Hidden interactions in financial markets

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

Contributed by H. Eugene Stanley, February 28, 2019 (sent for review February 7, 2019; reviewed by Grigoris Kalogeropoulos and Eugene Neduv)

The hidden nature of causality is a puzzling, yet critical notion for effective decision-making. Financial markets are characterized by fluctuating interdependencies which seldom give rise to emergent phenomena such as bubbles or crashes. In this paper, we propose a method based on symbolic dynamics, which probes beneath the surface of abstract causality and unveils the nature of causal interactions. Our method allows distinction between positive and negative interdependencies as well as a hybrid form that we refer to as "dark causality." We propose an algorithm which is validated by models of a priori defined causal interaction. Then, we test our method on asset pairs and on a network of sovereign credit default swaps (CDS). Our findings suggest that dark causality dominates the sovereign CDS network, indicating interdependencies which require caution from an investor's perspective.

financial markets | pattern causality | complex systems | sovereign CDS networks | pairs trading

The study of complex systems is hard-wired with the understanding of time series interdependencies. Specifically, financial markets have long been analyzed as complex systems with the asset pricing being the foundation upon which structures such as financial networks are studied (1–5). Knowing the type of asset price interactions is vital to planning effective investment strategies or microtrading tactics (6). Such knowledge is even more crucial if one considers the possibilities of switching regimes in financial markets (2).

A straightforward tool to gauge asset interactions is correlation, but is known to be misleading in nonlinear systems such as stock markets (7, 8). The robust notion is that of causality (9), which has been approached via methods such as Granger causality (10), cointegration (11), transfer entropy (12), convergent cross mapping (13), and S maps (14). Nevertheless, none of the aforementioned methods provide insight regarding more complex interactions, namely dark causality.

Our framework distinguishes between types of interactions (positive, negative, and dark) with the full details disclosed in *SI Appendix* and in Movie S1. In the case of positive causality, X consistently causes the same patterns in Y; inversely, in the case of negative causality, X consistently causes the opposite patterns in Y. However, this is a small subset of the cases of causal interaction, as seen at the end of Movie S1 and in *SI Appendix*, Tables S1 and S2. In many cases X consistently causes patterns in Y which are neither the same nor opposite. It is this third case where we observe complex causal interactions of temporal patterns that give rise to what we call dark causality, an obscure, yet substantial form of influence which is completely unapproachable through the lens of previous methods that gauge the nature of interactions (*SI Appendix*, Figs. S1–S4).

Positive causality (and positive correlation) is anticipated in cases of financial assets that are well-suited for asset pair trading (15). Intuitively such assets cause the same direction changes to one another. On the other hand, negative causality is expected in cases of competing financial entities such as equities and bonds (16). The hidden nature of financial market interdependencies may be the missing link leading to a deeper understanding of market crashes and other emergent phenomena. This is why the emergence of dark causality is important, to uncover types of asset interactions which cannot be classified as purely positive or negative.

To analyze the spectrum of causal interactions, both transparent (positive, negative) and opaque (dark) in complex systems, we introduce in this article a method that is based on interactions of symbolic dynamics (patterns) (17) in reconstructed attractors (18). Using basic models, we demonstrate that our method distinguishes between positive, negative, and dark causality. Then, we apply our approach to pairs of financial assets and expose the positive nature of causality between Microsoft and Apple stocks, and a competitive interaction between S&P 500 (as proxy of stock market performance) and US government 10-y bond yield. Lastly, we illustrate the prominence of a dark causality in the global network of sovereign credit default swaps (CDS).

Nature of Causality Through Contemporaneous Patterns

In complex systems comprising deterministic or stochastic components, spatiotemporal dynamics are sculpted into a distinct attractor. According to dynamical systems theory, when two time series X and Y are causally linked they coexist in a common attractor, which is an embedded manifestation of their joint dynamic system. Consequently, each variable is imbued with information of the other's state (18, 19).

Nevertheless, the sharing of a common attractor is not sufficient to assess the nature of causality. Such intricate information is imprinted in the interplay of local spatiotemporal dynamics between the X's and Y's attractors, M_X and M_Y, respectively, which is the focal point of our approach, Pattern Causality (PC). Therefore, positive (or negative) causality from X to Y is manifested when the patterns in M_Y can accurately recall patterns of M_X and are of the same (or opposite) nature. When dark causality emerges, the

Significance

The importance of understanding the structure of financial markets has been progressively coming to a head in public discourse. With the rise and prevalence of big datasets, it is a crucial time for academics to study the interdependency outside conventional analysis. Thus, the methodology proposed in this paper does not only allow distinction between positive and negative interdependence, but additionally identifies a yet unexplored form of interaction we coin as "dark causality." We benchmark this method on asset pairs and on a network of sovereign credit default swaps, where the dominant form of interaction is that of dark causality. Our results are relevant to finance practitioners who, by virtue of their profession, tend to engage on a regular basis with retail investors.

