

COMMUNITY DETECTION IN TEMPORAL MULTILAYER NETWORKS, AND ITS APPLICATION TO CORRELATION NETWORKS*

MARYA BAZZI[†], MASON A. PORTER[‡], STACY WILLIAMS[§], MARK MCDONALD[§],
DANIEL J. FENN[§], AND SAM D. HOWISON[¶]

Abstract.

Networks are a convenient way to represent complex systems of interacting entities. Many networks contain “communities” of nodes that are more densely connected to each other than to nodes in the rest of the network. In this paper, we investigate the detection of communities in temporal networks represented as multilayer networks. As a focal example, we study time-dependent financial-asset correlation networks. We first argue that the use of the “modularity” quality function—which is defined by comparing edge weights in an observed network to expected edge weights in a “null network”—is application-dependent. We differentiate between “null networks” and “null models” in our discussion of modularity maximization, and we highlight that the same null network can correspond to different null models. We then investigate a multilayer modularity-maximization problem to identify communities in temporal networks. Our multilayer analysis only depends on the form of the maximization problem and not on the specific quality function that one chooses. We introduce a diagnostic to measure *persistence* of community structure in a multilayer network partition. We prove several results that describe how the multilayer maximization problem measures a trade-off between static community structure within layers and higher values of persistence across layers. We also discuss some implementation issues that the popular “Louvain” heuristic faces with temporal multilayer networks and suggest ways to mitigate them.

Key words. Community structure, multilayer networks, temporal networks, modularity maximization, financial correlation networks.

AMS subject classifications. 62H30, 91C20, 94C15, 90C35

1. Introduction. In its simplest form, a network is simply a graph: it consists of a set of *nodes* that represent entities and a set of *edges* between pairs of nodes that represent interactions between those entities. One can consider *weighted graphs* (in which each edge has an associated edge weight that quantifies the interaction of interest) or *unweighted graphs* (weighted graphs with binary edge weights). Networks provide useful representations of complex systems across many disciplines [48]. Common types include social networks (which arise via offline and/or online interactions), information networks (e.g., hyperlinks between webpages in the World Wide Web), infrastructure networks (e.g., transportation routes between cities), and biological networks (e.g., metabolic interactions between cells or proteins, food webs, etc.).

Given a network representation of a system, it can be useful to apply a coarse-graining technique in order to investigate features that lie between features at the “microscale” (e.g., nodes and pairwise interactions) and the “macroscale” (e.g., total

*This work was supported by a CASE studentship award from the EPSRC (BK/10/41), HSBC Bank, and the FET-Proactive project PLEXMATH (FP7-ICT-2011-8; grant 317614) funded by the European Commission.

[†]Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute, Oxford OX2 6GG, UK (bazzi@maths.ox.ac.uk).

[‡]Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute, Oxford OX2 6GG, UK; and CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, UK (porter@maths.ox.ac.uk).

[§]Global Research, HSBC Bank, London E14 5HQ, UK.

[¶]Oxford Centre for Industrial and Applied Mathematics, Mathematical Institute, Oxford OX2 6GG, UK; and Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford OX2 6ED, UK (howison@maths.ox.ac.uk).

edge weight and degree distribution) [49,55]. One thereby studies “mesoscale” features such as core-periphery structure and (especially) community structure. Loosely speaking, a *community* (or *cluster*) in a network is a set of nodes that are “more densely” connected to each other than they are to nodes in the rest of the network [23,55]. Giving a precise definition of “densely connected” is, of course, necessary to have a method for community detection. It is important to recognize at the outset that this definition is subjective and in particular, may depend on the application in question. Correspondingly, community detection methods may need to be tailored. We restrict ourselves to *hard partitions*, in which each node is assigned to exactly one community, and we use the term “partition” to mean “hard partition”. It is also important, but beyond the scope of this paper, to consider “soft partitions”, in which communities can overlap [23,33,52,55].

Analysis of community structure has been very useful in a wide range of applications; many of which are described in [23,25,49,55]. In social networks, communities can reveal groups of people with common interests, places of residence, or other similarities [50,66]. In biological systems, communities can reveal functional groups that are responsible for synthesizing or regulating an important chemical product [29,40]. In the present paper, we use financial-asset correlation networks as examples [8,15]. Despite the diversity of markets, financial products, and geographical locations, financial assets can exhibit strong time-dependent correlations, both within and between asset classes. It is a primary concern for market practitioners (e.g., for portfolio diversification) to estimate the strengths of these correlations and to identify sets of assets that are highly correlated [43,67].

Most methods for detecting communities are designed for static networks. However, in many applications, entities and/or interactions between entities evolve in time. In such applications, one can use the formalism of *temporal networks*, where nodes and/or their edges weights vary in time [31]. This is important for numerous applications, including person-to-person communication [68], one-to-many information dissemination (e.g., Twitter networks [27] and Facebook networks [70]), cell biology [31], neuroscience [7], ecology [31], finance [20–22,51], and more.

Two main approaches have been adopted to detect communities in time-dependent networks. The first entails constructing a static network by aggregating snapshots of the evolving network at different points in time into a single network (e.g., by taking the mean or total edge weight for each edge across all time points, which can be problematic if the set of nodes varies in time and which also makes restrictive assumptions on the interaction dynamics between entities [30]). One can then use standard network techniques. The second approach entails using static community-detection techniques on each element of a time-ordered sequence of networks at different times or on each element of a time-ordered sequence of network aggregations¹ (computed as above) over different time intervals (which can be either overlapping or nonoverlapping) and then tracking the communities across the sequence [3,20,21,32,42,52].

A third approach consists of embedding a time-ordered sequence of networks in a larger network [18,44,54]. Each element of the sequence is a network *layer*, and nodes at different time points are joined by *inter-layer* edges. This approach was introduced

¹One needs to distinguish between this kind of aggregation and the averaging of a set of time series over a moving window to construct a correlation matrix, which one can then interpret as a fixed-time snapshot of a time-evolving network. Although both involve averaging over a time window, the former situation entails averaging a network, and the latter situation entails averaging over a collection of time series (one for each node) with no directly observable edge weights.

in [44] and the resulting network is a type of *multilayer network* [11, 35]. The main difference between this approach and the previous approach is that the presence of nonzero inter-layer edges introduces a dependence between communities identified in one layer and connectivity patterns in other layers. Thus far, most computations that have used a multilayer representation of temporal networks have assumed that inter-layer connections are “diagonal” (i.e., they exist only between copies of the same node) and “ordinal” (i.e., they exist only between consecutive layers) [35]. Diagonal is a natural model of the persistence of node identity in time, while ordinal preserves the time ordering.

The authors of [44] derived a generalization of *modularity maximization*, a popular clustering method for static networks, to multilayer networks. Modularity is a function that measures the “quality” of a network partition into disjoint sets of nodes by computing the difference between the total edge weight in sets in the observed network and the total expected edge weight in the same sets in a “null network” generated from some “null model” [23, 55]. Modularity maximization consists of maximizing the modularity quality function over the space of network partitions. (In practice, given the combinatorial complexity of this maximization problem, one uses some computational heuristic and finds a local maximum [28].) Intuitively, the null model controls for connectivity patterns that one anticipates finding in a network, and one uses modularity maximization to identify connectivity patterns in an observed network that are stronger than anticipated. We give a precise definition of the modularity function for single-layer networks in Section 2, where (importantly) we distinguish between a “null network” and a “null model” in modularity maximization. In Section 4, we discuss the choice of null network for a given application.

In Section 3, we describe the generalization of single-layer modularity to multilayer networks proposed in [44]. To date, almost no theory has explained how a multilayer partition obtained with zero inter-layer coupling (which reduces to single-layer modularity maximization on each layer independently) differs from a multilayer partition obtained with nonzero inter-layer coupling. In Section 5, we prove several theoretical properties of an optimal solution for the multilayer maximization problem to better understand how such partitions differ and how one can exploit this difference in practice. We also describe two implementation issues that arise when using the popular Louvain heuristic [10] to solve the multilayer maximization problem, and we suggest ways to mitigate them. The results of Section 5 are independent of the choice of quality function on individual layers and only depend on the form of the maximization problem. Section 6 contains a concluding discussion.

2. Single-layer modularity maximization.

2.1. The modularity function. Consider an N -node network \mathcal{G} and let the edge weights between pairs of nodes be $\{A_{ij} | i, j \in \{1, \dots, N\}\}$, so that $\mathbf{A} = (A_{ij}) \in \mathbb{R}^{N \times N}$ is the *adjacency matrix* of \mathcal{G} . In this paper, we only consider symmetric adjacency matrices (and hence undirected networks), so $A_{ij} = A_{ji}$ for all i and j . The *strength* of a node i is

$$k_i = \sum_{j=1}^N A_{ij} = \sum_{j=1}^N A_{ji}, \quad (2.1)$$

and it is given by the i^{th} row (or column) sum of \mathbf{A} .

When studying the structure of a network, it is useful to compare what is observed with what is anticipated. We define a *null model* to be a probability distribution on

the set of adjacency matrices and a *null network* to be the expected adjacency matrix under a specified null model. In a loose sense, null models play the role of *prior models*, as they control for features that one anticipates to find in the system under investigation. One can thereby take into account known (or suspected) connectivity patterns that might obscure unknown connectivity patterns that one hopes to discover via processes like community detection. For example, in social networks, one often takes the strength of a node in a null network to be its observed strength k_i [47, 49, 55]. We discuss the use of this null network for financial-asset correlation networks in Section 4. In spatial networks that represent the spread of a disease or information between different locations, some authors have used null networks in which edge weights between two locations scales inversely with the distance between them [19, 60].

As we discussed in Section 1, one uses modularity maximization to partition a network into sets of nodes called “communities” that have a larger total internal edge weight than the expected total internal edge weight in the same sets in a null network, generated from some null model [23, 49, 50, 55]. *Modularity maximization* consists of finding a partition that maximizes this difference [23, 55]. (As mentioned earlier, in practice one uses some computational heuristic and finds a local maximum [28]). In the present paper, we use the term “modularity” for an arbitrary choice of null network and we ignore any normalization constant that depends on the choice of null network but does not affect the solution of the modularity-maximization problem for a given null network. Modularity thus acts as a “quality function” $Q : \mathcal{C} \rightarrow \mathbb{R}$, where the set \mathcal{C} is the set of all possible N -node network partitions.

