A Structural Model of Market Dynamics, and Why it Matters

Jonathan Khedair and Reimer Kühn
Department of Mathematics, King's College London, Strand, London WC2R 2LS, UK
(Dated: March 19, 2019)

In this paper we explore an approach to understanding price fluctuations within a market via considerations of functional dependencies between asset prices. Interestingly, this approach suggests a class of models of a type used earlier to describe the dynamics of real and artificial neural networks. Statistical physics approaches turn out to be suitable for an analysis of their collective properties. In this paper, we first motivate the basic phenomenology and modelling arguments before moving on to discussing some major issues with inference and empirical verification. In particular, we focus on the natural creation of market states through the inclusion of interactions and how these then interfere with inference. This is primarily addressed in a synthetic setting. Finally we investigate real data to test the ability of our approach to capture some key features of the behaviour of financial markets.

PACS numbers: 02.50.r, 05.40.a, 89.65.Gh, 89.75.Da

I. INTRODUCTION

It is well known that the statistical properties of financial markets can be well summarized in terms of the so-called stylised facts; most notable of such are fat tailed return distributions, time dependent volatilities and long term memory in the form of volatility clustering [1]. The seemingly robust universality of these phenomena across various markets and time-scales provide the belief that there may in fact exist an underlying mechanism of market behaviour [2–5].

There have been many attempts in the past to formulate such laws, with efforts largely led by two distinct approaches. The first extends the classical Brownian Motion type models for describing the evolution of prices in a phenomenological manner. In this case, the statistical properties of interest are generated by using fat-tailed random price increments instead of Brownian ones [6], or processes with more complicated noise structures such as jump-diffusion type models [7]. Alternatively, models explicitly constructing a dynamical process for the evolution of volatilities have been proposed. Well known examples include a class of stochastic volatility models known as GARCH models [8] and the Heston model as their continuous time analogue [9].

The second approach on the other hand concentrates on the universal nature of the main stylised facts. In this regard, one appeals to powerful heuristics in statistical physics according to which universal macroscopic properties emerge quite naturally as a result of interactions between constituents in a large system. Prominent modelling examples along this line of reasoning are agent based models such as the Minority Game and related models [10–14]. See [5] for an up to date review of such approaches to finance in general. Whilst attempting to be mechanistic rather than descriptive, models of this type are still problematic in two particular ways. The first being that some of them require fine tuning in order to create the critical conditions needed for the models to exhibit the desired statistical properties one wishes to capture. Secondly, connecting such models to empirical data can be difficult. The Minority Game, for instance, deals with interacting strategies of agents, and their co-evolution in response to common information — a level of description that is difficult to link directly to real price dynamics; an empirical underpinning of that model would therefore be rather non-trivial.

In [15], a model was proposed which can be thought of as a hybrid of the two directions described above. The model recognises the importance of modelling at the level of the prices so as to connect with the vast existing literature on both theoretical and empirical studies of financial markets. It also builds on the idea that it is the collective nature of how fluctuations propagate through a market via mechanistic processes which is mainly responsible for the universal nature of market behaviour. Effective interactions between asset prices are therefore a key ingredient of that modelling approach.

Of course, the notion of dependencies between asset prices is not novel and it is in fact quite common to speak of them within a market. However, in this context it is very important to carefully distinguish between statistical measures of dependency such as correlations, and the functional dependencies we aim to discuss. The former typically feature in data-driven approaches to uncover market structure, whereas the latter are more deeply concerned with market mechanisms.

Measures of correlation have a long history in finance, dating back at least to the early work of Markowitz on portfolio optimization [16]. More recently they have been playing a key role in many studies originating in the Econophysics community. See [17] and references therein for a recent review providing an overview of such data driven studies. Measures of correlation have in particular played a central role in detecting structure in markets, including sectors of co-dependent assets or dynamically distinct regimes of market behaviour [18–27].

While many interesting and useful results have arisen from such studies, one should remember that they are limited to linear statistical dependencies, which in systems with high levels of non-linear feedback may miss important information. Moreover, results obtained are often critically dependent on the time scale of measurements and the window sizes used for estimation. One of the more significant advances in the field of correlation measurements in finance has been the advent of random matrix theory (RMT) in this area, which has provided the promise of truly capturing significant signals through proper filtering of noise, and has thus acted to stabilise available results [28–30].

As a final point we note that uncovering *mechanisms* responsible for the creation of correlations can provide a complementary avenue to quantifying market structures (beyond linear statistical dependencies). Thinking in particular in terms of interactions between asset prices, the wealth of existing literature on correlations provides an immediate starting point for comparison, when analysing and interpreting results. For example, one may look to compare the standard conclusions drawn from clustering or using minimum spanning trees, say, with those obtained when applied to interactions.

The rest of the paper is organised in the following way. In section II we provide an overview of the modelling approach alluded to above. Section III gives a brief account of the solution of a simple variant of a market model. From there on we will in Sect. IV move on to the main focus of the present paper, investigating the role of interactions in and their implications for inference. This will be first be done in a synthetic setting where connections to RMT can be made. Finally we will in Sect. V look at real data and demonstrate that key statistical properties of market dynamics can indeed be reproduced from an inferred model of interacting prices. We end with a summary and discussion in Sect. VI.

II. MODEL PHILOSOPHY

The idea that the universal statistical properties of price fluctuations in financial markets can be understood as a collective effect, for which effective interactions between asset prices could be ultimately responsible is based on the following argument. It starts out by considering a hypothetically *complete theory* of market dynamics. Such a complete theory would describe the co-evolution of *all* relevant degrees of freedom in a market (which would comprise not only traded volumes, the effects of supply and demand, of interest rates, of interacting trading strategies, of liquidity, or of various macro-economic indicators on prices and more, but also the influence of less tangible effects such as psychology of traders, the role of imitation, the effect of (random) news and so on). Suppose that with such a hypothetically complete theory in hand one could write down the complete set of equations of motion governing these degrees of freedom. Taking such a complete theory as a starting point to formulate a model of market dynamics in terms of prices alone would require integrating out all other degrees of freedom from the complete set of dynamical equations. This would lead to a reduced theory which would then *in general* exhibit two key features, viz. (i) effective interactions between prices and (ii) non-Markovian dynamics.