Author contributions: S.K.S., A.A.P., H.E.S., and K.M.Z. designed research; S.K.S., A.A.P., H.E.S., and K.M.Z. performed research; S.K.S., A.A.P., H.E.S., and K.M.Z. contributed new reagents/analytic tools; S.K.S. and A.A.P. analyzed data; and S.K.S., A.A.P., and H.E.S. wrote the paper.

Reviewers: G.K., National and Kapodistrian University of Athens, Greece; and E.N., Columbia University.

The authors declare no conflict of interest.

[1]To whom correspondence may be addressed. Email: stavros.stavroglou@liverpool.ac.uk, Athanasios.Pantelous@monash.edu, or hes@bu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1819449116/-/DCSupplemental.

Published online May 13, 2019.

patterns of M_X are imprinted and therefore recallable from M_Y, yet they are of neither the same nor opposite nature. This coupling of patterns is of complex nature, hence the naming convention "dark".

Fundamentally, the idea is to see whether spatiotemporal neighborhoods of M_Y can consistently predict the patterns of counterpart neighbors of M_X. If the recalled patterns of M_X are correctly predicted, then X is causing Y to the extent of M_Y's predictive accuracy. Whether the nature of causality is positive, negative, or dark depends on the correspondence of patterns between M_X and M_Y. Our approach is visually explained in Fig. 1 below, with the full methodology recorded in *SI Appendix* and Movie S1.

PC is a method for detecting and quantifying the nature of causality which is based on two essential properties. Firstly, PC is characterized by consistency in its inference of causality's nature. In order for a causal relationship to be found as positive, negative, or dark, the neighborhoods of M_Y have to systematically estimate correctly and in a consistent way the average patterns (signatures) in M_X. Thus, ephemeral correlations or evanescent causalities do not survive the PC trial. The notion of signature is defined in Eqs. **1–5** below. Let $Y(t_j), Y(t_j-\tau), \ldots, Y(t_j-(\mathrm{E}-1)\tau)$ be neighboring values of a stock price according to some metric (e.g., Manhattan distance) with E and τ being the embedding parameters. The signature of a given neighborhood in state space is defined as the average pattern of the nearest-neighbors' (NN) pattern:

$$P_{y(t)} = \text{signature}\left(S_{y(t)}\right), S_{y(t)} \in \mathbb{R}^E, \quad \textbf{[1]}$$

where

$$S_{y(t)} = \sum_{j=1}^{E+1} w_j^y s_j^y, w_j^y \in [0,1], s_j^y \in \mathbb{R}^E, \text{for all } \mathrm{NN}_{y(t)}, \quad \textbf{[2]}$$

$$w_j^y = \frac{e^{-d\left(y(t),y\left(t_j\right)\right)}}{\sum_j e^{-d\left(y(t),y\left(t_j\right)\right)}}, d: \text{Manhattan distance}, \quad \textbf{[3]}$$

$$s_j^y = \left(\frac{y_j^{(2)} - y_j^{(1)}}{y_j^{(1)}}, \ldots, \frac{y_j^{(E+1)} - y_j^{(E)}}{y_j^{(E)}}\right), y_j \in \mathbb{R}, \quad \textbf{[4]}$$

$$\begin{aligned} y(t_j) &= \left(Y\left(t_j\right), Y\left(t_j-\tau\right), \ldots, Y\left(t_j-(\mathrm{E}-1)\tau\right)\right) \\ &= \left(y_j^{(1)}, \ldots, y_j^{(E+1)}\right), Y \in \mathbb{R}. \end{aligned} \quad \textbf{[5]}$$

Consistency is what bestows PC with the ability to go beyond the abstract quantification of causality. Secondly, the use of symbolic dynamics aids in the suppression of noise, thus making up for any expected and knowable noise in the system.