Suppose that we have a partition C of a network into K disjoint sets of nodes $\{C_1, \dots, C_K\}$. We can then define a map $c(\cdot)$ from the set of nodes $\{1, \dots, N\}$ to the set of integers $\{1, \dots, K\}$ such that $c(i) = c(j) = k$ if and only if nodes i, j lie in C_k . We call $c(i)$ the *set assignment* (or *community assignment* when C is a global or local maximum) of node i in partition C . The value of modularity for a given partition C is then

$$Q(C|\mathbf{A}; \mathbf{P}) := \sum_{i,j=1}^N (A_{ij} - P_{ij})\delta(c_i, c_j), \quad (2.2)$$

where $\mathbf{P} = (P_{ij}) \in \mathbb{R}^{N \times N}$ is the adjacency matrix of the null network, c_i is shorthand notation for $c(i)$, and $\delta(c_i, c_j)$ is the Kronecker delta function. We state the modularity-maximization problem as follows:

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^N (A_{ij} - P_{ij})\delta(c_i, c_j), \quad (2.3)$$

which we can also write as $\max_{C \in \mathcal{C}} Q(C|\mathbf{B})$ or $\max_{C \in \mathcal{C}} \sum_{i,j}^N B_{ij}\delta(c_i, c_j)$, where $\mathbf{B} = \mathbf{A} - \mathbf{P}$ is the so-called *modularity matrix* [47]. It is clear from (2.3) that pairwise contributions to modularity are only counted when two nodes are assigned to the same set. These contributions are positive (respectively, negative) when the observed edge weight A_{ij} between nodes i and j is larger (respectively, smaller) than the expected edge weight P_{ij} between them. If $A_{ij} < P_{ij}$ for all i and j , then the optimal solution is N singleton communities. Conversely, if $A_{ij} > P_{ij}$ for all i and j , then the optimal solution is a single N -node community. To obtain a partition of a network with a high value of modularity, one hopes to have many edges within sets that satisfy $A_{ij} > P_{ij}$

and few edges within sets that satisfy $A_{ij} < P_{ij}$. As is evident from equation (2.3), what one regards as “densely connected” in this setting depends fundamentally on the choice of null network.

It can be useful to write the modularity-maximization problem using the trace of matrices [47]. As before, we consider a partition C of a network into K sets of nodes $\{C_1, \dots, C_K\}$. We define the *partition matrix* $\mathbf{S} \in \{0, 1\}^{N \times K}$ as

$$S_{ij} = \delta(c_i, j), \quad (2.4)$$

where $j \in \{1, \dots, K\}$ and $c_i = j$ means that node i lies in C_j . The columns of \mathbf{S} are orthogonal, and the j^{th} column sum of \mathbf{S} gives the number of nodes in C_j . This yields

$$\sum_{i,j=1}^N B_{ij} \delta(c_i, c_j) = \sum_{i,j=1}^N \sum_{k=1}^K S_{ik} B_{ij} S_{jk} = \text{Tr}(\mathbf{S}^T \mathbf{B} \mathbf{S}),$$

where the $(i, i)^{\text{th}}$ term of $\mathbf{S}^T \mathbf{B} \mathbf{S}$ is twice the sum of edge weights in C_i . (The $(i, j)^{\text{th}}$ off-diagonal term is twice the sum of edge weights between C_i and C_j .) It follows that one can restate the modularity-maximization problem in (2.3) as

$$\max_{\mathbf{S} \in \mathcal{S}} \text{Tr}(\mathbf{S}^T \mathbf{B} \mathbf{S}), \quad (2.5)$$

where \mathcal{S} is the set of all partition matrices in $\{0, 1\}^{N \times K}$ (with $K \leq N$).

Modularity maximization is one of myriad community-detection methods [23], and it has many limitations (e.g., a resolution limit on the size of communities [24] and a huge number of nearly degenerate local maxima [28]). Nevertheless, it is a popular method (which has been used successfully in numerous applications [23, 55]), and the ability to specify explicitly what one anticipates is a useful (and under-exploited) feature for users working on different applications. In Section 4, we make some observations on one’s choice of null network when using the modularity quality function.

2.2. The Louvain computational heuristic. For a given modularity matrix \mathbf{B} , a solution to the modularity-maximization problem is guaranteed to exist in any network with a finite number of nodes. However, the number of possible partitions in an N -node network, given by the Bell number [9], grows at least exponentially with N , so an exhaustive search of the space of partitions is infeasible. Modularity maximization was proven in [14] to be an NP-hard problem (at least for the null networks which we consider in this paper), so solving it requires the use of computational heuristics. In the present paper, we focus on the Louvain heuristic, which is a locally-greedy modularity-increasing sampling process over the set of partitions [10].

The Louvain heuristic consists of two phases, which are repeated iteratively. Initially, each node in the network constitutes a set, which gives an initial partition that consists of N singletons. During phase 1, one considers the nodes one by one (in some order), and one places each node in a set (including its own) that results in the largest increase of modularity. This phase is repeated until one reaches a local maximum (i.e., until one obtains a partition in which the move of a single node cannot increase modularity). Phase 2 consists of constructing a reduced network \mathcal{G}' from the sets of nodes in \mathcal{G} that one obtains after the convergence of phase 1. We denote the sets in \mathcal{G} at the end of phase 1 by $\{\hat{C}_1, \dots, \hat{C}_{\hat{N}}\}$ (where $\hat{N} \leq N$) and the set assignment of node i

in this partition by \hat{c}_i . Each set \hat{C}_k in \mathcal{G} constitutes a node k in \mathcal{G}' , and the reduced modularity matrix of \mathcal{G}' is

$$\mathbf{B}' = \hat{\mathbf{S}}^T \mathbf{B} \hat{\mathbf{S}},$$

where $\hat{\mathbf{S}}$ is the partition matrix of $\{\hat{C}_1, \dots, \hat{C}_N\}$. This ensures that the all-singleton partition in \mathcal{G}' has the same value of modularity as the partition of \mathcal{G} that we identified at the end of phase 1. One then repeats phase 1 on the reduced network and continues iterating until the heuristic converges (i.e., until phase 2 induces no further changes).

Because we use a nondeterministic implementation of the Louvain heuristic—in particular, the node order is randomized at the start of each iteration of phase 1—the network partitions that we obtain for a fixed modularity matrix can differ across runs.² To account for this, one can compute the frequency of co-classification of nodes into communities for a given modularity matrix \mathbf{B} across multiple runs of the heuristic instead of using the output partition of a single run. (See [39] for an application of such an approach to “consensus clustering” and [59] for an application of such an approach to hierarchical clustering.) We use the term *association matrix* for a matrix that stores the mean number of times that two nodes are placed in the same community across multiple runs of a heuristic, and we use the term *co-classification index* of nodes i and j to designate the $(i, j)^{\text{th}}$ entry of an association matrix.

There are many other heuristics that one can employ to maximize modularity [23, 48, 55], but the Louvain heuristic is a popular choice in practice [38]. It is very fast [23, 38], which is an important consideration in multilayer networks, for which the total number of nodes is the number of nodes in each layer multiplied by the number of layers. In Section 5, we point out two issues that the Louvain heuristic (independently of how it is implemented) faces with temporal multilayer networks.

2.3. Multiscale community structure. Many networks include community structure at multiple scales [23, 55], and some systems even have a hierarchical community structure of “parts-within-parts” [62]. In such a situation, although there are dense interactions within communities of some size (e.g., friendship ties between students in the same school), there are even denser interactions in subsets of nodes that lie inside these communities (e.g., friendship ties between students in the same school and in the same class year). Some variants of the modularity function have been proposed to detect communities at different scales. A popular choice is to scale the null network by a *resolution parameter* $\gamma \geq 0$ to yield a *multiscale modularity-maximization* problem [58]:

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^N (A_{ij} - \gamma P_{ij}) \delta(c_i, c_j). \quad (2.6)$$

In some sense, the value of the parameter γ determines the importance that one assigns to the null network relative to the observed network. The corresponding modularity matrix and modularity function evaluated at a partition C are $\mathbf{B} =$

²The implementation [1, 34] of the heuristic that we use in this paper is a generalized version of the implementation in [10]. It is independent of the null network—so it takes the modularity matrix as an input to allow an arbitrary choice of null network—and it randomizes the node order at the start of each iteration of the heuristic’s first phase to increase the search space of the heuristic. When one chooses the same null network that was assumed in [10] and uses a node order fixed to $\{1, \dots, N\}$ at each iteration of phase 1 (the value of N can change after each iteration of the heuristic’s second phase), then these implementations return the same output.

$\mathbf{A} - \gamma \mathbf{P}$ and $Q(C|\mathbf{A}; \mathbf{P}; \gamma) = \sum_{i,j=1}^N (A_{ij} - \gamma P_{ij}) \delta(c_i, c_j)$. The special case $\gamma = 1$ yields the modularity matrix and modularity function in the modularity-maximization problem (2.3). This formulation of multiscale modularity has a dynamical interpretation [36, 37] that we will discuss in the next subsection.

In most applications of community detection, the adjacency matrix of the observed and null networks have nonnegative entries. In these cases, the solution to (2.6) when $0 \leq \gamma \leq \gamma^- = \min_{i \neq j, P_{ij} \neq 0} (A_{ij}/P_{ij})$ is a single community regardless of any structure, however clear, in the observed network, because then

$$B_{ij} = A_{ij} - \gamma P_{ij} \leq 0 \quad \text{for all } i, j \in \{1 \dots N\}.$$

(We exclude diagonal terms because a node is always in its own community.) However, the solution to (2.6) when $\gamma > \gamma^+ = \max_{i \neq j, P_{ij} \neq 0} (A_{ij}/P_{ij})$ is N singleton communities because

$$B_{ij} = A_{ij} - \gamma P_{ij} < 0 \quad \text{for all } i, j \in \{1 \dots N\}.$$

Partitions at these boundary values of γ correspond to the coarsest and finest possible partitions of a network, and varying the resolution parameter between these bounds makes it possible to examine a network's community structure at intermediate scales.