Working with log-prices u_i of a system of N assets labelled $i = 1, \ldots, N$, the above line of reasoning motivates a general class of models with a dynamics of the following form

$$\mathrm{d}u_i(t) = f_i[\boldsymbol{u}]\mathrm{d}t + \sigma_i \mathrm{d}\phi_i(t) \ , \tag{1}$$

in which f_i is a drift functional for the i-th asset, which may depend functionally on the market-history of all other asset prices, collectively denoted by $\boldsymbol{u} = (u_i)$, and $\mathrm{d}\phi_i(t)$ is a random increment which may represent coloured noise. Expanding the drift functional in terms of the order of interaction involved, one could write it quite generally as

$$f_i[\boldsymbol{u}] = f_i^{(0)}(t) + f_i^{(1)}[u_i] + \sum_j f_{ij}^{(2)}[u_i, u_j] + \sum_{jk} f_{ijk}^{(3)}[u_i, u_j, u_k] + \ldots \tag{2}$$

with

$$f_i^{(1)}[u_i] = \int^t \mathrm{d}s\, M_i(t-s) g_i^{(1)}(u_i(s)) \ , \tag{3}$$

$$f_{ij}^{(2)}[u_i, u_j] = \int^t \mathrm{d}s\, M_{ij}(t-s) g_{ij}^{(2)}(u_i(s), u_j(s)) \ , \tag{4}$$

and so on. Here the M_i and the M_{ij} are memory kernels which would encode the precise non-Markovian nature of the dynamics, while the $g_i^{(1)}$, the $g_{ij}^{(2)}$ would be feedback functions, generally expected to be non-linear, encoding details of the interactions. We note in passing that other forms of the interactions may also occur, in which one would first obtain time-averages of log-price trajectories and use non-linear functions of these, rather than first evaluating non-linear functions of log-prices and taking time-averages of trajectories of these as in Eqs. (3), (4).

The above framework is kept fairly general, so as to be able to capture — at least in principle — the majority of mechanisms that might operate in a market, when described in terms of asset prices alone. As such though, it of little practical use and choices will have to be made to make further progress. As is usually the case in physics, one will start by selecting what can be regarded as the simplest model in the above class that is at the same time rich enough to produce the desired phenomenology.

In [15] and in the more recent analytical study [31] the authors restrict themselves to analysing a model with pairwise interactions only, and in the first instance also used a *Markovian approximation* of the dynamics. Drawing inspiration from similar interacting models found in the study of spin glasses, they made the further assumption that the non-linear feedback functions $g_{ij}^{(2)}$ are independent of the current state u_i of asset i on which it acts, by choosing $g_{ij}^{(2)}(u_i(t), u_j(t)) = g_j(u_j(t))$. Under these simplifying assumptions, the drift functional is in fact a function that is local in time, and can be written as

$$f_i(\boldsymbol{u}, t) = f_i^{(0)}(t) + f_i^{(1)}(u_i(t)) + \sum_j J_{ij} g_j(u_j(t)) \ . \quad (5)$$

In [15, 31], the non-linear feedback functions were chosen to be of sigmoid form, i.e. $g_j(u_j) = \tanh(\beta_j u_j)$ or $g_j(u_j) = \mathrm{erf}(\beta_j u_j)$, in which the $\beta_j > 0$ take the role of gain-parameters. The results presented in [31] and in the present paper use error function feedback, as it is the more convenient choice when evaluating integrals over these functions involving Gaussian weights.

With these choices for the drift function, the system described by the evolution equations (1) would in fact describe the evolution of a system of analogue neurons as proposed by Hopfield [32], provided one chooses

$$f_i^{(0)}(t) = I_i(t) \quad \text{and} \quad f_i^{(1)}(u_i) = -\kappa_i u_i \ . \quad (6)$$

In the neural networks context, the u_i would be post-synaptic potentials, the $I_i(t)$ would represent time-dependent external signals, the $\kappa_i > 0$ would denote trans-membrane conductances describing leakage across membranes, and the J_{ij} would represent synaptic couplings, which could be either excitatory ($J_{ij} > 0$) or inhibitory ($J_{ij} < 0$).

In the context of market dynamics investigated in [15, 31], a model with a drift function specified by Eqs. (5) and (6) would constitute an interacting generalization of the classical geometric Brownian motion model, formulated in terms of log-prices $u_i(t) = \log(S_i(t)/S_{i0})$, in which S_{i0} represents a price-scale, the κ_i would represent a mean-reversion effect in the log-price dynamics, which would stabilize the market in the long term, and the J_{ij} would quantify the strength of interactions. With the non-linear feedback functions g_j chosen to be sigmoid, positive couplings, $J_{ij} > 0$, would favour positively correlated asset prices whereas negative couplings, $J_{ij} < 0$, would favour negative correlations. In what follows we will refer to this model as interacting geometric Brownian motion model (iGBM).

The standard geometric Brownian motion model would be recovered for the simple choice $f_i(\boldsymbol{u}, t) = I_i = \mu_i - \frac{1}{2}\sigma_i^2$, with μ_i and σ_i quantifying the drift and volatility in the model. Moreover the random increments $\mathrm{d}\phi_i(t)$ in Eq. (1) would for the standard geometric Brownian motion model chosen to be Brownian (Gaussian white) increments $\mathrm{d}\phi_i(t) = \mathrm{d}W_i(t)$.

In [31], the authors further decompose the state-independent contribution $I_i(t)$ to the drift as

$$I_i(t) = I_i + \sigma_{0i} u_0(t) \ , \quad (7)$$

consisting of a time independent part I_i and and a slow fluctuating component driven by an external process. The authors interpret u_0 as describing a global risk component felt by all assets in the market, which could be interpreted as an indicator of macroscopic economic conditions. They model it as an Ornstein-Uhlenbeck process

$$\mathrm{d}u_0(t) = -\gamma u_0(t)\mathrm{d}t + \sqrt{2\gamma}\, \mathrm{d}W_0(t) \ , \quad (8)$$

assuming $\gamma \ll 1$, such that the u_0 process becomes considerably slower than the microscopic asset dynamics.

III. MODEL SOLUTION AND PHENOMENOLOGY

In [31] a Markovian approximation of a model of interacting prices with drift-function as specified in Eqs. (5) – (8) is studied analytically, using random realizations for (most of) the model parameters, and Gaussian white noise for the random forces in Eq. (1). Specifically, interactions are assumed to be Gaussian and of the form

$$J_{ij} = \frac{J_0}{N} + \frac{J}{\sqrt{N}} x_{ij} \quad (9)$$

with $x_{ij} \sim \mathcal{N}(0,1)$ and $\langle x_{ij} x_{ji} \rangle = \eta$. Here J_0 and J parametrize the mean and standard deviation of the coupling distribution, while η is a measure of forward-backward correlations between interactions of pairs of asset prices.