Positive Causality (Mutualism) Model

Positive causality suggests that two variables interact in such a way that changes in X cause consistently the same changes in Y. We use a mutualism model (*SI Appendix*, Eq. S19) which describes two dynamically coupled variables with variable X exerting more influence on Y than vice versa (Fig. 2*A*). As we can see in Fig. 2*D*, PC detects the positive nature of causality and is much higher from X to Y, attesting to the asymmetric influence. Positive feedback loops can be described by such a system. These loops are very crucial to the understanding of many fields, including the dynamics of ecosystems (20), physiology for cardio-excitation–contraction coupling of the heart (21), and also finance as indications of systemic risk (22), to name a few.

Negative Causality (Competition) Model

When negative causality is dominant this translates as X causing the opposite change in Y. For theoretical validation, we use a competition model (*SI Appendix*, Eq. S20) of two dynamically coupled variables with X's impact being more intense than Y's (Fig. 2*B*). It is obvious from Fig. 2*E* that our approach correctly identifies the conflicting interaction and also reveals the asymmetry between X and Y. This case is also relevant to negative feedback loops, which are present in complex food webs (23) and biological oxygen-dependent functions (24). Such loops are also

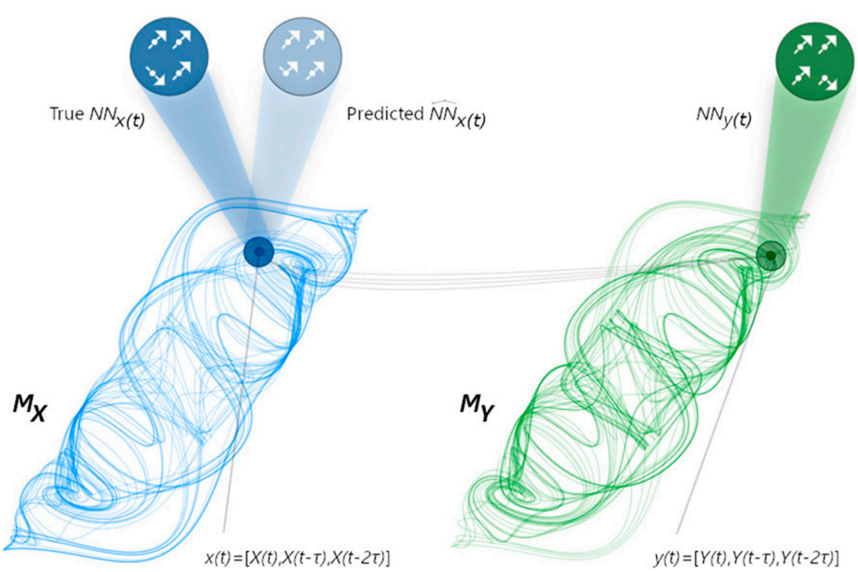

Fig. 1. PC (say from X to Y) is based on the accuracy of the "memories" (patterns up to time t) about M_X's patterns which are embedded in M_Y's counterpart spatiotemporal neighbors. If X causes Y, the NNs of point $y(t)$ ($\mathrm{NN}_{y(t)}$) will correspond temporally to the NNs of $x(t)$ ($\mathrm{NN}_{x(t)}$). This enables us to "predict" patterns in the neighborhoods of M_X using patterns from M_Y. As a next step we calculate the average pattern or signature of the predicted patterns of M_X and validate it with the actual signature. The more accurate our predictions are, the higher the PC from X to Y is. To put it simply, if X causes Y, then patterns from M_X leave their "footprints" on M_Y; thus, we can use patterns from M_Y to estimate the driving patterns from M_X.