For an observed and/or null network with signed edge weights, the intuition behind the effect of varying γ in (2.6) on an optimal solution is not straightforward. A single community and N singleton communities do not need to be optimal partitions for any value of $\gamma \geq 0$. In particular, B_{ij} has the same sign as A_{ij} for sufficiently small values of γ , and B_{ij} has the opposite sign to P_{ij} for sufficiently large values of γ . For further discussion, see Section 4, where we explore the effect of varying the resolution parameter on an optimal partition for an observed and null network with signed edge weights. We take $\gamma^- = 0$ in numerical experiments with signed networks because optimal partitions can vary in the interval $\gamma \in [0, \gamma^-]$.³

It is important to differentiate between a “resolution limit” on the smallest community size that is imposed by a community-detection method [24] and inherently multiscale community structure in a network [23, 55, 62]. For the formulation of multiscale modularity in (2.6), the resolution limit described in [24] applies to any fixed value of γ . By varying γ , one can identify communities that are smaller than the limit for any particular γ value. In this sense, multiscale formulations of modularity help “mitigate” the resolution limit, though there remain issues [2, 28, 36]. In this paper, we do not address the issue of how to identify communities at different scales, though we note in passing that the literature includes variants of multiscale modularity (e.g., see [2]). We make observations on null networks in Section 4, and we illustrate how our observations can manifest in practice using the formulation of multiscale modularity in (2.6). (Our observations hold independently of the formulation of multiscale modularity that one adopts, but the precise manifestation can be different for different variants of multiscale modularity.)

We use the term *multiscale community structure* to refer to a set $\mathcal{C}_{\text{local}}(\gamma)$ of local optima that we obtain with a computational heuristic for a set of (not necessarily all distinct) resolution-parameter values $\gamma = \{\gamma_1, \dots, \gamma_l\}$, where $\gamma^- = \gamma_1 \leq \dots \leq \gamma_l = \gamma^+$. We use the term *multiscale association matrix* for an association matrix

³For $\gamma \geq \gamma^+$, one can show that all modularity contributions no longer change signs: these are negative (respectively, positive) between pairs of nodes with $P_{ij} \geq 0$ (respectively, $P_{ij} \leq 0$).

$\hat{\mathbf{A}} \in [0, 1]^{N \times N}$ that stores the co-classification index of all pairs of nodes for partitions in this set:

$$\hat{A}_{ij} = \frac{\sum_{C \in \mathcal{C}_{\text{local}}(\gamma)} \delta(c_i, c_j)}{|\mathcal{C}_{\text{local}}(\gamma)|}. \quad (2.7)$$

We use this matrix repeatedly in our computational experiments of Section 4.

2.4. Null models and null networks. In this section, we describe three null networks. We make several observations on the interpretation of communities that we obtain from Pearson correlation matrices using each of these null networks in the computational experiments of Section 4.

2.4.1. Newman-Girvan (NG) null network. A popular choice of null network for networks with positive edge weights is the Newman-Girvan (NG) null network, whose adjacency-matrix entries are $P_{ij} = k_i k_j / (2m)$, where k_i are the observed node strengths [46, 50]. This yields the equivalent maximization problems

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^N \left(A_{ij} - \frac{k_i k_j}{2m} \right) \delta(c_i, c_j) \Leftrightarrow \max_{S \in \mathcal{S}} \text{Tr} \left[S^T \left(\mathbf{A} - \frac{\mathbf{k} \mathbf{k}^T}{2m} \right) S \right], \quad (2.8)$$

where $\mathbf{k} = \mathbf{A} \mathbf{1}$ is the $N \times 1$ vector of node strengths and $2m = \mathbf{1}^T \mathbf{A} \mathbf{1}$ is the total edge weight of the observed network. This null network can be derived from a variety of null models. One way to generate an unweighted network with expected adjacency matrix $\mathbf{k} \mathbf{k}^T / (2m)$ is to generate each of its edges and self-edges with probability $k_i k_j / (2m)$ (provided $k_i k_j \leq 2m$ for all i, j). That is, the presence and absence of edges and self-edges is a Bernoulli random variable with probability $k_i k_j / (2m)$ [12, 13]. More generally, any probability distribution on the set of adjacency matrices that satisfies $\mathbb{E}(\sum_{j=1}^N W_{ij}) = k_i$ (i.e., the expected strength equals the observed strength, see for e.g., [16]) and $\mathbb{E}(W_{ij}) = f(k_i) f(k_j)$ for some real-valued function f has an expected adjacency matrix of $\mathbb{E}(\mathbf{W}) = \mathbf{k} \mathbf{k}^T / (2m)$.⁴ The adjacency matrix of the NG null network is symmetric and positive semidefinite.

We briefly describe a way of deriving the NG null network from a model on time-series data (in contrast to a model on a network). The *partial correlation* $\text{corr}(a, b | c)$ between a and b given c is the Pearson correlation between the residuals that result from the linear regression of a with c and b with c , and it is given by

$$\text{corr}(a, b | c) = \frac{\text{corr}(a, b) - \text{corr}(a, c) \text{corr}(b, c)}{\sqrt{1 - \text{corr}^2(a, c)} \sqrt{1 - \text{corr}^2(b, c)}}. \quad (2.9)$$

Suppose that the data used to construct the observed network is a set of time series $\{z_i | i \in \{1, \dots, N\}\}$, where $z_i = \{z_i(t) | t \in T\}$ and T is a discrete set of time points. The authors in [41] pointed out that when $A_{ij} = \text{corr}(z_i, z_j)$ then $k_i = \text{cov}(\hat{z}_i, \hat{z}_{\text{tot}})$ and thus that

$$\frac{k_i k_j}{2m} = \text{corr}(\hat{z}_i, \hat{z}_{\text{tot}}) \text{corr}(\hat{z}_j, \hat{z}_{\text{tot}}), \quad (2.10)$$

⁴The linearity of the expectation and the assumptions $\mathbb{E}(\sum_{j=1}^N \mathbf{W}_{ij}) = k_i$ and $\mathbb{E}(\mathbf{W}_{ij}) = f(k_i) f(k_j)$ imply that $f(k_i) = k_i / \sum_{j=1}^N f(k_j)$ and $\sum_{j=1}^N f(k_j) = \sqrt{2m}$. Combining these equations gives the desired result.

where $\hat{z}_i(t) = (z_i(t) - \langle z_i \rangle) / \sigma(z_i)$ is a standardized time series and $\hat{z}_{\text{tot}}(t) = \sum_{i=1}^N \hat{z}_i(t)$ is the sum of the standardized time series.⁵ Taking $a = \hat{z}_i$, $b = \hat{z}_j$, and $c = \hat{z}_{\text{tot}}$, equation (2.9) implies that if $\text{corr}(\hat{z}_i, \hat{z}_j | \hat{z}_{\text{tot}}) = 0$ then $\text{corr}(\hat{z}_i, \hat{z}_j) = k_i k_j / (2m)$. That is, Pearson correlation coefficients between pairs of time series that satisfy $\text{corr}(\hat{z}_i, \hat{z}_j | \hat{z}_{\text{tot}}) = 0$ are precisely the adjacency-matrix entries of the NG null network. One way of generating a set of time series in which pairs of distinct time series satisfy this condition is to assume that each standardized time series depends linearly on the mean time series and that residuals are mutually uncorrelated (i.e., $\hat{z}_i = \alpha_i \hat{z}_{\text{tot}} / N + \beta_i + \epsilon_i$ for some $\alpha_i, \beta_i \in \mathbb{R}$ and $\text{corr}(\epsilon_i, \epsilon_j) = 0$ for $i \neq j$).

The multiscale modularity-maximization problem in (2.6) was initially introduced in [58] using an *ad hoc* approach. Interestingly, one can derive this formulation of the maximization problem for sufficiently large values of γ by considering a quality function based on a continuous-time Markov process $X(t)$ on an observed network [36, 37]. The probability density of a continuous-time Markov process with exponentially distributed waiting times at each node parametrized by $\lambda(i)$ satisfies

$$\dot{\mathbf{p}} = \mathbf{p} \mathbf{\Lambda} \mathbf{M} - \mathbf{p} \mathbf{\Lambda}, \quad (2.11)$$

where the vector $\mathbf{p}(t) \in [0, 1]^{1 \times N}$ is the probability density of a random walker at each node [i.e., $p_i(t) := \mathbb{P}(X(t) = i)$ for each i], $\mathbf{\Lambda}$ is a diagonal matrix with the rate $\lambda(i)$ on its i^{th} diagonal entry, and \mathbf{M} is the transition matrix of a random walker (i.e., $M_{ij} := A_{ij} / k_i$). The solution to equation (2.11) is $\mathbf{p}(t) = \mathbf{p}_0 e^{\mathbf{\Lambda}(\mathbf{M} - \mathbf{I})t}$ and its stationary distribution is $\boldsymbol{\pi} = \mathbf{k}^T \mathbf{\Lambda}^{-1} / (2m)$. The *stability* of a partition is a quality function defined by [17, 36, 37]

$$r(\mathbf{S}, t) = \text{Tr} \left[\mathbf{S}^T \left(\boldsymbol{\Pi} e^{\mathbf{\Lambda}(\mathbf{M} - \mathbf{I})t} - \boldsymbol{\pi}^T \boldsymbol{\pi} \right) \mathbf{S} \right],$$

where $\Pi_{ij} = \delta(i, j) \pi_i$. Equivalently, the stability is

$$r(C, t) = \sum_{i,j=1}^N \left[\pi_i \left(e^{\mathbf{\Lambda}(\mathbf{M} - \mathbf{I})t} \right)_{ij} - \pi_i \pi_j \right] \delta(c_i, c_j). \quad (2.12)$$

Taking $\mathbf{p}_0 = \boldsymbol{\pi}$, the term in brackets on the left-hand side of (2.12) is $\mathbb{P}(X(0) = i \cap X(t) = j)$ and the term in brackets on the right-hand side is $\mathbb{P}(X(0) = i \cap X(t \rightarrow \infty) = j)$ (provided the system is ergodic). The intuition behind the stability quality function is that a good partition at a given time before reaching stationarity corresponds to one in which the time that a random walker spends within communities is large compared with the time that it spends transiting between communities. In other words, a random walker that starts out at a community ends up there again in the early stages of the random walk, long before stationarity. The resulting maximization problem is $\max_{\mathbf{S} \in \mathcal{S}} r(\mathbf{S}, t)$, or equivalently $\max_{C \in \mathcal{C}} r(C, t)$. By linearizing $e^{\mathbf{\Lambda}(\mathbf{M} - \mathbf{I})t}$ at $t = 0$ and taking $\mathbf{\Lambda} = \mathbf{I}$, one obtains the multiscale modularity-maximization problem in (2.6) at short timescales with $\gamma = 1/t$ and $P_{ij} = k_i k_j / (2m)$. This approach provides a dynamical interpretation of the resolution parameter γ as the inverse (after linearization) of the time used to explore a network by a random walker.