The all-to-all nature of the couplings allows one, using the generating functional method, to describe the typical dynamical behaviour of the system in terms of a family of non-interacting processes which are self-consistently coupled through a set of dynamical order-parameters — a description that becomes exact in the thermodynamic limit of infinite system size. The effective single-node dynamics is of the form

$$\begin{aligned} \dot{u}_\vartheta(t) = &- \kappa\, u_\vartheta(t) \, + I + \, J_0\, m(t) \, + \sigma_0 u_0(t) \\ &+ \eta\, J^2 \int_0^t \mathrm{d}s\, G(t,\, s)\, n_\vartheta(s) \, + \, \phi_\vartheta(t) \, , \end{aligned} \quad (10)$$

with $\vartheta \equiv (I, \kappa, \sigma, \sigma_0)$ used as a short-hand for the collection of site-random quantities. In Eq. (10) we use $n_\vartheta(s)$ to denote a non-linear function of $u_\vartheta(s)$, i.e. $n_\vartheta(s) = g(u_\vartheta(s))$, and ϕ_ϑ is a Gaussian coloured noise process with

$$\langle \phi_\vartheta(t) \rangle = 0 \ , \quad (11)$$

$$\langle \phi_\vartheta(t) \phi_{\vartheta'}(s) \rangle = \delta_{\vartheta,\vartheta'} [\sigma^2 \delta(t-s) + J^2 q(t,s)] \ . \quad (12)$$

The order parameters $m(t)$, $G(t,s)$ $q(t,s)$ appearing in the equation of motion (10) and in the specification of the noise statistics must be determined self-consistently to satisfy

$$m(t) = \left\langle \langle n_\vartheta(t) \rangle_{\phi_\vartheta} \right\rangle_\vartheta \ , \quad (13)$$

$$q(t,s) = \left\langle \langle n_\vartheta(t) n_\vartheta(s) \rangle_{\phi_\vartheta} \right\rangle_\vartheta \ , \quad (14)$$

$$G(t,s) = \left\langle \frac{\delta \langle n_\vartheta(t) \rangle_{\phi_\vartheta}}{\delta \phi_\vartheta(s)} \right\rangle_\vartheta \ , \quad t > s \ . \quad (15)$$

Here inner averages are averages over the coloured noise ϕ_ϑ while outer averages are over distributions of site-random parameters.

In [31] the model is solved using separation of time scales argument, assuming that the system becomes statistically stationary at given values of the slowly evolving external u_0 process. This allows the authors to evaluate, e.g. distributions of log-returns for various time-scales. For a broad class of distributions of mean-reversion parameters κ_i they are indeed found to be fat tailed in a way broadly in line with empirical facts. We refer to [31] for details and further results.

In what follows, though, we wish to highlight the fact that — with interactions of the form (9) — one expects the model to exhibit glassy properties characterized by a large number of meta-stable states, if the variance of the coupling distribution is sufficiently large compared to its mean. These meta-stable states provide a simple and transparent mechanism for intermittent market dynamics and volatility clustering through the interplay of dynamics "inside states" (with different degrees of stability, hence different degrees of effective volatility in each state), and occasional transitions between them. In the present set-up the slow OU-process (8) was found to be the key driver for transitions between meta-stable states, hence the rate of such transitions is controlled by the rate γ of of that process.

To test the hypothesis of a link between volatility clustering and meta-stable state structure, a numerical experiment was performed in [31], in which the Gaussian couplings of Eq. (9) were modified by the addition of Hebb-Hopfield couplings which stabilize a set of attractors aligned with P unbiased binary random patterns $\{\xi_i^\mu\}$ as follows

$$J_{ij} \leftarrow J_{ij} + \frac{1}{N}\sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu \ . \tag{16}$$

The ξ_i^μ are chosen from $\{\pm 1\}$, independently and with equal probability for $i = 1, \ldots, N$ and $\mu = 1, \ldots P$. The choice of the modification is inspired by models of neural networks, where specific patterns are encoded in couplings of the type added. The combination of Gaussian and Hebb-Hopfield couplings is expected to stabilize slightly distorted versions of the $\{\xi_i^\mu\}$ as attractors.

In a simulation the time evolution of returns $\Delta I(t) = I(t+\tau) - I(t)$ of the market-index

$$I(t) = \frac{1}{N}\sum_i \mathrm{e}^{u_i(t)} \tag{17}$$

was monitored alongside the evolution of overlaps

$$m_\mu(t) = \frac{1}{N}\sum_i \xi_i^\mu\, g(u_i(t)) \ . \tag{18}$$

Figure 1 clearly demonstrates a strong correlation between the volatility of index returns and the evolution of the overlaps, which measure the similarity of the system states with the three attractors embedded in the couplings.

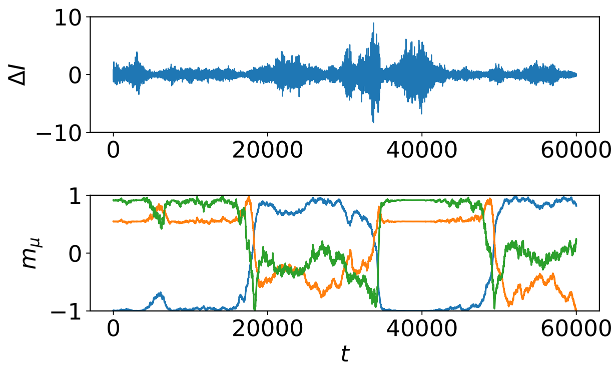

FIG. 1. The top panel shows the index returns ΔI over time intervals $\tau = 10$ for a system of size $N = 50$. The bottom panel shows the corresponding overlaps of the system state with each of the $P = 3$ embedded attractors. Other system parameters are $J_0 = J = 0.5 = \eta = 0.5$, $I_i \sim \mathcal{N}(0, 0.5)$, $\sigma_i \equiv 0.1$, and $\gamma = 10^{-4}$. The κ_i are exponentially distributed with mean $\kappa_0 = 0.2$

Given the importance of the underlying interaction network in creating such effects, practical applications of the model would of course be tackled by first attempting to infer this underlying structure. Once this had been achieved, predictive abilities of the model could be tested and its usefulness in terms of capturing the collective behaviour of asset dynamics could be analysed.

IV. INFERENCE

Having motivated the need for interactions in modelling the dynamics of asset prices, we now concern ourselves with the *inverse* problem of inferring the parameters of such models. We will do this in a parametric framework, using the maximum likelihood method to infer parameters of an iGBM specified by Eqs. (5)-(7) from observed trajectories of log-returns of a set of (interacting) assets.

Before turning to actual market data though, we will attempt to infer model parameters using synthetic data generated by a model in the very same class, for which results can thus be compared with a known ground truth. This is an important exercise that will allow us to understand the limitations of inference in systems of the type considered here. As we will discuss in detail below, two main issues are encountered.

The first is related to sampling problems. In order to fully specify an iGBM, one has to fix a large number of parameters. This number scales as N^2 for a system of size N. In order to reliably infer the model parameters, sufficiently large data samples must be taken. We will find below that the required sample sizes are intimately related to those required to reliably estimate elements of large random correlation matrices.

The second issue is a direct consequence of the existence of metastable states in the dynamics of iGBM type models. Given that a system may spend large amounts of time in any of the meta-stable states, even very long dynamical trajectories may only explore a small fraction of the complete set of states that a system may in principle support. In such a case it is conceivable that only some low-dimensional projection of the coupling matrix can be reliably estimated despite the fact that the amount of data used for estimation is very large indeed. We will demonstrate below in a synthetic setting, that this hypothesis is actually correct. Analogous effects have indeed been noted when inferring interactions in glassy Ising type models in the low temperature regime [33–35]. Assuming then that real market dynamics may be beset by the same type of problems, one will have to be cautious when drawing conclusions from inferred model parameters, and ideally devise suitable measures which could indicate whether an inference attempt has or has not overcome the problem associated with slow exploration of a sufficiently large number of long-lived states.