APPLIED PHYSICAL SCIENCES

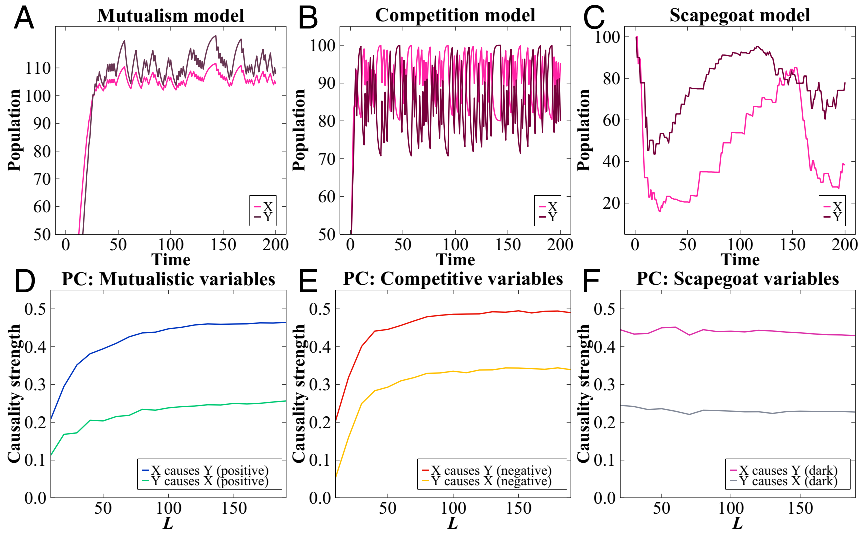

Fig. 2. Nature of causality in theoretical models. L is the time series (library) length. (*A*) Positive case: variables beneficial to each other; (*B*) negative case: variables competitive to each other; (*C*) dark case: variables involved in a persistent yet neither beneficial nor harmful relationship. (*D*) PC between mutualistic variables. (*E*) PC between competitive variables. (*F*) PC in a scapegoat relationship. Color scheme: Blue and green are used for positive causality. Red and yellow are used for negative causality. Purple and gray are used for dark causality.

encountered in the trading activity during financial crises and market crashes when investors are more pessimistic (25).

Dark Causality (Scapegoat) Model

For the case of dark causality, we employ a model (*SI Appendix*, Eq. S21) which describes a scapegoat relationship (Fig. 2*C*). More specifically we simulate the interaction between two different prey populations (e.g., lambs and rabbits) under the presence of a common predator (e.g., wolves). The design of our model allows the predator to hunt, at any given time unit, one type of prey exclusively. Each type of prey population may reproduce only when the other type is hunted in its place (e.g., when the wolves hunt lambs then the rabbits do reproduce). By calculating PC, we expose the hidden interaction between the two prey types (Fig. 2*F*), a relationship whose meaning is neither positive nor negative and falls into the category of dark causality. In finance, such relationships cannot be defined as beneficial or detrimental (as they can in ecology), but in the CDS example they are abundant and need to be further scrutinized.

Models exhibit almost ideal circumstances. This is when consistency is unhindered by observational error and system noise, which is not the case for shadow attractors reconstructed by real data. Nevertheless, even though noise and errors constrains the level of consistency, neighbors in shadow attractors can still recall significant amounts of spatiotemporal dynamics.

Pairs Trading Candidate Assets

Our first financial application of PC is on daily time series data of Apple (AAPL) and Microsoft (MSFT) retrieved from Datastream. The time span is from March 13, 1986 to August 6, 2018.

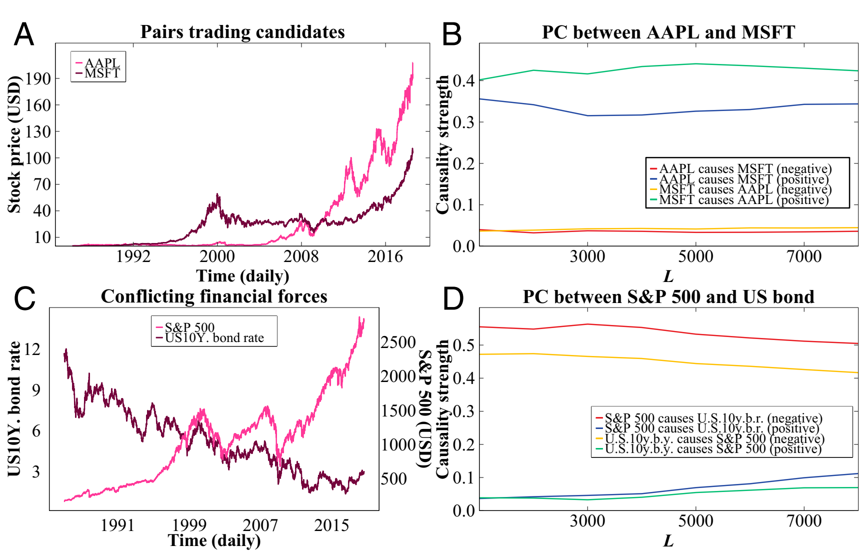

Fig. 3. Nature of causality in financial data. L is the time series (library) length. (*A*) Daily time series of AAPL and MSFT. (*B*) PC reveals their positive causal interactions. (*C*) Daily time series of S&P 500 (right y axis) and US 10-y government bond yield (left *y* axis). (*D*) PC confirms their negative interaction. Color scheme: Blue and green are used for positive causality. Red and yellow are used for negative causality. Purple and gray are used for dark causality.

