However, to demonstrate the worth of our contribution, the mathematical formalities (i.e., lemmas, theorems, and their proofs) are delineated in *SI Appendix*, sections 6 and 7. In particular, in what follows, according to *SI Appendix*, Lemma 1, M_X is said to strongly influence M_Y in an absolute way if all values of M_Y are affected by M_X, which will be tested each time we accurately predict a future pattern of M_Y, i.e., when Eq. **4** equals Eq. **9**. Furthermore, the strength of the influence is calculated by the intensity ratio, see Eq. **10**, and we expect *SI Appendix*, Lemma 3 to hold, which states that some (and not all) of M_Y's values are affected by M_X. *SI Appendix*, Lemmas 2 and 4 suggest that, if M_X influences M_Y, then subsequently X influences Y, effectively allowing conclusions from attractor analysis to be interpreted for raw time series as well. Finally, *SI Appendix*, Theorems 1, 2, and 3 separate the nature of influence into positive, negative, and dark, respectively, and they are included at the end of our method when we use the PC matrix (*SI Appendix*, Tables S6 and S7) to support the visualization of our treatment.

Shadow Attractors Reconstruction. We create the shadow attractors, M_X and M_Y, for X and Y, respectively, by finding the optimal pair (E, τ). In particular, we initially compare the predicting accuracy for a whole range of reasonable embedding values of E and τ, and then we calculate the distance matrices, D_X and D_Y (e.g., either using the L_1 norm if we want to treat all distances equally or L_2 if we want to penalize bigger distances), among all vectors in M_X and M_Y:

$$X = \{X(1), \ldots, X(L)\} \Rightarrow M_X$$
$$= \begin{pmatrix} \underline{x}(1) = <X(1), X(1+\tau), \ldots, X(1+(E-1)\tau)> \\ \underline{x}(2) = <X(2), X(2+\tau), \ldots, X(2+(E-1)\tau)> \\ \vdots \\ \underline{x}(L-(E-1)\tau) = <X(L-(E-1)\tau), X(L-(E-2)\tau), \ldots, X(L)> \end{pmatrix},$$

and

$$D_X = \begin{pmatrix} d(\underline{x}(1), \underline{x}(1)) & \cdots & d(\underline{x}(1), \underline{x}(L-(E-1)\tau)) \\ \vdots & \ddots & \vdots \\ d(\underline{x}(L-(E-1)\tau), \underline{x}(1)) & \cdots & d(\underline{x}(L-(E-1)\tau), \underline{x}(L-(E-1)\tau)) \end{pmatrix}. \quad [1]$$

We derive M_Y and D_Y similarly.

Once the shadow attractors are derived, we obtain access to the reconstructed topology of the complex system. In the next step, we parse the local areas in the attractors and extract useful information for the prediction and the causality inference.

The Nearest Neighbors and Their Future Projections. For each point $\underline{x}(t)$ in M_X, we find its $E+1$ nearest neighbors $NN_{\underline{x}(t)}$, which is the minimum number of points needed for a bounded simplex in an E-dimensional space. From these $E+1$ nearest neighbors, we need to keep the time indices, find the corresponding points on M_Y, and project them ahead by h steps to determine the future states:

a. The projected time indices $t_{\underline{x}_1}, t_{\underline{x}_2}, ..., t_{\underline{x}_{E+1}}$:

$$NN_{\underline{x}(t)} = argmin_{(E+1)}\{d(\underline{x}(t), \underline{x}(1)), \ldots, d(\underline{x}(t), \underline{x}(t-(E-1)*\tau-h))\}$$
$$= \{NN_{\underline{x}(t_1)}, NN_{\underline{x}(t_2)}, \ldots, NN_{\underline{x}(t_{E+1})}\} \Rightarrow t_1, t_2, \ldots, t_{E+1} \overset{\text{projecting } h \text{ steps ahead}}{\Rightarrow} \quad [2]$$
$$t_1+h, t_2+h, ..., t_{E+1}+h = t_{\underline{x}_1}, t_{\underline{x}_2}, ..., t_{\underline{x}_{E+1}}.$$