⁵The equality (2.10) holds for signed correlation networks. The strength of a node i is given by the i^{th} (signed) column or row sum of the correlation matrix.

2.4.2. Generalization of Newman-Girvan null network to signed networks (NGS). In [26], Gómez et al. proposed a generalization of the NG null network to signed networks. They separated \mathbf{A} into its positive and negative edge weights:

$$\mathbf{A} = \mathbf{A}^+ - \mathbf{A}^-,$$

where \mathbf{A}^+ denotes the positive part of \mathbf{A} and $-\mathbf{A}^-$ denotes its negative part. Their generalization of the NG null network to signed networks (NGS) is $P_{ij} = k_i^+ k_j^+ / (2m^+) - k_i^- k_j^- / (2m^-)$. This yields the maximization problem

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^N \left[\left(A_{ij}^+ - \frac{k_i^+ k_j^+}{2m^+} \right) - \left(A_{ij}^- - \frac{k_i^- k_j^-}{2m^-} \right) \right] \delta(c_i, c_j), \quad (2.13)$$

where k_i^+ and $2m^+$ (respectively, k_i^- and $2m^-$) are the strengths and total edge weight in \mathbf{A}^+ (respectively, \mathbf{A}^-). The intuition behind this generalization is to use an NG null network on both unsigned matrices \mathbf{A}^+ and \mathbf{A}^- but to count contributions to modularity from negative edge weights (i.e., the second group of terms in (2.13)) in an opposite way to those from positive edge weights (i.e., the first group of terms in (2.13)). Negative edge weights that exceed their expected edge weight are penalized (i.e., they decrease modularity) and those that do not are rewarded (i.e., they increase modularity). One can generate a network with edge weights 0, 1, or -1 and expected edge weights $k_i^+ k_j^+ / (2m^+) - k_i^- k_j^- / (2m^-)$ by generating one network with expected edge weights $W_{ij}^+ = k_i^+ k_j^+ / (2m^+)$ and a second network with expected edge weights $W_{ij}^- = k_i^- k_j^- / (2m^-)$ using the procedure described for the NG null network in Section 2.4.1, and then defining a network whose edge weights are given by the difference between the edge weights of these two networks. More generally, any probability distribution on the set of signed adjacency matrices $\{\mathbf{W} \in \mathbb{R}^{N \times N}\}$ with the same properties as those for the NG null network for \mathbf{W}^+ and \mathbf{W}^- (where $\mathbf{W} = \mathbf{W}^+ - \mathbf{W}^-$ defined as above) will have expected edge weights of $W_{ij} = k_i^+ k_j^+ / (2m^+) - k_i^- k_j^- / (2m^-)$ for all $i, j \in \{1, \dots, N\}$ (by linearity of the expectation).

The authors of [44] derived a variant of the multiscale formulation of modularity in (2.6) for the NGS null network at short time scales by building on the random-walk approach used to derive the NG null network.⁶ They considered the function

$$\hat{r}(C, t) = \sum_{i,j=1}^N \left(\pi_i [\delta_{ij} + t \Lambda_{ii}(M_{ij} - \delta_{ij})] - \pi_i \rho_{i|j} \right) \delta(c_i, c_j), \quad (2.14)$$

where the term in brackets on the left-hand side is a linearization of the exponential term in (2.12), \mathbf{M} and π_i are as defined in (2.12) on a network with adjacency matrix $|\mathbf{A}| := \mathbf{A}^+ + \mathbf{A}^-$, and $\rho_{i|j}$ is the probability of jumping from node i to node j at stationarity conditional on the network structure [44]. If the network is unipartite, unsigned, and undirected, then $\rho_{i|j}$ reduces to the stationary probability π_j .

⁶In particular, they derived the multiscale formulation of modularity obtained using a Potts-model approach in [64]. This multiscale formulation results in one resolution parameter γ_1 for the term $(k_i^+ k_j^+) / (2m^+)$ and a second resolution parameter γ_2 for the term $(k_i^- k_j^-) / (2m^-)$ in (2.13) (see [42] for an application of this multiscale formulation to the United Nations General Assembly resolution networks). Without an application-driven justification for how to choose these parameters, this increases the parameter space substantially, so we only consider the case $\gamma_1 = \gamma_2$ in this paper.

2.4.3. Uniform (U) null network. A third null network that we consider is a uniform (U) null network, with adjacency-matrix entries $P_{ij} = \langle k \rangle^2 / (2m)$, where $\langle k \rangle := (\sum_{i=1}^N k_i) / N$ denotes the mean strength in a network. We thereby obtain the equivalent maximization problems

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^N \left(A_{ij} - \frac{\langle k \rangle^2}{2m} \right) \delta(c_i, c_j) \Leftrightarrow \max_{S \in \mathcal{S}} \text{Tr} \left[S^T \left(\mathbf{A} - \frac{\langle k \rangle^2}{2m} \mathbf{1}_N \right) S \right], \quad (2.15)$$

where \mathbf{A} is an unsigned adjacency matrix and $\mathbf{1}_N$ is an $N \times N$ matrix in which every entry is 1.⁷ The expected edge weight in (2.15) is constant and satisfies

$$\frac{\langle k \rangle^2}{2m} = \frac{\left(\sum_{i=1}^N k_i / N \right)^2}{\sum_{i=1}^N k_i} = \frac{2m}{N^2} = \langle \mathbf{A} \rangle,$$

where $\langle \mathbf{A} \rangle$ denotes the mean value of the adjacency matrix.⁸ One way to generate an unweighted network with adjacency matrix $\langle \mathbf{A} \rangle \mathbf{1}_N$ is to generate each edge with probability $\langle \mathbf{A} \rangle$ (provided $\langle \mathbf{A} \rangle \leq 1$). That is, the presence and absence of an edge (including self-edges) are independent and identically distributed (i.i.d.) Bernoulli random variables with probability $\langle \mathbf{A} \rangle$. More generally, any probability distribution on the set of adjacency matrices that satisfies $\mathbb{E}(\sum_{i,j=1}^N W_{ij}) = 2m$ and $\mathbb{E}(W_{ij}) = \mathbb{E}(W_{i'j'})$ for all i, j, i', j' has an expected adjacency matrix $\mathbb{E}(\mathbf{W}) = \langle \mathbf{A} \rangle \mathbf{1}_N$. The adjacency matrix of the U null network is symmetric and positive semidefinite. One can derive the multiscale formulation in (2.6) for the U null network from the stability quality function in precisely the same way as it is derived for the NG null network, except that one needs to consider exponentially distributed waiting times at each node with rates proportional to node strength (i.e., $\Lambda_{ij} = \delta(i, j) k_i / \langle k \rangle$) [37].

3. Multilayer modularity maximization.

3.1. Multilayer representation of temporal networks. We restrict our attention to temporal networks in which only edges vary in time. (Thus, each node is present in all layers.) We use the notation \mathbf{A}_s for a *layer* in a sequence of adjacency matrices $\mathcal{T} = \{\mathbf{A}_1, \dots, \mathbf{A}_{|\mathcal{T}|}\}$, and we denote node i in layer s by i_s . We use the term *multilayer network* for a network defined on the set of nodes $\{1_1, \dots, N_1; 1_2, \dots, N_2; \dots; 1_{|\mathcal{T}|}, \dots, N_{|\mathcal{T}|}\}$ [35].

Thus far, computations that have used a multilayer framework for temporal networks have almost always assumed (1) that inter-layer connections exist only between nodes that correspond to the same entity (i.e., between nodes i_s and i_r for some i and $s \neq r$) and (2) that the network layers are “ordinal” (i.e., inter-layer edges exist only between consecutive layers) [7, 35, 44, 45, 57]. It is also typically assumed that (3) inter-layer connections are uniform (i.e., inter-layer edges have the same weight). In a recent review article on multilayer networks [35], condition (1) was called “diagonal” coupling, and condition (2) implies that a network is “layer-coupled”. We refer to the

⁷For a network in which all nodes have the same strength, the uniform and Newman-Girvan null networks are equivalent because $k_i = k_j$ for all $i, j \Leftrightarrow k_i = 2m/N = \langle k \rangle$ for all i . This was pointed out for an application to foreign exchange markets in [20, 21].

⁸Although we use the uniform null network on unsigned adjacency matrices in this paper, the expected edge weight in the uniform null network is always nonnegative for correlation matrices, as positive semidefiniteness guarantees that $\langle \mathbf{A} \rangle = \mathbf{1}^T \mathbf{A} \mathbf{1} / (N^2) \geq 0$.




FIG. 3.1. Example of (left) a multilayer network with unweighted intra-layer connections (solid lines) and uniformly weighted inter-layer connections (dashed curves) and (right) its corresponding adjacency matrix. (The adjacency matrix that corresponds to a multilayer network is sometimes called a “supra-adjacency matrix” in the network-science literature [35].)

type of coupling defined by (1), (2), and (3) as *ordinal diagonal* and *uniform* inter-layer coupling and we denote the value of the inter-layer edge weight by $\omega \in \mathbb{R}$. We show a simple illustration of a multilayer network with ordinal diagonal and uniform inter-layer coupling in Fig. 3.1. One can consider more general inter-layer connections (e.g., nonuniform ones). Although we restrict our attention to uniform coupling in our theoretical and computational discussions, we give an example of a nonuniform choice of inter-layer coupling in Section 5. Results similar to those of subsection 5.2 also apply in this more general case.