With reliably estimated system parameters in hand, one could proceed to investigate structural properties of inferred interactions which could reveal market structures and dependencies beyond those detectable using correlation measures. Moreover, one could of course go on to generate market trajectories via Monte Carlo simulations, using these, e.g., as inputs for risk analysis, or to identify specific market states. The list could go on.

A. Maximum Likelihood

The starting point for our analysis is the likelihood which we write as the probability of a trajectory of observed prices, after discretizing time in steps of size Δ as

$$P[\boldsymbol{u}] = \prod_{i,t} \frac{1}{\sqrt{2\pi\sigma_i^2/\Delta}} \exp\Big(- \frac{\Delta}{2\sigma_i^2} \big[\dot{u}_i(t) - f_i(\boldsymbol{u},t)\big]^2 \Big) ,$$

with $\dot{u}$ a short-hand for the discrete time derivative $\dot{u}(t) = [u(t+\Delta) - u(t)]/\Delta$. As is usually the case, we work with the negative log likelihood $\mathcal{L}$, which one can write as

$$\mathcal{L} = \sum_{i,t} \frac{\big[\dot{u}_i(t) - f_i(\boldsymbol{u},t)\big]^2}{2\sigma_i^2/\Delta} \tag{19}$$

up to an additive constant independent of the couplings. Estimates of parameters $\vartheta = \{\kappa_i, J_{ij}, I_i, \sigma_{0i}\}$ are then found as usual by solving $\nabla_\vartheta \mathcal{L} = 0$. This equation is easily seen to generate a system of linear equations for the parameters to be estimated. The system of equations actually decomposes into sets of equations which are decouple for different i. We note in passing that estimates of ϑ can be found without knowledge of the σ_i, and that estimates of the latter can be obtained by estimating the variance of $\dot{u}_i(t) - f_i(\boldsymbol{u},t)$.

In practical applications, one would like to estimate the hidden u_0 process as well, which due to the linear Gaussian structure can be done in closed form too. In general, however, one may expect that there may be more than one latent process relevant for asset dynamics. In what follows, we will assume that there is just a single such process, and in order not to overburden the following discussion, we will also assume it (and the σ_{0i}) to be known.

To formulate the equations for the parameters associated with node i we introduce the vector $\mathbf{\Phi}^{(i)}$ of parameters

$$\mathbf{\Phi}^{(i)} = \begin{pmatrix} \kappa_i \\ I_i \\ J_{i1} \\ \vdots \\ J_{iN} \end{pmatrix} = \begin{pmatrix} \kappa_i \\ I_i \\ \mathbf{J}_i \end{pmatrix} , \tag{20}$$

The maximum likelihood equation $\nabla_\vartheta \mathcal{L} = 0$ then stipulates that $\mathbf{\Phi}^{(i)}$ satisfies the equation

$$\mathbf{A}^{(i)} \mathbf{\Phi}^{(i)} = \mathbf{b}^{(i)} , \tag{21}$$

where $\mathbf{b}^{(i)}$ is the vector $2+N$ dimensional vector

$$\mathbf{b}^{(i)} = \left(\begin{array}{c} \sigma_{0i}\langle u_i u_0\rangle - \langle u_i \dot{u}_i\rangle \\ \sigma_{0i}\langle u_0\rangle - \langle \dot{u}_i\rangle \\ \hline \sigma_{0i}\langle g_1 u_0\rangle - \langle g_1 \dot{u}_i\rangle \\ \vdots \\ \sigma_{0i}\langle g_N u_0\rangle - \langle g_N \dot{u}_i\rangle \end{array}\right) \tag{22}$$

and $\mathbf{A}^{(i)}$ is a $(2+N)\times(2+N)$ block-matrix of the form

$$\mathbf{A}^{(i)} = \left(\begin{array}{cc|c} \langle u_i^2\rangle & -\langle u_i\rangle & -[\langle u_i g_k\rangle]_k \\ \langle u_i\rangle & -1 & -[\langle g_k\rangle]_k \\ \hline [\langle g_j u_i\rangle]_j & -[\langle g_j\rangle]_j & -[\langle g_j g_k\rangle]_{jk} \end{array}\right) , \tag{23}$$

in which $\langle \dots \rangle = \frac{1}{T}\sum_{t=1}^{T}(\dots)$ is shorthand for a sample average taken over a sample of T time points. That is, $\mathbf{A}^{(i)}$ is a block matrix with matrix elements determined by various sample correlations and sample averages.

The values of the unknown parameters are then given by

$$\Phi^{(i)} = (\mathbf{A}^{(i)})^{-1}\mathbf{b}^{(i)} , \tag{24}$$

It is convenient to write the $2+N$ dimensional vector $\mathbf{b}^{(i)}$ in the form

$$\mathbf{b}^{(i)} = \begin{pmatrix} \mathbf{b}_0^{(i)} \\ \mathbf{b}_1^{(i)} \end{pmatrix} \tag{25}$$

and, in a similar vein the block-matrix $\mathbf{A}^{(i)}$ as

$$\mathbf{A}^{(i)} = \begin{pmatrix} \mathbf{A}_0^{(i)} & \mathbf{A}_1^{(i)} \\ \mathbf{A}_2^{(i)} & \mathbf{A}_3^{(i)} \end{pmatrix} .$$

Note that $\mathbf{A}_3^{(i)}$ is actually *independent* of i, $\mathbf{A}_3^{(i)} = \mathbf{A}_3$. Using this representation, one can use standard block-matrix inversion formulae to obtain, e.g.,

$$\mathbf{J}_i = \hat{\mathbf{A}}_2^{(i)}\mathbf{b}_0^{(i)} + \hat{\mathbf{A}}_3^{(i)}\mathbf{b}_1^{(i)}, \tag{26}$$

with

$$\hat{\mathbf{A}}_2^{(i)} = \left(\mathbf{A}_1^{(i)} - \mathbf{A}_0^{(i)}(A_2^{(i)})^{-1}\mathbf{A}_3\right)^{-1}, \tag{27}$$

$$\hat{\mathbf{A}}_3^{(i)} = \left(\mathbf{A}_3 - \mathbf{A}_2^{(i)}(A_0^{(i)})^{-1}\mathbf{A}_1^{(i)}\right)^{-1}. \tag{28}$$

Estimating the couplings therefore translates to estimating the block correlation matrices appearing in $\hat{\mathbf{A}}_2^{(i)}$ and $\hat{\mathbf{A}}_3^{(i)}$.

B. Transition to Learning

Given the set of maximum likelihood (ML) equations (26)-(28) for the parameters of interest, one is in principle able to recover these with arbitrary precision, given enough data. In practice only finite data sets will be available for estimation, and the question of the quality of estimates naturally arises. This can of course only be properly tested for in a synthetic case where the parameters of the process generating the data are known.