b. The distance of the projected neighbors from $\underline{y}(t)$:

$$d_{\underline{y}_1} = d\left(\underline{y}(t), \underline{y}(t_{\underline{x}_1})\right), d_{\underline{y}_2} = d\left(\underline{y}(t), \underline{y}(t_{\underline{x}_2})\right), d_{\underline{y}_{E+1}} = d\left(\underline{y}(t), \underline{y}(t_{\underline{x}_{E+1}})\right). \quad [3]$$

In order to avoid any data snooping, the following must hold for all of the projections of the nearest neighbors: $t_n < t$, where $t_n \in \{t_{\underline{x}_1}, t_{\underline{x}_2}, ..., t_{\underline{x}_{E+1}}\}$. In this step, we extract the projected time indices of $\underline{x}(t)$'s neighbors' projections and use them to calculate the distances of their cotemporals $\underline{y}(t_{\underline{x}_n})$, where $t_{\underline{x}_n} := t_{\underline{x}_1}, t_{\underline{x}_2}, ..., t_{\underline{x}_{E+1}}$.

The Affected Variable's Predicted Pattern h Steps Ahead. We use the relevant information from Eqs. **2** and **3** to estimate the predicted pattern $\hat{P}_{\underline{y}(t+h)}$ of $\underline{y}(t+h)$:

$$\hat{P}_{\underline{y}(t+h)} = signature\left(\hat{S}_{\underline{y}(t+h)}\right), \quad [4]$$

where

$$\hat{S}_{\underline{y}(t+h)} = \sum_{t_n=t_{\underline{x}_1}}^{t_{\underline{x}_{E+1}}} w_{t_n}^x s_{t_n}^y, \quad [5]$$

$$w_{t_n}^x = \frac{e^{d_{t_n}}}{\sum_{t_n} e^{d_{t_n}}}. \quad [6]$$

Here, the d_{t_n} represent the distances from Eq. **3**, and

$$s_{t_n}^y = \left(\frac{Y(t_{\underline{x}_2}) - Y(t_{\underline{x}_1})}{Y(t_{\underline{x}_1})}, ..., \frac{Y(t_{\underline{x}_{E+1}}) - Y(t_{\underline{x}_E})}{Y(t_{\underline{x}_E})}\right). \quad [7]$$

Remark. $t_{\underline{x}_1}, t_{\underline{x}_2}, ..., t_{\underline{x}_{E+1}}$ correspond to the ones calculated in Eq. **2**.

Here, we are using information from M_X in order to predict M_Y's future pattern $\underline{y}(t+h)$.

The Driver Variable's Pattern. Then, we keep the current pattern of $\underline{x}(t)$, which is $P_{\underline{x}(t)}$:

$$P_{\underline{x}(t)} = signature(S_{\underline{x}(t)}), \quad [8]$$

where the signature is the way of extracting patterns from vectors, as described in *SI Appendix*.

By holding the current signature of $\underline{x}(t)$, we are able to assess both the intensity and the type of the causality from X to Y.

The Affected Variable's Real Pattern (Backtesting Process). Then, we keep the real pattern of $\underline{y}(t+h)$, which is $P_{\underline{y}(t+h)}$:

$$P_{\underline{y}(t+h)} = signature\left(S_{\underline{y}(t+h)}\right). \quad [9]$$

Here, we extract the real signature of $\underline{y}(t+h)$ and we are able to test our hypothesis for causality. In order for that to be true, the predicted pattern from Eq. **4** must be the same as the real pattern from Eq. **9**. This process is in accordance with *SI Appendix*, Lemmas 1 and 3.