3.2. The multilayer modularity function. The authors of [44] generalized the single-layer multiscale modularity-maximization problem in (2.6) to a multilayer network using a similar approach as the one used to derive the NGS null network from a stochastic Markov process on the observed network. For simplicity, we express intra-layer and inter-layer connections in an $N|\mathcal{T}|$ -node multilayer network using a single $N|\mathcal{T}| \times N|\mathcal{T}|$ matrix. Each node i_s in layer s has the unique index $i' := i + (s - 1)N$, and we use \mathcal{A} to denote the multilayer adjacency matrix, which has entries $\mathcal{A}_{i'j'} = A_{ijs}\delta(s, r) + \omega\delta(|s - r|, 1)$ when the inter-layer coupling is ordinal diagonal and uniform. (As discussed in [35], one can use either an adjacency tensor or an adjacency matrix to represent a multilayer network.) The generalization in [44] consists of applying the function in (2.14) to the $N|\mathcal{T}|$ -node multilayer network:

$$\hat{r}(C, t) = \sum_{i,j=1}^{N|\mathcal{T}|} \left(\pi_i [\delta_{ij} + t\Lambda_{ii}(\mathcal{M}_{ij} - \delta_{ij})] - \pi_i \rho_{i|j} \right) \delta(c_i, c_j), \quad (3.1)$$

where C is now a multilayer partition (i.e., a partition of an $N|\mathcal{T}|$ -node multilayer network), Λ is the $N|\mathcal{T}| \times N|\mathcal{T}|$ diagonal matrix with the rates of the exponentially distributed waiting times at each node of each layer on its diagonal, \mathcal{M} (with entries $\mathcal{M}_{ij} := \mathcal{A}_{ij} / \sum_j \mathcal{A}_{ij}$) is the $N|\mathcal{T}| \times N|\mathcal{T}|$ transition matrix for the $N|\mathcal{T}|$ -node multilayer network with adjacency matrix \mathcal{A} , π_i is the corresponding stationary distribution (with the strength of a node and the total edge weight now computed from the multilayer adjacency matrix \mathcal{A}), and $\rho_{i|j}$ is the probability of jumping from node i to node j at stationarity conditional on the structure of the network within and between layers. The authors’ choice of $\rho_{i|j}$, which accounts for the sparsity pattern of inter-layer edges in the multilayer network, leads to the *multilayer modularity-maximization*

problem

$$\max_{C \in \mathcal{C}} \sum_{i,j=1}^{N|\mathcal{T}|} \mathcal{B}_{ij} \delta(c_i, c_j). \quad (3.2)$$

which we can also write as $\max_{C \in \mathcal{C}} Q(C|\mathcal{B})$, where \mathcal{B} is the *multilayer modularity matrix*

$$\mathcal{B} = \begin{bmatrix} \mathbf{B}_1 & \omega \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \omega \mathbf{I} & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \ddots & \ddots & \ddots & \mathbf{0} \\ \vdots & \ddots & \ddots & \ddots & \omega \mathbf{I} \\ \mathbf{0} & \dots & \mathbf{0} & \omega \mathbf{I} & \mathbf{B}_{|\mathcal{T}|} \end{bmatrix}, \quad (3.3)$$

and \mathbf{B}_s is a single-layer modularity matrix computed on layer s . (For example, $\mathbf{B}_s = \mathbf{A}_s - \langle \mathbf{A}_s \rangle \mathbf{1}_N$ if one uses the U null network and sets $\gamma = 1$.) We rewrite the multilayer modularity-maximization problem in [44] as

$$\max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \sum_{s=1}^{|\mathcal{T}|-1} \sum_{i=1}^N \delta(c_{i_s}, c_{i_{s+1}}), \quad (3.4)$$

where B_{ijs} denotes the $(i, j)^{\text{th}}$ entry of \mathbf{B}_s . Equation (3.4) clearly separates intra-layer contributions (left term) from inter-layer contributions (right term) to the multilayer quality function.

In practice, one can solve this multilayer modularity-maximization problem with the Louvain heuristic in subsection 2.2 by using the multilayer modularity matrix \mathcal{B} instead of the single-layer modularity matrix \mathbf{B} as an input (the number of nodes in the first iteration of phase 1 becomes $N|\mathcal{T}|$ instead of N). It is clear from (3.4) that inter-layer merges decrease the value of the multilayer quality function when $\omega < 0$, so we only consider $\omega \geq 0$. Furthermore, although merging pairs of nodes with $\mathcal{B}_{ij} = 0$ into the same set does not change the value of the quality function, we will assume in the rest of the paper that sets in globally optimal partitions of the multilayer modularity-maximization problem (3.4) do not contain disconnected components in the $N|\mathcal{T}|$ -node weighted graph with adjacency matrix \mathcal{B} .

In Section 5, we try to gain some insight into how to interpret a globally optimal multilayer partition by proving several properties that it satisfies. The results that we show are *independent* of the choice of matrices $\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}$, so (for example) they still apply when one uses the stability quality function in (2.12) on each layer instead of the modularity quality function. For ease of writing (and because modularity is the quality function that we use in our computational experiments of Section 5), we will continue to refer to the maximization problem (3.4) as a multilayer modularity maximization problem.

4. Interpretation of community structure in correlation networks with different null networks. It is clear from the structure of \mathcal{B} in equation (3.3) that the choice of quality function within layers (i.e., diagonal blocks in the multilayer adjacency matrix) and the choice of coupling between layers (i.e., off-diagonal blocks) for a given quality function affect the solution of the maximization problem in (3.4). In

this section, we make some observations on the choice of null network for correlation networks when using the modularity quality function. To do this, we consider the multilayer modularity-maximization problem (3.4) with zero inter-layer coupling (i.e., $\omega = 0$), which is equivalent to performing single-layer modularity maximization on each layer independently.

4.1. Toy examples. We describe two simple toy networks to illustrate some features of the NG (2.8) and NGS (2.13) null networks that can be misleading for asset correlation networks.

4.1.1. NG null network. Assume that the nodes in a network are divided into K nonoverlapping categories (e.g., asset classes) such that all intra-category edge weights have a constant value $a > 0$ and all inter-category edge weights have a constant value b , with $0 \leq b < a$. Let κ_i denote the category of node i , and rewrite the strength of node i as

$$k_i = |\kappa_i|a + (N - |\kappa_i|)b = |\kappa_i|(a - b) + Nb.$$

The strength of a node in this network scales linearly with the number of nodes in its category. Suppose that we have two categories κ_1, κ_2 that do not contain the same number of nodes. Taking $|\kappa_1| > |\kappa_2|$ without loss of generality, it follows that

$$P_{i,j \in \kappa_1} = \frac{1}{2m} \left[|\kappa_1|(a - b) + Nb \right]^2 > \frac{1}{2m} \left[|\kappa_2|(a - b) + Nb \right]^2 = P_{i,j \in \kappa_2}, \quad (4.1)$$

where $P_{i,j \in \kappa_i}$ is the expected edge weight between pairs of nodes in κ_i in the NG null network. That is, pairs of nodes in an NG null network that belong to larger categories have a larger expected edge weight than pairs of nodes that belong to smaller categories.

To see how equation (4.1) can lead to misleading results, we perform a simple experiment. Consider the toy network in Fig. 4.1(a) that contains 100 nodes divided into four categories of sizes 40, 30, 20, and 10. We set intra-category edge weights to 1 and inter-category edge weights to 0.3 (i.e., $a = 1$ and $b = 0.3$ in equation (4.1)). In Fig. 4.1(b) (respectively, Fig. 4.1(c)), we show the multiscale association matrix defined in (2.7) using an NG null network (respectively, a U null network). Colors scale with the frequency of co-classification of pairs of nodes into the same community across resolution-parameter values. Because the nodes are ordered by category, diagonal blocks in Fig. 4.1(b,c) indicate the co-classification index of nodes in the same category, and off-diagonal blocks indicate the co-classification index of nodes in different categories. We observe in Fig. 4.1(b) that larger categories are identified as a community across a smaller range of resolution-parameter than smaller categories when using an NG null network. In particular, category κ is identified as a single community when $\gamma < a/P_{i,j \in \kappa}$ (with $a/P_{i,j \in \kappa_1} < a/P_{i,j \in \kappa_2}$ when $|\kappa_1| > |\kappa_2|$ by equation (4.1)). When $\gamma \geq a/P_{i,j \in \kappa}$, category κ is identified as $|\kappa|$ singleton communities. However, we observe in Fig. 4.1(c) that all four categories are identified as a single community across the same range of resolution-parameter values when using the U null network. In particular, category κ is identified as a single community when $\gamma < a/\langle \mathbf{A} \rangle$ and as $|\kappa|$ singleton communities when $\gamma \geq a/\langle \mathbf{A} \rangle$.

The intuition behind multiscale modularity maximization is that the communities that one obtains for larger values of γ reveal “more densely” connected nodes in the observed network. Although all diagonal blocks in Fig. 4.1(a) have the same internal




FIG. 4.1. (a) Toy unsigned block matrix with constant diagonal and off-diagonal blocks that take the value indicated in the block. (b) Multiscale association matrix of (a) that gives the frequency of co-classification of nodes across resolution-parameter values using an NG null network. (c) Multiscale association matrix of (a) that uses a U null network. (d) Toy signed block matrix with constant diagonal and off-diagonal blocks that take the value indicated in the block. (e) Multiscale association matrix of (d) that uses an NGS null network. (f) Multiscale association matrix of (d) that uses a U null network. For the NG and U (respectively, NGS) null networks, our sample of resolution-parameter values is the set $\{\gamma^-, \dots, \gamma^+\}$ (respectively, $\{0, \dots, \gamma^+\}$) with a discretization step of 10^{-3} between each pair of consecutive values.

connectivity, different ones are identified as communities for different values of γ when using the NG null network—as γ increases, nodes in the largest category split into singletons first, followed by those in the second largest category, etc. One would need to be cautious in using multiscale community structure to gain information about connectivity patterns in the observed network in this example.