The general analysis of ML inference presented above entails that inference of the couplings requires us to have good estimates for several correlation matrices. The most problematic of these turns out to be the $N \times N$ submatrix $\mathbf{A}_3$ defined above, with elements

$$(\mathbf{A}_3)_{jk} = \langle g(u_j)g(u_k)\rangle \,. \tag{29}$$

If interactions and consequently correlations are to exist, then estimating the shared variation across assets relies on correct inference of the principle components of this correlation matrix. It is well known that unless the ratio $\alpha = N/T$ of matrix dimension N and sample size T used to measure it is sufficiently small, the estimated correlation matrix will be indistinguishable from a sample correlation matrix of a collection of i.i.d. random vectors, and as a consequence true couplings can not be inferred at all. The limiting distribution for the spectral density of such sample correlation matrices is the Marčenko-Pastur (MP) distribution, and it can be used as a null model, to assess whether principle components of $\mathbf{A}_3$ corresponding to collective modes of the underlying dynamical system have been identified.

We do indeed observe such a transition to learning in the present case. In order to make contact with the standard theory which stipulates a comparison with independent identically distributed random vectors, we restrict ourselves to observing the system in a *single* ergodic component so as to guarantee homogeneous statistics. This can be achieved by turning off the driving force that normally triggers transitions between ergodic components, i.e. by keeping the the global macro-economic noise-component u_0 constant.

We report results in Fig. 2. Rather than using matrix inversions as described in the previous section, though, we use an iterative stochastic gradient descent method on the basis of a data sample of size $T = N/\alpha$ to learn the couplings, using mini-batches of fixed size to evaluate gradients. Measuring estimation errors as after convergence as functions of α, we find that unless $\alpha < \alpha_c \simeq 0.06$, i.e. sufficiently many observations of the system are made to estimate correlation matrices, the inferred interactions are essentially independent of the true couplings. In this case, one is not able to recover any of the collective statistical features of the data set. As α is decreased below α_c, correlations between inferred and true couplings gradually emerge, as do deviations between the MP distribution and the observed eigenvalue density of centred versions of correlation matrices $\mathbf{A}_3$. The same transition is noted in the learning of interaction networks in Ising models[36].

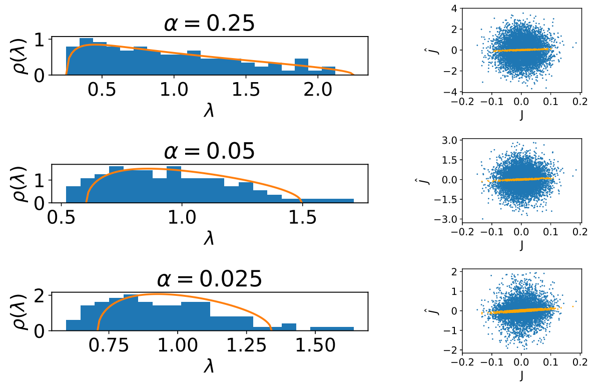

FIG. 2. Left panels: Eigenvalue distributions of centred versions of correlation matrices $\mathbf{A}_3$, compared to MP distributions (full lines) for several values of $\alpha = N/T$ in a system of size $N = 150$. Right panels: Scatter-plots showing estimated couplings $\hat{J}_{ij}$ as functions of the true couplings for the same values of α. Estimated couplings within a (relative) error margin of $\epsilon = 0.01$ of true couplings are marked in a lighter (orange) shade. As α decreases deviations of eigenvalue spectra from the MP law become noticeable, and we observe a clear separation of the leading eigenvalues from the bulk of the spectra. Noting the different vertical scales in the right set of panels, we conclude that at the same time, correlations between estimated and true couplings become more pronounced (i.e., the fraction of estimated couplings within the error margin of $\epsilon = 0.01$ of the true couplings increases) with decreasing α.

C. The Curse of Non-ergodicity

With the dynamics stuck in a single ergodic component, couplings will, however never be fully learnt, independently of the size of the data set used to estimate pertinent correlation matrices, as only a low dimensional projection of the full coupling matrix will be relevant for the dynamics in any single such ergodic component. Only that same low dimensional projection can as a result be inferred.

In order to be able to fully learn the true couplings, the system dynamics must therefore be able to explore sufficiently many of the long-lived states that the system is in principle able to support, in order to estimate the couplings with reasonable precision.

In the synthetic case, the rate at which transitions between long-lived states occur is controlled by γ, the rate of change of the slow u_0 process introduced to mimic slowly evolving macro-economic conditions. By tuning γ one can therefore influence the degree to which the full state space will be explored in a given fixed time horizon, and thus the precision with which the true couplings can be inferred.

Figures 3 and 4 illustrate this phenomenology in a qualitative fashion. In both cases we use a time horizon of $T = 10^5$ for the dynamics of a system of $N = 150$ assets. In Fig. 3 the parameter γ is chosen to be very small, so that typically only a single ergodic component can be explored. As a consequence only a small subset of couplings is inferred, i.e. within a given error margin ϵ of the true couplings. The information in the left panel is only available for synthetic systems, for which the true couplings are known. The data format in the right panel on the other hand — inferred couplings versus random initial conditions for them — is available even when true couplings are not known. The results show that the subset of couplings which are inferred with reasonable precision turns out to be largely independent of the initial couplings. In Fig. 4 the experiment is repeated with a larger value of γ chosen such that over the same time horizon, the system will typically make $\mathcal{O}(1000)$ transitions between different long-lived states of the system. As a consequence, the overwhelming majority of inferred couplings is now within the error margin $\epsilon = 0.01$ of the true couplings, and is found to be independent of initial values used for inference.

In Fig. 5 we explore the learning process in a more quantitative fashion by showing normalized learning curves for various $\gamma's$. The error is calculated as

$$E(\tau) = \frac{||\hat{\mathbf{J}}(\tau) - \mathbf{J}||}{||\mathbf{J}||} \tag{30}$$

where τ is defined in terms of the number of gradient descent updates. For the figure we plot $\hat{E}(\tau) = E(\tau)/E(0)$ as this eliminates the fluctuations of the initial normalized error across initial conditions. The dependence of the relative asymptotic error on the ergodic time scales is clearly visible, with large residual errors observed for cases for which system exploration is not possible.

While we have here looked at the problem of inference for a specific structural model of market dynamics, the lessons learnt would be relevant for inference of system parameters from observed dynamical trajectories for *any* complex system exhibiting a multiplicity of meta-stable states.