The specific equities are chosen on the one hand for their popularity and on the other hand because they are usually studied in tandem (26).

Causal interactions among these two assets (Fig. 3*A*) are distinctly positive (Fig. 3*B*) which renders them ideal candidates for pairs trading, but an ill-advised combination for a diversified portfolio strategy. Furthermore, the higher causative force from *MSFT* to *AAPL* suggests that the trading activity of *MSFT* is more influential on *AAPL* than vice versa, a fact that can aid in modeling forecasting.

Conflicting Financial Forces

Next, we apply PC in a classic example of opposing forces in finance (27), that of S&P 500 (as proxy of stock market performance) and US government 10-y bond yield (Fig. 3*C* below). S&P 500 and bond yield data are available from Datastream. The time span is from January 2, 1985 to August 6, 2018.

The results in Fig. 3*D* validate the clasp between S&P 500 and US 10-y bond, as negative causality, the previously assumed norm for decades, however, with a diminishing intensity since the year 2000. Contrary to the common view that government

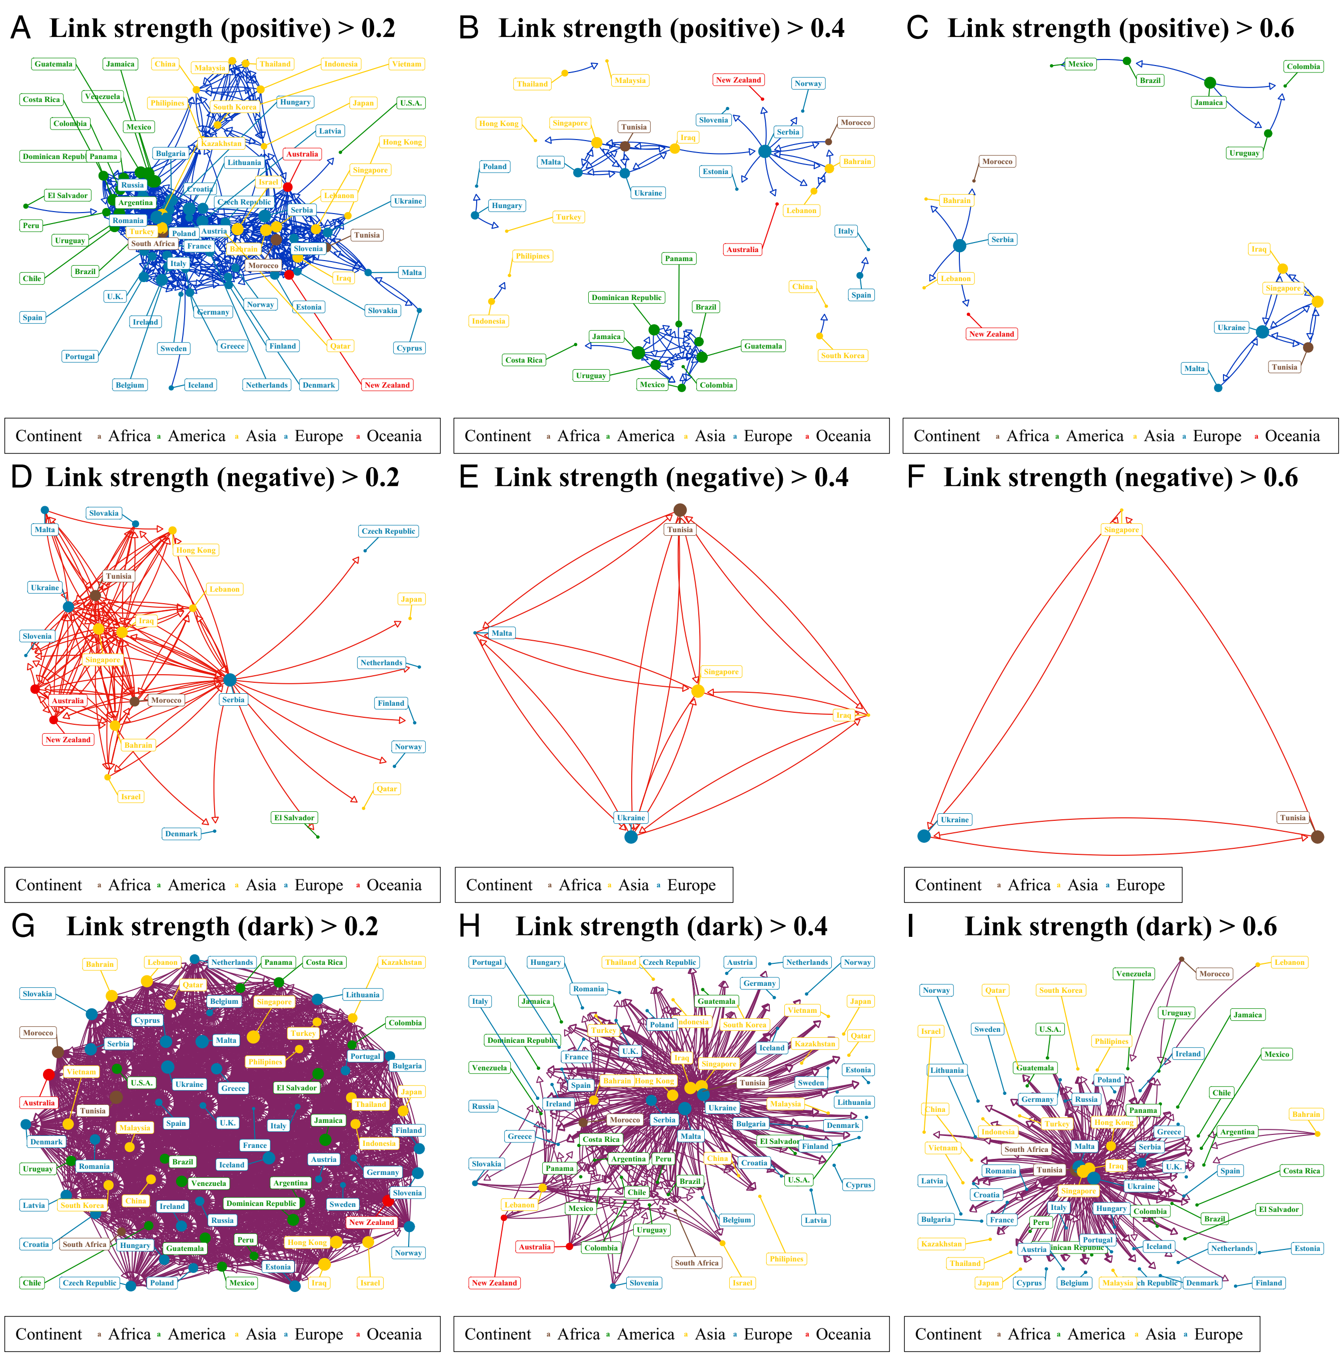

Fig. 4. Nature of causality in CDS network. (*A*–*C*) Positive aspect of the network, gradual elimination of links below 0.2 (*A*), 0.4 (*B*), and 0.6 (*C*) PC. (*D*–*F*) Negative aspect of the network, gradual elimination of links below 0.2 (*D*), 0.4 (*E*), and 0.6 (*F*) PC. (*G*–*I*) Dark aspect of the network, gradual elimination of links below 0.2 (*G*), 0.4 (*H*), and 0.6 (*I*) PC. Overall dark causality is the most persistent type of PC. Link color scheme: Blue is used for positive causality. Red is used for negative causality. Purple is used for dark causality.

bond policy drives the stock market, our method supports the opposite. Policymakers may need to reconsider whether government bond yield leveraging affects the stock market or the assumption that they cause a negative feedback loop.

Dark Causality in Global Sovereign CDS Networks

As a final and more complex system we analyze a dataset consisting of 69 sovereign CDS. The daily time series data were downloaded from Datastream and the time period spans from May 4, 2010 to August 6, 2018. CDS are relatively new derivates and academic research is still characterizing their mathematics and trying to understand their relation to financial crises (28) and determinants to the market (29).

Using PC as link weight we build the three emergent aspects of the CDS network (positive, negative, and dark) (see Fig. 4 below). To get a broad view of the dominant nature of causality in this network we sequentially eliminate the weakest links from 0.2 link strength up to 0.6 by step of 0.2. Negative causality produces the most fragile network and when we eliminate links of up to 0.4 and 0.6 weight, the network is decimated down to three assets. The positive causality aspect of the network is slightly stronger with 15 assets remaining after the final elimination process. On the contrary, the dark aspect of the causality network seems to be quite robust, since even after our final elimination step it remains with 65 assets and the connections remain rather dense.