4.1.2. NGS null network. A key difference between an NG null network (2.8) and an NGS null network (2.13) is that the expected edge weight between two nodes must be positive in the former but can be negative in the latter. Consider a signed variant of the example in Section 4.1.1 in which intra-category edge weights equal a constant $a > 0$ and inter-category edge weights equal a constant $b < 0$. The strengths of node i in the κ^{th} category are

$$k_i^+ = |\kappa|a \quad \text{and} \quad k_i^- = (N - |\kappa|)b.$$

We consider two categories κ_1, κ_2 with different numbers of nodes. Taking $|\kappa_1| > |\kappa_2|$ without loss of generality, it follows that

$$\begin{aligned} P_{i,j \in \kappa_1} &= \frac{1}{2m^+} \left(|\kappa_1|a \right)^2 - \frac{1}{2m^-} \left[(N - |\kappa_1|)b \right]^2 \\ &> \frac{1}{2m^+} \left(|\kappa_2|a \right)^2 - \frac{1}{2m^-} \left[(N - |\kappa_2|)b \right]^2 = P_{i,j \in \kappa_2}, \end{aligned}$$

where $P_{i,j \in \kappa_i}$ is the expected edge weight between pairs of nodes in κ_i in the NGS null network. As was the case for an NG null network, pairs of nodes in an NGS null network that belong to larger categories have a larger expected edge weight than pairs of nodes that belong to smaller categories.

However, the fact that the expected edge weight can be negative can further complicate interpretations of multiscale community structure. A category κ for which $P_{i,j \in \kappa} < 0$ and $P_{i \in \kappa, j \notin \kappa} \geq 0$ is identified as a community when $A_{ij} < -\gamma P_{ij}$ for all $i, j \in \kappa$ (this inequality must hold for sufficiently large γ because $P_{i,j \in \kappa} < 0$) and does not split further for larger values of γ . This poses a particular problem in the interpretation of multiscale community structure obtained with the NGS null network because nodes with negative expected edge weights do not need to be “densely connected” in the observed network to contribute positively to modularity. In fact, if one relaxes the assumption of uniform edge weights across categories, one can ensure that nodes in the category with *lowest* intra-category edge weight will *never* split. This is counterintuitive to standard interpretations of multiscale community structure [36].

In Fig. 4.1(d,e), we illustrate the above feature of the NGS null network using a simple example. The toy network in Fig. 4.1(d) contains 100 nodes divided into three categories: one of size 50 and two of size 25. The category of size 50 and one category of size 25 have an intra-category edge weight of 1 between each pair of nodes. The other category of size 25 has an intra-category edge weight of 0.4 between each pair of nodes. All inter-category edges have weights of -0.05 . (We choose these values so that the intra-category expected edge weight is negative for the third category but positive for the first two and so that inter-category expected edge weights are positive.) We observe in Fig. 4.1(e) that the first and second categories split into singletons for sufficiently large γ , that the smaller of the two categories splits into singletons for a larger value of the resolution parameter, and that the third category never splits. Repeating the same experiment with the U null network in Fig. 4.1(f) (after a linear shift of the adjacency matrix to the interval $[0, 1]$, i.e., $A_{ij} \mapsto \frac{1}{2}(A_{ij} + 1)$ for all i, j), we observe that the co-classification index of nodes reflects the value of the edge weight between them. It is highest for pairs of nodes in the first and second category, and it is lowest for pairs of nodes in the third category.

4.2. Data sets. We show how the features discussed in Section 4.1 can manifest in real data. We use two data sets of financial time series for our computational experiments. The first is a multi-asset data set and consists of weekly price time series for $N = 98$ financial assets during the time period 01 Jan 99–01 Jan 10 (resulting in 574 prices for each asset). The assets are divided into seven asset classes: 20 government bond indices (Gov.), 4 corporate bond indices (Corp.), 28 equity indices (Equ.), 15 currencies (Cur.), 9 metals (Met.), 4 fuel commodities (Fue.), and 18 commodities (Com.). This data set was studied in [22] using principal component analysis and a detailed description of the financial assets can be found in that paper.

The second data set is a single-asset data set that consists of daily price time series for $N = 859$ financial assets from the Standard & Poor’s (S&P) 1500 during the time period 01 Jan 99–01 Jan 13 (resulting in 3673 prices for each asset).⁹ The financial assets are all equities and are divided into ten sectors: 62 materials, 141 industrials, 150 financials, 142 information technology, 55 utilities, 47 consumer staples, 138 consumer discretionary, 48 energy, 68 health care, and 6 telecommunication

⁹We consider fewer than 1500 nodes because we only include nodes for which data is available at all time points to avoid issues associated with choices of data-cleaning techniques.

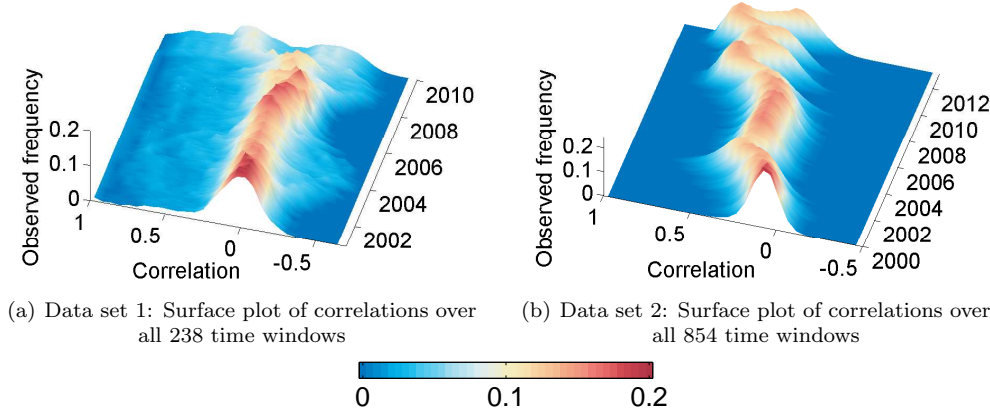


FIG. 4.2. Surface plots of the correlations over all time windows for (a) the first data set and (b) the second data set. The colors in each panel scale with the value of the observed frequency.

services.

The precise way that one chooses to compute a measure of similarity between pairs of time series and the subsequent choices that one makes (e.g., uniform or nonuniform window length, and overlap or no overlap if one uses a *rolling time window*) affect the values of the similarity measure. There are myriad ways to define similarity measures—the best choices depend on facets such as application domain, time-series resolution, and so on—and this is an active and contentious area of research [56, 61, 63, 69]. Constructing a similarity matrix from a set of time series and investigating community structure in a given similarity matrix are separate problems, and we are concerned with the latter in the present paper. Accordingly, in all of our experiments, we use Pearson correlation coefficients for our measure of similarity. We compute them using a rolling time window with a uniform window length and uniform amount of overlap.

We adopt the same network representation for both data sets. We use the term *time window* for a set of discrete time points and divide each time series into overlapping time windows that we denote by $\mathcal{T} = \{T_s\}$. The length of each time window $|T|$ and the amount of overlap between consecutive time windows $|T| - \delta t$ are uniform.¹⁰ We fix $(|T|, \delta t) = (100, 2)$ for the first data set (which amounts to roughly two years of data in each time window) and $(|T|, \delta t) = (260, 4)$ for the second data set (which amounts to roughly one year of data in each time window). Every network layer with adjacency matrix \mathbf{A}_s is a Pearson correlation matrix between the time series of logarithmic returns during the time window T_s . For each data set, we study the sequence of matrices

$$\{\mathbf{A}_s \in [-1, 1]^{N \times N} | s \in \{1, \dots, |\mathcal{T}|\}\}.$$

We show a surface plot of the observed frequency of correlations in each layer for each data set in Fig. (4.2).

¹⁰The amount of overlap determines the number of data points that one adds and removes from each time window. It thus determines the number of data points that can alter the connectivity patterns in each subsequent correlation matrix (i.e., each subsequent layer).




FIG. 4.3. Multiscale association matrix for the U, NG, and NGS null networks for the entire correlation matrix and a subset of the correlation in the last layer of the first data set. In panel (a), we show the entire matrix; in panels (b,c,d), we show the multiscale association matrix that we obtain from this matrix using each of the three null networks. In panel (e), we show the first 35×35 block of the correlation matrix from panel (a); and in panels (f,g,h), we show the multiscale association matrix that we obtain from this subset of the correlation matrix using each of the three null networks. The colors scale with the entries of the multiscale association and the entries of the correlation matrix. Black squares on the diagonals correspond to government and corporate bond assets, and white squares correspond to equity assets.

4.3. Multiscale community structure in asset correlation networks. We perform the same experiments as in Fig. 4.1 on the correlation matrices of both data sets. Our resolution-parameter sample is the set $\{\gamma^-, \dots, \gamma^+\}$ (respectively, $\{0, \dots, \gamma^+\}$) for the U and NG (respectively, NGS) null networks with a discretization step of the order of 10^{-3} . We store the co-classification index of pairs of nodes averaged over all resolution-parameter values in the sample. We use the U and NG null networks for a correlation matrix that is linearly shifted to the interval $[0, 1]$. For each null network, we thereby produce $|\mathcal{T}|$ multiscale association matrices with entries between 0 and 1 that indicate how often pairs of nodes are in the same community across resolution-parameter values.

We show the multiscale association matrices for a specific layer of data set 1 in Fig. 4.3. The matrix in Fig. 4.3(a) corresponds to the correlation matrix during the interval 08 Feb–01 Dec 10. In accord with the results in [22], this matrix reflects the increase in correlation between financial assets that took place after the Lehman bankruptcy in 2008. (One can also see this feature in the surface plot of Fig. 4.2(a).) The matrices in Fig. 4.3(b,c,d) correspond, respectively, to the multiscale association matrix for the U, NG, and NGS null networks. We reorder all matrices (identically) using a node ordering based on the partitions that we obtain with the U null network that emphasizes block-diagonal structure in the correlation matrix. We observe that the co-classification indices in the multiscale association matrix of Fig. 4.3(b) are a

better reflection of the strength of correlation between assets in Fig. 4.3(a) than the multiscale association matrices in Fig. 4.3(c,d). As indicated by the darker shades of red in the upper left corner in Fig. 4.3(c,d), we also observe that the government and corporate bond assets (black squares on the diagonal) remain in the same community across a larger range of resolution-parameter values than the equity assets (white squares on the diagonal). In fact, when we use an NGS null network, the expected weight between two government or corporate bonds is negative (it is roughly -0.1), and these assets remain in the same community across arbitrarily large values of the resolution parameter. One would need to be cautious in using the multiscale association matrices in Fig. 4.3(c,d) to gain insight about the connectivity between assets in Fig. 4.3(a).