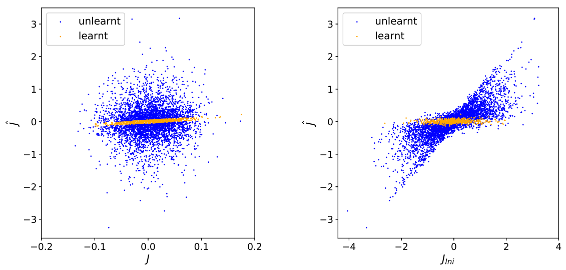

FIG. 3. Left panel: Scatter-plot of inferred couplings versus true couplings for a system of size $N = 150$, with $\gamma = 10^{-7}$, $T = 10^5$ and $N = 150$. Inferred couplings within an error margin of $\epsilon = 0.01$ of the true couplings are displayed in lighter (orange) shade. In the right panel we display results of the same experiment in a different fashion, namely as a scatter-plot of inferred couplings vs. initial choices for the same couplings. In this representation, points on the diagonal have not moved in the course of attempting to infer the couplings.

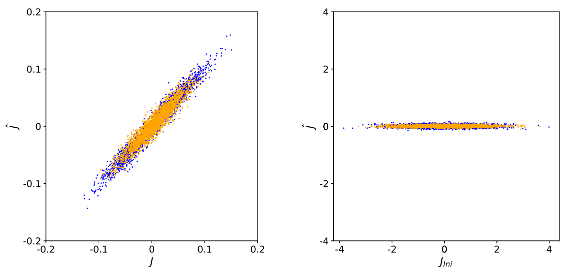

FIG. 4. The same as figure 3 but this time for $\gamma = 0.01$. The ability to explore many ergodic components means couplings can be well estimated.

This is, in particular, also the case for inference of statistical properties of real financial markets, if indeed our hypothesis is correct that salient features of market dynamics are due to the existence of a multiplicity of market states.

V. REAL DATA

Having explored the consequences of market states and interactions in a synthetic setting, we now turn our attention to inference in a real market. As our starting point, we use S&P500 constituent stock prices recorded every minute. In particular, we will look at the $N = 200$ assets which remain continuously in the top 500 ranked by market capitalization as a proxy for our market, and we sample prices every 5 minutes to smooth micro-structure noise.

Two pre-processing steps are employed before data are used to infer parameters of an iGBM model. First, prices are transformed to a co-moving frame — i.e. log-prices are de-trended over time windows of a given length. Sec-

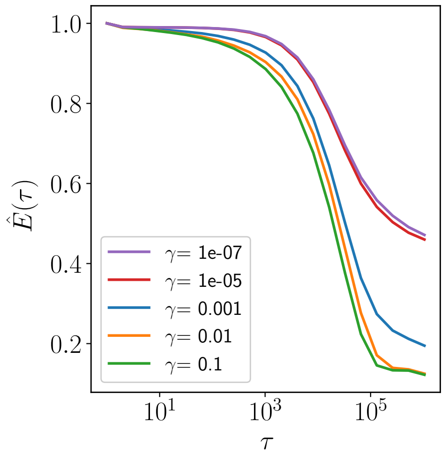

FIG. 5. Average normalized error of inferred couplings as a function of the number τ of updates of a stochastic gradient descent dynamics for various values of γ. Averages are taken over 200 runs. We use mini-batches of size 100 to evaluate gradients and a fixed learning rate $\eta = 0.5$. Curves from top to bottom correspond to increasing γ, thus increasing rate of transitions between long-lived states.

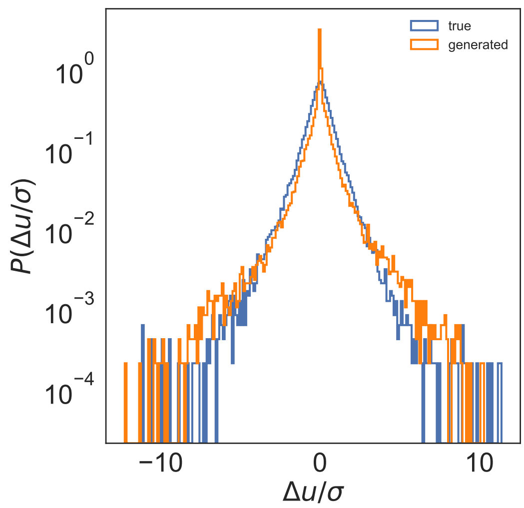

FIG. 6. True and generated ensemble distribution of normalized log-returns (i.e., log-returns measured in units of their variance) for $N = 200$ assets selected from the S&P500.

ondly de-trended log-prices are translated and scaled to fit into a standard interval, which is equivalent to introducing gain parameters and thresholds in the non-linear feedback functions. An iGBM type model is then fitted to the evolution of these de-trended and normalized data. After fitting the model, the pre-processing steps can be reversed to interpret whatever analysis can be extracted from predictions of the inferred model.

Our analysis of real data is still at its beginnings. Here we will mainly report first results concerning collective properties such as return distributions or cross correlations between log-returns as generated by the calibrated model. Connecting with our earlier discussions of market states and their relation to volatility clustering or market instabilities, we will also be interested in how such states or instabilities might be detected in a calibrated iGBM model.

The first requirement for any structural model of financial fluctuations is that it adequately captures the statistical properties of these fluctuations. Beyond the scientific perspective of thereby supporting the modelling approach itself, this is also a formal requirement of any model under the Basel III accords [37] for it to be used in the banking industry.

In Figure 6 we show a comparison between true and generated ensemble 5-minute normalized log-return distributions for a 4 month period in 2017, with the model calibrated on data taken from a 6 month period prior to the period for which the comparison is made. We note that we have so-far only used Gaussian white noise in simulations of the inferred model, whereas the original data clearly show jumps. The latter are detected during the inference stage and interaction parameters are inferred only on the continuous parts of system trajectories. Yet it seems that even without jumps in the generative model we are still able to capture the shape of the fat tailed return distributions, though estimates of the variance are found to be off by a factor of 5 in this case. We expect that this discrepancy can be overcome by estimating jump-size distributions and frequencies and using a proper jump diffusion model to generate market trajectories from an inferred model.

One can take a closer look also at the structure of re-

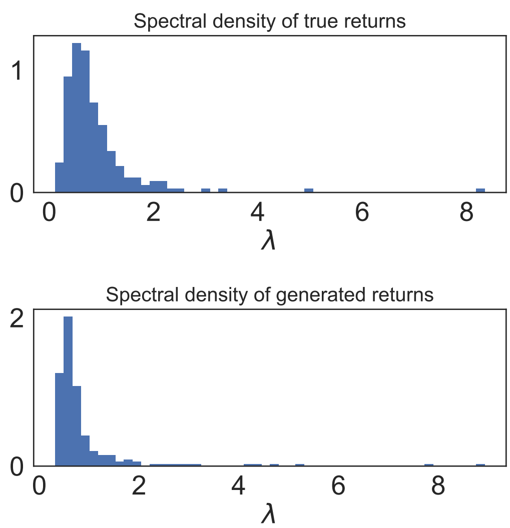

FIG. 7. Spectral density for true and generated return correlations corresponding to same time period as in figure 6.

turns across the market, by looking at cross correlations between them. It is well known, that these can be used e.g. to identify industrial sectors in a market in terms of the dominant principle components of the correlation matrix of returns [38–40]. In Fig. 7 we only report the eigenvalue density of the correlation matrices, comparing results obtained directly from the data with those obtained by generating the same amount of data using the inferred model.