The asymmetrical domination of dark causality means that the CDS market is strongly interconnected yet the correspondences in the shadow attractors are not clear about similar or opposite temporal patterns. This rise of complex interactions imprinted in the CDS attractors indicates nontrivial dynamics. Without a pure (positive or negative) form of causality the practice of few major dealers concentrating portfolios of large volumes of CDS (29) is not advisable to minimize systemic risk and credit risk exposure. In other words, the preeminence of dark causality can be seen as a sign of "*terra incognita*" (in Latin for unexplored land) in the already notorious CDS market (30).

Concluding Remarks

Greek philosophers such as Plato and Aristotle cerebrated the concept of causality (31, 32). Through this study, we quantify the nature of causality among time series by gauging the correspondence of patterns in contemporaneously embedded neighborhoods and particularly the detection and quantification of dark causality. The more accurate the recalling ability of M_X's patterns about M_Y's patterns, the higher the causality from X to Y is. Whether the nature of causality is positive, negative, or dark depends on the coupling of patterns between M_X and M_Y.

Causal networks are abundant in natural ecosystems, biological processes, and financial markets. More often than not, the mere quantification of causality is not enough. Species interact in complex and varied ways (e.g., symbiosis, competition, or scapegoat relationship). Physiological functions are subject to underlying synergies which are not straightforward. The spectrum of causalities among financial assets is bountiful and their insights would be a great boon for economists and policymakers alike. By unveiling the innermost mechanics of dynamical systems, PC offers an insight into the variety of causal interactions.

Method

We propose an algorithm to analyze the nature of causality by using the theories of symbolic dynamics and attractor reconstruction, both of which refer to time series. According to symbolic dynamics theory, we can represent each time series through patterns that account for one step ahead percentage changes. Symbolic dynamics allow the patterns to expose the nature of causality but first we need to identify a valid causal relationship. To that end we use attractor reconstruction theory. Time series X, Y, and Z that belong to the same dynamical system, e.g., ecosystems, stock markets, or human body, are considered to be parts of the common attractor M, which corresponds to the states of that system. Thus, M evolves in three dimensions where each dimension corresponds to the X, Y, and Z, values of the time series on the right.

From a dynamical systems perspective, the time series X, Y, and Z, are 1D manifestations of the 3D attractor M. Takens' theorem allows us to use lagged values of those 1D manifestations to reconstruct a shadow version of the original attractor by using lagged values of a single time series. M_X being the shadow attractor reconstructed from time series X, is also a library of the past symbolic dynamics of X. The shadow attractor M_X maintains the topology of the original attractor M and is essentially a one-to-one mapping from M to M_X. If we also use time series Y, to reconstruct another shadow version of the original attractor, we know that M_Y maps also one-to-one to M. Therefore, M_X and M_Y map one-to-one between themselves (*SI Appendix*, Eqs. S6–S9). The same can hold also with M_Z.

If a causal relationship exists, then the NNs of each point on M_Y must correspond to NNs of each contemporaneous point on M_X. Revealing the patterns of each NN, we can then calculate the "average pattern" or signature of each neighborhood (*SI Appendix*, Eq. S7). Signature is the weighted average pattern, of a current point's NNs; where the weight for each NN's pattern is higher, the closer it is to the current point. When signatures in M_Y, accurately predict signatures in M_X, then we can say that X left its markings on Y and thus causality is analogous to the prediction accuracy. However, the very correspondence of signatures is what eventually allows us to assess the nature of causality. *SI Appendix*, Tables S1 and S2 describe the three types of causality that encompass all possible causal relationships. Positive causality corresponds to same pattern changes (*SI Appendix*, Eq. S16). Negative causality corresponds to opposite pattern changes (*SI Appendix*, Eq. S17). However, as shown in the matrix there is a third case which cannot be classified as positive or negative. Thus, we use the term dark causality to refer to signature couplings that are complex to interpret (*SI Appendix*, Eq. S18). For the visualization of our method, see Movie S1.

ACKNOWLEDGMENTS. We thank Dimitri K. Pandeleakis for his valuable support in developing the animations for Movie S1, and Nicholas S. Spyrison for his comments that significantly improved the presentation of this paper. S.K.S. and A.A.P. acknowledge the gracious support of this work by the Engineering and Physical Sciences Research Council and Economic and Social Research Council Centre for Doctoral Training on Quantification and Management of Risk and Uncertainty in Complex Systems and Environments (Grant EP/L015927/1).