When studying correlation matrices of multi-asset data sets, one may wish to vary the size of the asset classes included in the data (e.g., by varying the ratio of equity and bond assets). We show how doing this can lead to further misleading conclusions. By repeating the same experiment using only a subset of the correlation matrix (the first 35 nodes), we consider an example where we have inverted the relative sizes of the bond asset class and the equity asset class. As indicated by the darker shades of red in the lower right corner in Fig. 4.3(g,h), equity assets now have a higher co-classification index than government and corporate bond assets. If one uses the co-classification index in the multiscale association matrices of Fig. 4.3(c,d) (respectively, Fig. 4.3(g,h)) to gain information about the observed correlation between equity and bond assets in Fig. 4.3(a) (respectively, Fig. 4.3(e)), one may draw different conclusions despite the fact that these have not changed. However, the multiscale association matrix with a U null network in Fig. 4.3(f) reflects the observed correlation between equity and bond assets in Fig. 4.3(e).¹¹

To quantify the sense in which a multiscale association matrix of one null network “reflects” the values in the correlation matrix, we compute the Pearson correlation between the upper triangular part of each multiscale association matrix and its corresponding adjacency matrix across all time layers of both data sets for the U, NG, and NGS null networks. We show these correlation plots in Fig. 4.4. Observe that the correlation between the adjacency and multiscale association matrix in Fig. 4.4(a,b) is highest in each layer for the U null network and lowest in (almost) each layer for the NGS null network.

The above observation can be explained as follows. Recall from (2.5) that we can write the modularity-maximization problem as $\max_{S \in \mathcal{S}} \text{Tr}(\mathbf{S}^T \mathbf{B} \mathbf{S})$, where \mathcal{S} is the set of partition matrices. When one uses a U null network, the entries of the modularity matrix are the entries of the adjacency matrix shifted by a constant $\gamma \langle \mathbf{A} \rangle$, and the quality function reduces to

$$\max_{S \in \mathcal{S}} [\text{Tr}(\mathbf{S}^T \mathbf{A} \mathbf{S}) - \gamma \langle \mathbf{A} \rangle \|c(\mathbf{S})\|_2] , \quad (4.2)$$

where $\|c(\mathbf{S})\|_2 = \|\text{Tr}(\mathbf{S}^T \mathbf{1}_N \mathbf{S})\|_2$ is the 2-norm of the vector of set sizes in \mathbf{S} (i.e., $c(\mathbf{S})$ is the vector whose k^{th} entry is $\sum_{i=1}^N S_{ik}$). It follows that modularity maximization

¹¹The authors of [65] showed that a globally optimal partition for a null network called the “constants Potts model” (CPM), in which the edge weights are given by a constant that is *independent* of the network, is “sample-independent”. Their result can be generalized as follows for the U null network (in which expected edge weights are constant but are not independent of the observed network). Suppose that C_{\max} is a partition that maximizes $Q(C|\mathbf{A}; \mathbf{P}; \gamma_1)$ and consider the subgraph induced by the network on a set of communities $C_1, \dots, C_l \in C_{\max}$. Then $\{C_1 \cup C_2 \dots \cup C_l\}$ maximizes $Q(C|\hat{\mathbf{A}}; \mathbf{P}; \gamma_2)$, where $\hat{\mathbf{A}}$ is the adjacency matrix of the induced subgraph and $\gamma_2 = \gamma_1 \langle \mathbf{A} \rangle / \langle \hat{\mathbf{A}} \rangle$. For the CPM null network, the same result holds with $\gamma_1 = \gamma_2$.




FIG. 4.4. Correlation between the adjacency matrix and the multiscale association matrix for the U (solid curve), NG (dashed curve), and NGS (dotted curve) null networks over all time layers for (a) data set 1 and (b) data set 2. We compute the Pearson correlation coefficients between entries in the upper diagonal blocks in each matrix (to avoid double counting, as the matrices are symmetric), and we exclude diagonal entries (which, by construction, are equal to 1 in both matrices).

with a U null network is equivalent to a block-diagonalization of the adjacency matrix \mathbf{A} (the first term in (4.2)) with a penalty on the size of communities (the second term). As one increases the resolution parameter, one favors smaller sets of nodes with stronger internal connectivity. Note that one could also apply equation (4.2) on adjusted adjacency matrices $\mathbf{A}' = \mathbf{A} - \tilde{\mathbf{A}}$. For example, one can let $\tilde{\mathbf{A}}$ be a matrix that controls for random fluctuations in a correlation matrix \mathbf{A} (e.g., the “random component” \mathbf{C}^r in [41]).

For a general null network, equation (4.2) takes the form

$$\max_{S \in \mathcal{S}} [\text{Tr}(\mathbf{S}^T \mathbf{A} \mathbf{S}) - \text{Tr}(\mathbf{S}^T (\gamma \mathbf{P}) \mathbf{S})] ,$$

where \mathbf{P} is the adjacency matrix of the null network. That is, modularity maximization finds block-diagonal structure in \mathbf{A} (first term) that is not in $\gamma \mathbf{P}$ (second term). It is common to avoid using the U null network in applications because “it is not a good representation of most real-world networks” [47]. The extent to which one wants a null network to be a good representation of an observed network depends on the features in an application for which one wants to control. We argue that whether an NG null network is more appropriate than a U null network for a given situation depends at least in part on one’s interpretation of node strength for that application. As we discussed in Section 2.4, the strength of a node in correlation matrices is given by the covariance between its standardized time series and the mean time series. When using the NG null network, it thus follows that pairwise differences $B_{ij} - B_{i'j'}$ in the modularity quality function depend on $\text{corr}(\hat{z}_i, \hat{z}_j)$, $\text{corr}(\hat{z}_{i'}, \hat{z}_{j'})$, and $\text{corr}(\hat{z}_k, \hat{z}_{\text{tot}})$, where $k \in \{i, j, i', j'\}$, the quantity \hat{z}_i is the standardized time series of asset i defined in subsection 2.4.1, and $\hat{z}_{\text{tot}} = \sum_{i=1}^N \hat{z}_i$. When using the U null network, pairwise differences in the modularity quality function depend only on the observed edge weights $\text{corr}(\hat{z}_i, \hat{z}_j)$ and $\text{corr}(\hat{z}_{i'}, \hat{z}_{j'})$. In contrast, the term $\text{corr}(\hat{z}_k, \hat{z}_{\text{tot}})$ introduces a dependence between the communities that one finds with the NG null network and the extent to which nodes in those communities are representative [as measured by $\text{corr}(\hat{z}_k, \hat{z}_{\text{tot}})$] of the mean time series for the sample. In situations in which one may wish to vary one’s node sample (e.g., by changing the size of asset classes), one

needs to bear such dependencies in mind when interpreting the communities that one obtains.

5. Effect of inter-layer coupling on a multilayer partition. In Section 4, we set the inter-layer connection weights to 0 in the multilayer network. The solution to the multilayer modularity-maximization problem (3.4) then depends solely on the values in the modularity matrix of each time layer, and the multilayer modularity-maximization problem reduces to performing single-layer modularity maximization on each layer independently.

Recall the multilayer modularity-maximization problem

$$\max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \sum_{s=1}^{|\mathcal{T}|-1} \sum_{i=1}^N \delta(c_{i_s}, c_{i_{s+1}}).$$

A solution to this problem is a partition of an $N|\mathcal{T}|$ -node multilayer network. Its communities can contain nodes from the same layer and nodes from different layers. Nodes from different layers can be the same node at different times $((i_s, i_r)$ with $s \neq r$) or different nodes at different times $((i_s, j_r)$ with $i \neq j$ and $s \neq r$). We say that a node i_s remains in the same community (respectively, changes communities) between consecutive layers s and $s+1$ if $\delta(c_{i_s}, c_{i_{s+1}}) = 1$ (respectively, $\delta(c_{i_s}, c_{i_{s+1}}) = 0$).

Positive ordinal, diagonal, and uniform inter-layer connections favor nodes remaining in the same community between consecutive layers. Every time a node does not change communities between two consecutive layers (i.e., $\delta(c_{i_s}, c_{i_{s+1}}) = 1$), a positive contribution of 2ω is added to the multilayer quality function. One thereby favors communities that do not to change in time because community assignments are transitive: if $\delta(c_{i_s}, c_{j_s}) = 1$ and $\delta(c_{i_s}, c_{i_{s+1}}) = \delta(c_{j_s}, c_{j_{s+1}}) = 1$, then $\delta(c_{i_{s+1}}, c_{j_{s+1}}) = 1$. We define the *persistence* of a multilayer partition to be the total number of nodes that do not change communities between layers:

$$\text{Pers}(C) := \sum_{s=1}^{|\mathcal{T}|-1} \sum_{i=1}^N \delta(c_{i_s}, c_{i_{s+1}}) \in \{0, \dots, N(|\mathcal{T}| - 1)\}. \quad (5.1)$$

As indicated in equation (5.1), $\text{Pers}(C)$ is an integer between 0, which occurs when no node ever remains in the same community across layers, and $N(|\mathcal{T}| - 1)$, which occurs when every node always remains in the same community. (See [7] for a closely related measure called “flexibility” that has been applied to functional brain networks.) Let $\text{Pers}(C)|_s$ denote the number of nodes that remain in the same community between two consecutive layers s and $s+1$:

$$\text{Pers}(C)|_s := \sum_{i=1}^N \delta(c_{i_s}, c_{i_{s+1}}) \in \{0, \dots, N\}, \quad (5.2)$$

so that $\text{Pers}(C) = \sum_{s=1}^{|\mathcal{T}|-1} \text{Pers}(C)|_s$. Persistence provides an insightful way of rewriting the multilayer modularity-maximization problem:

$$\max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \text{Pers}(C). \quad (5.3)$$

The multilayer maximization problem thus measures a trade-off between static community structure within layers (the first term in (5.3)) and temporal persistence across layers (the second term in (5.3)).