It appears that broad features of the empirical spectrum are recovered by the inferred model. Looking at eigenvalues corresponding to the largest principle components, we note that, while they are approximately of the right magnitude, some additional substructure appears in the synthetically generated spectrum. There is in both spectra a band of small eigenvalues, though in the empirical market data the width of of this band is somewhat larger. Whether the discrepancies are related to the absence of jumps in the generative inferred model is currently under investigation.

Finally, to explore the idea of market states and their relation to the evolution of market volatilities we perform an experiment similar in spirit to that reported in Fig. 1, in which overlaps of system states with attractors embedded in the network are reported. In a real market, we do not have a-priori knowledge of the nature of attractors. It turns out, however, that a singular value decomposition of the inferred couplings provides a decomposition of the contribution of interactions to the drift function of a similar structure as that generated by Hebb-Hopfield type couplings. The singular value decomposition of the inferred coupling is of the form

$$\hat{\mathbf{J}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \tag{31}$$

with orthogonal matrices $\mathbf{U} = (U_{i\mu})$ and $V = (V_{i\mu})$, and the diagonal matrix $\mathbf{\Sigma} = \mathrm{diag}(\sigma_\mu)$ of singular values of $\hat{\mathbf{J}}$.

In terms of this decomposition, one can write the interaction contribution to the drift function as

$$\sum_j \hat{J}_{ij} g(u_j(t)) = \sum_\mu \sigma_\mu U_{i\mu} m_\mu(t) \tag{32}$$

in which overlaps $m_\mu(t)$ with singular vectors now given by

$$m_\mu(t) = \sum_j V_{\mu j} g(u_j(t)) \ . \tag{33}$$

We show in Fig. 8 that this essentially provides a low dimensional visualisation of how the market evolves in terms of the dominant regime changes. The figure covers a time period of 5 years, from early 2012 to late 2016. The top panel shows evolutions of overlaps of market data with three selected singular vectors of the interaction matrix inferred on the basis of data of the year preceding the period shown. The period shown exhibits two phases characterized by comparatively large volatility and overlaps changing significantly, the first marking the height of the sovereign debt crisis, the second beginning around June 2015. One is tempted to speculate whether the transition to the less volatile intermediate phase might have been triggered by by Draghi's famous 'whatever it takes' speech of 26/07/2012, which is widely believed to have contributed to calming down financial markets. The second period of notice contains the aftermath of a major flash crash of 24/08/2015 which is accompanied by a large reconfiguration of the internal dynamics and an accompanying volatility burst. Interestingly, in terms of overlap dynamics, it appears that precursors are visible already two months before the actual flash crash.

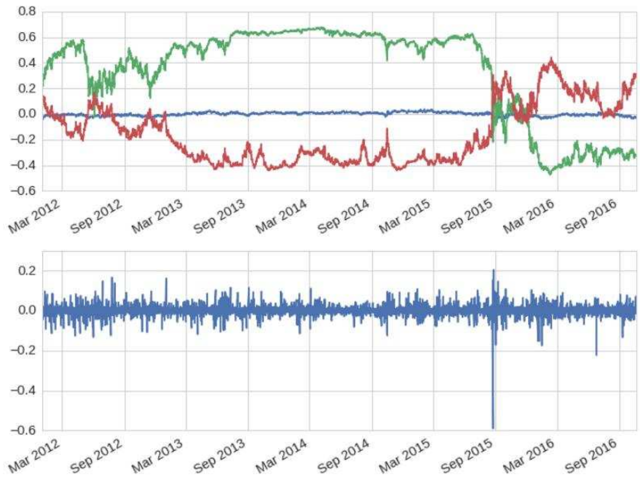

FIG. 8. The top panel shows the projection of the system state onto three selected singular vectors of the inferred interaction matrix. The overlaps are bounded by ± 1, with specific market states typically defined by finite overlap with only a few such directions. All others are of order $\sim 1/\sqrt{N}$ as is illustrated by the middle relatively flat line representing one of these. The lower panel shows the corresponding evolution of the index returns for the same time period.

While the analysis of S&P500 data in terms of an iGBM model has produced evidence that such a model of interacting asset prices can go some way in reproducing statistical properties of the real market dynamics from which it was inferred, results are still far from perfect, and further analysis is clearly required. We expect that some of the inadequacies encountered can be overcome by properly taking effects of jumps into account.

VI. SUMMARY AND DISCUSSION

The main purpose of this paper was to motivate the need for including effective interactions between asset prices when talking about their co-evolution. We introduced a general class of models which would incorporate this idea, before specializing to one of the simplest representatives in this model class to be studied in greater detail. We then discussed the principal difficulties one will encounter when attempting to calibrate a model in this class.

The first problem is directly related to sampling issues for large random matrices, most notably when estimating large correlation matrices. The second issue is related to the possible existence of a large number distinct long-lived market states. Within an iGBM model they are created by random interactions which produce 'conflicting ordering instructions' for asset prices in a manner known from spin-glasses. We show that due to the existence of these long-lived states, large systematic errors are unavoidable when inferring interactions based on data sampled only across a a few of these market states.

One can fairly argue that the analysis of these two principal obstacles for model calibration can hold important lessons for other data driven approaches to analysing financial markets.

Assuming that proper calibration *can* be achieved, the most obvious question to ask is whether the inferred model is sufficient to capture essential statistical properties of asset dynamics of real markets. While we were able to reproduce qualitative aspects of collective behaviour of a market such as fat tailed return-distributions and spectra of correlation matrices generated by a calibrated model which do show the existence a set of collective modes broadly in line with what is seen in real data, quantitative shortcomings remained. We believe that these are mainly due to the fact that we have so far not included jump processes in our modelling. This point clearly deserves further investigation.

The possible existence of market states is clearly relevant from a practical point of view. Apart from their role in creating obstacles for proper calibration mentioned above, detecting and possibly predicting regime changes accompanied by phases of increased market volatility could be useful both for traders and for regulators and central banks. In Sect. V we studied a very simple way based on projecting market states on singular vectors of an inferred interaction matrix, which appears to be able to detect such regime changes, which are typically accompanied by changes in market volatility.

Future directions of our work would include (i) a more thorough study of market states and looking at further ways to detect and possibly to predict them; this would in particular require to create statistical measures of predictability that go beyond the anecdotal evidence presented in Fig. 8, (ii) systematically including jump processes in the model, (iii) exploring whether memory and perhaps higher order interactions, as stipulated by the general line of reasoning presented in Sect. II, need to be taken into account, (iv) exploring the need for a richer multi-factor model of external slow noise sources, and finally (v) assessing whether and how our modelling approach might be used in practice for risk management purposes. The list could possibly go on.

Acknowledgements: J.K. is supported by the EPSRC Centre for Doctoral Training "Cross-Disciplinary Approaches to Non-Equilibrium Systems" (CANES, EP/L015854/1).

[1] R. Cont. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 1:223–236, 2001.