1. Gai P, Kapadia S (2010) Contagion in financial networks. *Proc R Soc London Ser A* 466: 2401–2423.
2. Preis T, Schneider JJ, Stanley HE (2011) Switching processes in financial markets. *Proc Natl Acad Sci USA* 108:7674–7678.
3. Gabaix X, Gopikrishnan P, Plerou V, Stanley HE (2003) A theory of power-law distributions in financial market fluctuations. *Nature* 423:267–270.
4. Boginski V, Butenko S, Pardalos PM (2005) Statistical analysis of financial networks. *Comput Stat Data Anal* 48:431–443.
5. Preis, T, Reith D, Stanley HE (2010) Complex dynamics of our economic life on different scales: Insights from search engine query data. *Philos Trans R Soc London Ser A* 368:5707–5719.
6. Markowitz H (1952) Portfolio selection. *J Finance* 7:77–91.
7. Yule GU (1926) Why do we sometimes get nonsense-correlations between time-series? A study in sampling and the nature of time-series. *J R Stat Soc* 89:1–63.
8. Pearson K (1895) Notes on regression and inheritance in the case of two parents. *Proc R Soc Lond* 58:240–242.
9. Pearl J (2009) *Causality: Models, Reasoning, and Inference* (Cambridge Univ Press, Cambridge, UK).
10. Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37:424–438.
11. Engle RF, Granger CW (1987) Co-integration and error correction: Representation, estimation, and testing. *Econometrica* 55:251–276.
12. Schreiber T (2000) Measuring information transfer. *Phys Rev Lett* 85:461–464.
13. Sugihara G, et al. (2012) Detecting causality in complex ecosystems. *Science* 338:496–500.
14. Deyle ER, May RM, Munch SB, Sugihara G (2016) Tracking and forecasting ecosystem interactions in real time. *Proc R Soc B* 283:2015–2258.
15. Gatev E, Goetzmann WN, Rouwenhorst KG (2006) Pairs trading: Performance of a relative-value arbitrage rule. *Rev Financ Stud* 19:797–827.
16. Chow G, Jacquier E, Kritzman M, Lowry K (1999) Optimal portfolios in good times and bad. *Financ Anal J* 55:65–73.
17. Morse M, Hedlund GA (1938) Symbolic dynamics. *Am J Math* 60:815–866.
18. Takens F (1981) Detecting strange attractors in turbulence. *Dynamical Systems and Turbulence, Warwick 1980*. Lecture Notes in Mathematics, eds Rand D, Young LS (Springer, Berlin), Vol. 898, pp 366–381.

19. Deyle ER, Sugihara G (2011) Generalized theorems for nonlinear state space reconstruction. *PLoS One* 6:e18295.
20. Crespi BJ (2004) Vicious circles: Positive feedback in major evolutionary and ecological transitions. *Trends Ecol Evol* 19:627–633.
21. Hall JE (2015) *Guyton and Hall Textbook of Medical Physiology* (Elsevier Health Sciences, Amsterdam, The Netherlands).
22. Arthur WB (1990) Positive feedbacks in the economy. *Sci Am* 262:92–99.
23. Neutel AM, et al. (2007) Reconciling complexity with stability in naturally assembling food webs. *Nature* 449:599–602.
24. Raven PH, Johnson GB (1999) *Biology* (McGraw-Hill Companies Inc., Boston).
25. Prechter RR, Jr (2001) Unconscious herding behavior as the psychological basis of financial market trends and patterns. *J Psychol Financial Mark* 2:120–125.
26. Da Z, Engelberg J, Gao P (2011) In search of attention. *J Finance* 66:1461–1499.
27. Kwan SH (1996) Firm-specific information and the correlation between individual stocks and bonds. *J Financ Econ* 40:63–80.
28. Gündüz Y, Kaya O (2014) Impacts of the financial crisis on Eurozone sovereign CDS spreads. *J Int Money Financ* 49:425–442.
29. Peltonen TA, Scheicher M, Vuillemey G (2014) The network structure of the CDS market and its determinants. *J Financ Stab* 13:118–133.
30. Acharya V, Philippon T, Richardson M, Roubini N (2009) The financial crisis of 2007-2009: Causes and remedies. *Financ Mark Inst Instrum* 18:89–137.
31. Hamilton E, Cairns H, Cooper L (1961) *The Collected Dialogues of Plato* (Princeton Univ Press, Princeton).
32. Hennig B (2009) The four causes. *J Philos* 106:137–160.

APPLIED PHYSICAL SCIENCES