The Nature and Intensity of Influence at Every Time Step t. We repeat this procedure, see Eqs. **2**–**9**, for every point of the shadow manifold M_X and fill in the PC matrix (*SI Appendix*, Tables S6 and S7) for every time step t whose influence is valid as described above. Otherwise, the PC matrix for the current t is left empty. We fill in the PC matrix, when the prediction is valid, by calculating the norms of the signatures, which are the representations of the pattern's strength, and divide the cause's norm $\left\|S_{\underline{x}(t)}\right\|$ by the effect's norm $\left\|S_{\underline{y}(t+h)}\right\|$:

$$PC[P_X, P_Y, t] = \frac{\left\|S_{\underline{y}(t+h)}\right\|}{\left\|S_{\underline{x}(t)}\right\|}. \quad [10]$$

For a normalized output, we can instead fill in the PC matrix by filtering first with the Gauss error function:

$$PC[P_X, P_Y, t] = erf\left(\frac{\left\|S_{\underline{y}(t+h)}\right\|}{\left\|S_{\underline{x}(t)}\right\|}\right), \quad [11]$$

where

$$erf(x) = \frac{1}{\sqrt{\pi}} \int_{-x}^{x} e^{-t^2} dt. \quad [12]$$

The Overall (for All t) Nature and Intensity of Causality. At this point, the produced results contain three time series, one for each type of influence (positive, negative, and dark), labeled $P(t)$, $N(t)$, and $D(t)$, respectively, indicating at each time step the intensity of the influence (from 0 to 1). Notice that, for a given t, only one of the three can be different from zero, meaning that we cannot have more than one type of influence at the same time.

Stavroglou et al.

Causal Network Analytics. Doing research in the era of big data involves the analysis of interdependencies among many time series variables. Thus, instead of just X and Y, we have N variables, i.e., $X_1, \ldots, X_N$. The variables are heretofore referred to as "nodes" of a network. Hence, the maximum number of causal interactions to be put under scrutiny is $N(N-1)$, not accounting for loops. Now, we can have a total of $N(N-1)$ resulting time series of each type [referring to $P(t), N(t)$, and $D(t)$], effectively creating three dynamic causal networks, one for each aspect (positive, negative, and dark), or symbolically: $P_k^l(t)$, referring to the intensity of positive influence at time t, from node k to node l; $N_k^l(t)$, referring to the intensity of negative influence at time t, from node k to node l; $D_k^l(t)$, referring to the intensity of dark influence at time t, from node k to node l.

Ultimately, $P_k^l(t), N_k^l(t), D_k^l(t) \forall k, l$ are the positive, negative, and dark aspects, respectively, of the causal network at time t, and can be seen as three concurrent networks, of the same nodes but with mutually exclusive links (no link can exist at the same time for more than one of the three aspects). Optionally, we can filter the network to keep only the strongest relationships by using algorithms such as the minimum/maximum spanning tree (14, 35) or the planar maximally filtered graph (36).

Strength Centrality. This metric refers to the aggregation of the weights of the links from and to the node (37). Out-strength denotes the weighted influence exerted directly on other nodes, and in-strength denotes the weighted influence received directly from other nodes. Weights, here, are calculated from Eq. 11.

Link Persistence. This measures the overall weight of a given link from node X to node Y by aggregating cumulatively across time to rank time series interdependencies on strength and persistence (38).

Complexity. The proposed method is computationally efficient for long time series (large L). The only parameters that impact our method are the time series length L and its embedding dimension E. The higher is L and/or E, the longer it will take for the distance matrices D_X and D_Y to be calculated. To extract the candidate neighbors of a point $\underline{x}(t)$, we only need the $D_X[t, 1:(t-1)]$ part of D_X (same for D_Y). Computing D_X and D_Y costs L^2E for each, and the iteration part of the main algorithm is of order $O(L)$. The total cost of our algorithm is of order $O(L^2E+L)$, with the main bulk of the calculations being that of the initial distance matrices. More details about the complexity can be found in *SI Appendix*, section 3.

Method's Validation Using Simulation. Our method has been validated using 100,000 simulations and different lengths of chains for the three types of interactions, positive, negative, and dark. Analysis and discussion of this simulation-based validation is provided in detail in *SI Appendix*, section 4. Particularly for short chain lengths, the results derived are rather impressive.