[2] J. P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management.* Cambridge University Press, Cambridge, 2006.

[3] R. N. Mantegna and H. E. Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance.* Cambridge University Press, New York, NY, USA, 2000.

[4] J. Feigenbaum. Financial Physics. *Rep. Progr. Phys.*, 66(10):1611, 2003.

[5] D. Sornette. Physics and Financial Economics (1776–2014): Puzzles, Ising and Agent-Based Models. *Rep. Progr. Phys.*, 77:062001, 2014.

[6] R. N. Mantegna and H. E. Stanley. Stochastic Process with Ultraslow Convergence to a Gaussian: The Truncated Lévy Flight. *Phys. Rev. Lett.*, 73:2946–2949, 1994.

[7] R. Cont and P. Tankov. *Financial Modelling with Jump Processes.* CRC PressINC, 2003.

[8] T Bollersev. Generalized Autoregressive Conditional Heteroscedasticity. *J. Econometrics*, 31:307–327, 1986.

[9] S. L. Heston. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Rev. Fin. Stud.*, 6(2):327–43, 1993.

[10] D. Stauffer and D. Sornette. Self-Organized Percolation Model for Stock Market Fluctuations. *Physica A*, 271:496–506, 1999.

[11] R. Cont and J.-P. Bouchaud. Herd Behavior and Aggregate Fluctuations in Financial Markets. *Macroec. Dyn.*, 4:170–196, 2000.

[12] S. Bornholdt. Expectation Bubbles in a Spin Model of Market Intermittency from Frustration Across Scales. *Int. J. Mod. Phys. C*, 12:667–674, 2001.

[13] D. Challet, M. Marsili, and Y.-C. Zhang. *Minority Games: Interacting Agents in Financial Markets.* Oxford University Press, 2004.

[14] A. C. C. Coolen. *The Mathematical Theory of Minority Games–Statistical Mechanics of Interacting Agents.* Oxford University Press, Oxford, 2005.

[15] R. Kühn and P. Neu. Intermittency in a Minimal Interacting Generalization of the Geometric Brownian Motion Model. *J. Phys. A*, 41:324015 (12pp), 2008.

[16] H. M. Markowitz. Portfolio Selection. *The Journal of Finance*, 7:77–91, 1952.

[17] G. Marti, F. Nielsen, M. Binkowski, and P. Donnat. A Review of Two Decades of Correlations, Hierarchies, Networks and Clustering in Financial Markets. *arXiv preprint arXiv:1703.00485.*, 2018.

[18] M. C. Münnix, T. Shimada, R. Schäfer, F. Leyvraz, Th. H. Seligman, Th. Guhr, and H. E. Stanley. Identifying States of a Financial Market. *Sci. Rep.*, 2:00644, 2012.

[19] M. Marsili. Dissecting Financial Markets: Sectors and

States. *Quant. Fin.*, 2:297–304, 2002.

[20] R. N. Mantegna. Hierarchical Structure in Financial Markets. *Eur. Phys. J. B*, 11:193–197, 1999.

[21] M. Tumminello, F. Lillo, and R. N. Mantegna. Correlation, Hierarchies, and Networks in Financial Markets. *J. Econ. Beh. and Org.*, 75:40 – 58, 2010. Transdisciplinary Perspectives on Economic Complexity.

[22] D.-M. Song, M. Tumminello, W.-X. Zhou, and R. N. Mantegna. Evolution of Worldwide Stock Markets, Correlation Structure, and Correlation-Based Graphs. *Phys. Rev. E*, 84:026108, 2011.

[23] M. MacMahon and D. Garlaschelli. Community Detection for Correlation Matrices. *Phys. Rev. X*, 5:021006, Apr 2015.

[24] A. Kocheturov, M. Batsyn, and P. M. Pardalos. Dynamics of Cluster Structures in a Financial Market Network. *Physica A*, 413:523 – 533, 2014.

[25] D. Y. Kenett, Y. Shapira, A. Madi, S. Bransburg-Zabary, G. Gur-Gershgoren, and E. Ben-Jacob. Dynamics of Stock Market Correlations. *AUCO Czech Economic Review*, 4:330–341, 2010.

[26] A. Sensoy and B. M. Tabak. Dynamic Spanning Trees in Stock Market Networks: The Case of Asia-Pacific. *Physica A*, 414:387 – 402, 2014.

[27] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertész, and A. Kanto. Dynamics of Market Correlations: Taxonomy and Portfolio Analysis. *Phys. Rev. E*, 68:056110, 2003.

[28] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, T. Guhr, and H. E. Stanley. Random Matrix Approach to Cross Correlations in Financial Data. *Phys. Rev. E*, 65:066126, 2002.

[29] J.P. Bouchaud and M. Potters. *Financial Applications of Random Matrix Theory: A Short Review*, chapter 40, pages 824–848. Oxford University Press, Oxford, 2011.

[30] J. Bun, J.-P. Bouchaud, and M. Potters. Cleaning Large Correlation Matrices: Tools from Random Matrix Theory. *Phys. Rep.*, 666:1 – 109, 2017.

[31] K. Anand, J. Khedair, and R. Kühn. A Structural Model for Fluctuations in Financial Markets. *Phys. Rev. E*, 97:052312, 2018.

[32] J. J. Hopfield. Neurons with Graded Responses Have Collective Computational Properties Like Those of Two-State Neurons. *Proc. Natl. Acad. Sci., USA*, 81:3088–3092, 1984.

[33] M Mézard and J Sakellariou. Exact Mean-Field Inference in Asymmetric Kinetic Ising Systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(07):L07001, 2011.

[34] H. C. Nguyen and J. Berg. Mean-Field Theory for the Inverse Ising Problem at Low Temperatures. *Phys. Rev. Lett.*, 109:050602, 2012.

[35] A. Braunstein, A. Ramezanpour, R. Zecchina, and P. Zhang. Inference and learning in sparse systems with multiple states. *Phys. Rev. E*, 83:056114, May 2011.

[36] S. Cocco, R. Monasson, L. Posani, S. Rosay, and J. Tubiana. Statistical Physics and Representations in Real and Artificial Neural Networks. *Physica A*, 504:45 – 76, 2018. Lecture Notes of the 14th International Summer School on Fundamental Problems in Statistical Physics.

[37] Basel Committee on Banking Supervision. Minimum Capital Requirements for Market Risk. www.bis.org, 2016.

[38] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters. Noise Dressing of Financial Correlation Matrices. *Phys. Rev. Lett.*, 83:1467–1470, 1999.

[39] P. Gopikrishnan, B. Rosenow, V. Plerou, and H. E. Stanley. Quantifying and Interpreting Collective Behavior in Financial Markets. *Phys. Rev. E*, 64:035106, Aug 2001.

[40] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. Nunes Amaral, and H. E. Stanley. Universal and Nonuniversal Properties of Cross Correlations in Financial Time Series. *Phys. Rev. Lett.*, 83:1471–1474, 1999.