Data Availability. For the ecosystem analysis, the dataset for the Chihuahuan ecosystem near Portal, Arizona, can be accessed in ref. 39. For the EEG analysis, we use data of 20 subjects from a dataset made available publicly by Henri Begleiter of the Neurodynamics Laboratory of the State University of New York Health Center in Brooklyn, New York (17). Each subject has undergone five trials, and for each trial there are recordings in time series ($L = 256$) from the 64 electrodes' voltage measurements. Finally, the banking CDS data are available at Thomson Reuters Datastream. The R code can be accessed at https://github.com/skstavroglou/pattern_causality.

1. S. E. Toulmin, *Foresight and Understanding: An Inquiry into the Aims of Science* (Indiana University Press, Bloomington, IN, 1961).
2. G. Crawford, B. Sen, *Derivatives for Decision Makers: Strategic Management Issues* (Wiley, 1996), vol. 1.
3. Aristotle, *Aristotle's Metaphysics*, H. G. Apostle, Transl. (Indiana University Press, Bloomington, IN, 1966).
4. R. E. Bellman, L. A. Zadeh, Decision-making in a fuzzy environment. *Manage. Sci.* **17**, B141–B273 (1970).
5. J. E. Russo, P. J. Schoemaker, E. J. Russo, *Decision Traps: Ten Barriers to Brilliant Decision-Making and How to Overcome Them* (Doubleday/Currency, New York, NY, 1989).
6. C.A. Hidalgo, R. Hausmann, The building blocks of economic complexity. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 10570–10575 (2009).
7. M. Kitsak *et al.*, Identification of influential spreaders in complex networks. *Nat. Phys.* **6**, 888–893 (2010).
8. S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, S. Havlin, Catastrophic cascade of failures in interdependent networks. *Nature* **464**, 1025–1028 (2010).
9. A. Vespignani, Complex networks: The fragility of interdependency. *Nature* **464**, 984–985 (2010).
10. F. S. Chapin, III, P. A. Matson, P. Vitousek, *Principles of Terrestrial Ecosystem Ecology* (Springer Science and Business Media, 2011).
11. S. K. Morgan Ernest *et al.*, Long-term monitoring and experimental manipulation of a Chihuahuan desert ecosystem near Portal, Arizona (1977–2013). *Ecology* **97**, 1082 (2016).
12. T. J. Valone, J. Balaban-Feld, Impact of exotic invasion on the temporal stability of natural annual plant communities. *Oikos* **127**, 56–62 (2018).
13. T. J. Valone, J. Balaban-Feld, An experimental investigation of top–down effects of consumer diversity on producer temporal stability. *J. Ecol.* **107**, 806–813 (2019).
14. T. C. Hu, Letter to the editor—the maximum capacity route problem. *Oper. Res.* **9**, 898–900 (1961).
15. G. R. Breese, R. Sinha, M. Heilig, Chronic alcohol neuroadaptation and stress contribute to susceptibility for alcohol craving and relapse. *Pharmacol. Ther.* **129**, 149–171 (2011).
16. American Psychiatric Association, Diagnostic and statistical manual of mental disorders. *BMC Med.* **17**, 133–137 (2013).
17. H. Begleiter, Data from "EEG database data set." UCI Machine Learning Repository. http://archive.ics.uci.edu/ml/datasets/EEG+Database. Deposited 13 October 1999.
18. J. G. Snodgrass, M. Vanderwart, A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *J. Exp. Psychol. Hum. Learn.* **6**, 174–215 (1980).
19. X. L. Zhang, H. Begleiter, B. Porjesz, W. Wang, A. Litke, Event related potentials during object recognition tasks. *Brain Res. Bull.* **38**, 531–538 (1995).
20. G. Tett, The dream machine: Invention of credit derivatives. *Financial Times*, 24 March 2006. https://www.ft.com/content/7886e2a8-b967-11da-9d02-0000779e2340. Accessed 12 September 2019.

21. G. R. Allington, D. N. Koons, S. K. Morgan Ernest, M. R. Schutzenhofer, T. J. Valone, Niche opportunities and invasion dynamics in a desert annual community. *Ecol. Lett.* **16**, 158–166 (2013).
22. D. D. Ignace, P. Chesson, Removing an invader: Evidence for forces reassembling a Chihuahuan desert ecosystem. *Ecology* **95**, 3203–3212 (2014).
23. E. M. Christensen, D. J. Harris, S. K. M. Ernest, Long-term community change through multiple rapid transitions in a desert rodent community. *Ecology* **99**, 1523–1529 (2018).
24. A. Pfefferbaum, E. V. Sullivan, D. H. Mathalon, K. O. Lim, Frontal lobe volume loss observed with magnetic resonance imaging in older chronic alcoholics. *Alcohol. Clin. Exp. Res.* **21**, 521–529 (1997).
25. H. F. Moselhy, G. Georgiou, A. Kahn, Frontal lobe changes in alcoholism: A review of the literature. *Alcohol Alcohol.* **36**, 357–368 (2001).
26. M. T. Ratti, P. Bo, A. Giardini, D. Soragna, Chronic alcoholism and the frontal lobe: Which executive functions are impaired? *Acta Neurol. Scand.* **105**, 276–281 (2002).
27. H. Begleiter, B. Porjesz, W. Wang, A neurophysiologic correlate of visual short-term memory in humans. *Electroencephalogr. Clin. Neurophysiol.* **87**, 46–53 (1993).
28. S. Hertz, B. Porjesz, H. Begleiter, D. Chorlian, Event-related potentials to faces: The effects of priming and recognition. *Electroencephalogr. Clin. Neurophysiol.* **92**, 342–351 (1994).
29. R. Babihuga, M. Spaltro, Bank funding costs for international banks (No. 14–71). International Monetary Fund. https://www.imf.org/external/pubs/ft/wp/2014/wp1471.pdf. Accessed 12 September 2019.
30. N. Kroner, *A Blueprint for Better Banking: Svenska Handelsbanken and a Proven Model for More Stable and Profitable Banking* (Harriman House Limited, 2011).
31. C. M. Buch, M. Koetter, J. Ohls, Banks and sovereign risk: A granular view. *J. Financ. Stab.* **25**, 1–15 (2016).
32. D. H. Erkens, M. Hung, P. Matos, Corporate governance in the 2007–2008 financial crisis: Evidence from financial institutions worldwide. *J. Corp. Finance* **18**, 389–411 (2012).
33. T. Grammatikos, R. Vermeulen, Transmission of the financial and sovereign debt crises to the EMU: Stock prices, CDS spreads and exchange rates. *J. Int. Money Finance* **31**, 517–533 (2012).
34. S. K. Stavroglou, A. A. Pantelous, H. E. Stanley, K. M. Zuev, Hidden interactions in financial markets. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 10646–10651 (2019).
35. J. C. Gower, G. J. Ross, Minimum spanning trees and single linkage cluster analysis. *J. Appl. Stat.* **18**, 54–64 (1969).
36. M. Tumminello, T. Aste, T. Di Matteo, R. N. Mantegna, A tool for filtering information in complex systems. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 10421–10426 (2005).
37. A. Barrat, M. Barthélemy, R. Pastor-Satorras, A. Vespignani, The architecture of complex weighted networks. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 3747–3752 (2004).
38. S. K. Stavroglou, A. A. Pantelous, K. Soramaki, K. Zuev, Causality networks of financial assets. *J. Net. Theory Financ.* **3**, 17–67 (2017).
39. S. K. Morgan Ernest *et al.*, Long-term monitoring and experimental manipulation of a Chihuahuan desert ecosystem near Portal, Arizona (1977–2013). *Ecology*, 10.1890/15-2115.1 (2016).

APPLIED PHYSICAL SCIENCES