To better understand the effect of nonzero inter-layer coupling on partitions one obtains without inter-layer coupling (i.e., $\omega = 0$), we introduce notation that helps to compare a multilayer partition to single-layer partitions. We denote by $\mathcal{N}_s := \{1_s, \dots, N_s\}$ the set of nodes in layer s . The restriction of a set of nodes $C_l \subseteq \{1_1, \dots, N_1; 1_2, \dots, N_2; \dots; 1_{|\mathcal{T}|}, \dots, N_{|\mathcal{T}|}\}$ to a layer s is $C_l|_s := C_l \cap \mathcal{N}_s$, and we define the *partition induced by a multilayer partition* $C \in \mathcal{C}$ on layer s by

$$C|_s := \{C_l|_s, C_l \in C\}.$$

We refer to a “globally optimal partition” as an “optimal partition” in this section for ease of writing. In the next two subsections, we illustrate how the set of partitions induced by a multilayer partition with $\omega > 0$ on individual layers can differ from intra-layer partitions obtained with $\omega = 0$.

5.1. Toy examples.

5.1.1. Changes in connectivity patterns. This toy examples illustrates how inter-layer coupling can enable us to detect and differentiate between changes in connectivity patterns across layers. In Fig. 5.1, we show an unweighted multilayer network with $|\mathcal{T}| = 10$ layers and $N = 8$ nodes in each layer. Every layer except for layers 3 and 6 contains two 4-node cliques. In layer 3, node 5_3 is connected to nodes $\{1_3, 2_3\}$ instead of nodes $\{6_3, 7_3, 8_3\}$. In layer 6, node 5_6 is connected to nodes $\{1_6, 2_6, 3_6, 4_6\}$ instead of nodes $\{6_6, 7_6, 8_6\}$. We show the layers of the multilayer network in panels (a)–(c) of Fig. 5.1. We examine its communities using a U null network with a resolution-parameter $\gamma = 1$. Layer s then has the following single-layer modularity matrix:

$$B_{ijs} = \begin{cases} 1 - \langle \mathbf{A}_s \rangle, & \text{if } i \text{ is connected to } j \\ -\langle \mathbf{A}_s \rangle, & \text{otherwise.} \end{cases}$$

The optimal partition in each layer is unique and is $C_s = \{\{1_s, 2_s, 3_s, 4_s\}, \{5_s, 6_s, 7_s, 8_s\}\}$ in layer s for $s \notin \{3, 6\}$ and is $C_s = \{\{1_s, 2_s, 3_s, 4_s, 5_s\}, \{6_s, 7_s, 8_s\}\}$ in layers 3 and 6. When the value of inter-layer coupling is 0, the optimal multilayer partition is the union of $|\mathcal{T}|$ disconnected optimal single-layer partitions. The resulting multilayer partition, which we show in panel (d) of Fig. 5.1, has a persistence $\text{Pers}(C) = 0$. We denote this partition by \mathcal{C}_0 , where $\mathcal{C}_0 = \bigcup_{i=1}^{10} C_s$. For any $\omega > 0$, any partition with the same intra-layer partitions as \mathcal{C}_0 and a nonzero value of persistence yields a higher value of multilayer modularity than \mathcal{C}_0 . This follows immediately from the expression of the multilayer quality function:

$$Q(C|\mathcal{B}) = \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \text{Pers}(C).$$

Increasing persistence without changing intra-layer partitions increases the second term of $Q(C|\mathcal{B})$ without changing the first. (In Section 5.2, we prove that $\omega > 0$ is both necessary and sufficient for an optimal partition to have a positive value of persistence.) To obtain the multilayer partition in panel (e), we combine all of the sets in panel (d) that contain 1_s into one set and all of the sets that contain N_s into another set. This partition has a persistence equal to $N(|\mathcal{T}| - 1) - 4$, and any other way of combining the sets in \mathcal{C}_0 yields a lower value of persistence.

Let’s examine Fig. 5.1 further. We now consider the multilayer partitions in panel (e), where both changes in network structure occurs; panel (f), where only

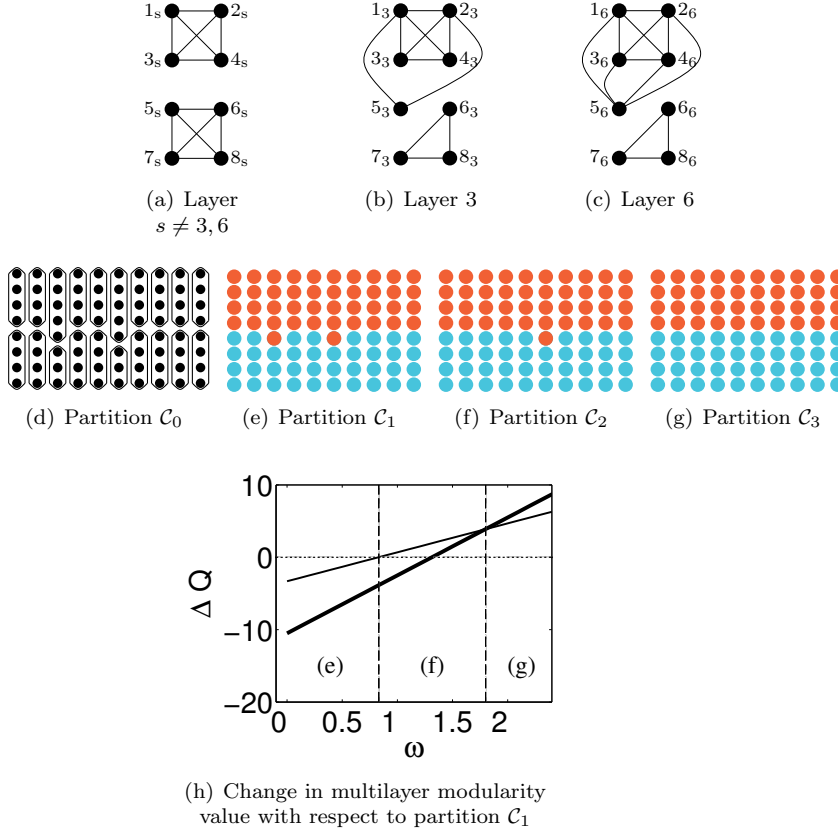


FIG. 5.1. Toy example illustrating the use of ordinal diagonal and uniform inter-layer coupling for detecting changes in community structure across layers. We consider ten layers ($|\mathcal{T}| = 10$) with eight nodes ($N = 8$) in each layer. We show the network structures in (a) layers $s \notin \{3, 6\}$, (b) layer 3, and (c) layer 6. Panels (d)–(g) illustrate four different multilayer partitions. In each panel, the s^{th} column of circles represents the nodes in the s^{th} layer, which we order from 1 to 8. We show sets of nodes in the same community using solid curves in panel (d) (to avoid having to use 20 distinct colors) and using colors in panels (e)–(g). In panel (h), we show the difference between the multilayer modularity value between the partition in panels (f) (thin line) and (g) (thick line) and the partition in panel (e) for different values of ω . We include the horizontal dotted line to show the point at which the thin line intercepts the horizontal axis. The panel labels in the regions defined by the area between two consecutive vertical lines in panel (h) indicate which of the multilayer partitions in panels (e), (f), and (g) has a higher value of modularity.

the stronger change occurs; and panel (g), where neither change occurs. We denote these multilayer partitions by \mathcal{C}_1 , \mathcal{C}_2 , and \mathcal{C}_3 , respectively, and we note that $\text{Pers}(\mathcal{C}_1) < \text{Pers}(\mathcal{C}_2) < \text{Pers}(\mathcal{C}_3)$. The value ω of inter-layer coupling determines which partition of these three has the highest value of multilayer modularity. To see this, we compute the modularity cost of changing static community structure within layers in partition \mathcal{C}_1 in favor of persistence. (Such a computation is a multilayer version of the calculations for static networks in [28].) The intra-layer modularity cost in \mathcal{C}_1 of moving node 5_s from

the community $\{1_s, 2_s, 3_s, 4_s, 5_s\}$ to the community $\{6_s, 7_s, 8_s\}$ in layers $s \in \{3, 6\}$ is

$$\begin{aligned}\Delta Q(s) &= 2 \left(\sum_{j \in \{6, 7, 8\}} B_{5js} - \sum_{j \in \{1, 2, 3, 4\}} B_{5js} \right) \\ &= \begin{cases} -4 + 2\langle \mathbf{A} \rangle_3 \approx -3.3, & \text{if } s = 3 \\ -8 + 2\langle \mathbf{A} \rangle_6 \approx -7.2, & \text{if } s = 6. \end{cases}\end{aligned}$$

The inter-layer modularity cost from this move is $+4\omega$ in both cases; the first factor of 2 follows by symmetry of \mathcal{B} , and the second factor of 2 follows from the fact that either move increases persistence by $+2$. Consequently, for $0 < 4\omega < \Delta Q(3)$, the partition in panel (e) yields a higher multilayer modularity value than the partitions in (f) and (g). When $\Delta Q(3) < 4\omega < \Delta Q(6)$, the multilayer modularity value of the partition in (f) is higher than those of (e) or (g). Finally, when $4\omega > \Delta Q(6)$, the partition in panel (g) has the highest multilayer modularity value. When $4\omega = \Delta Q(3)$ (respectively, $4\omega = \Delta Q(6)$), the multilayer partition in panels (e) and (f) (respectively, (f) and (g)) have the same value of multilayer modularity. We illustrate these results in Fig. 5.1(h) by plotting $Q(\mathcal{C}_2|\mathcal{B}) - Q(\mathcal{C}_1|\mathcal{B})$ and $Q(\mathcal{C}_3|\mathcal{B}) - Q(\mathcal{C}_1|\mathcal{B})$ against ω . This example is a simple illustration of how inter-layer connections can help distinguish between changes in connectivity patterns: stronger changes (in terms of modularity cost) persist across larger values of inter-layer coupling (see [4, 53] for other approaches to “change point detection” in temporal networks).

5.1.2. Shared connectivity patterns. In the previous toy example, the intra-layer partitions induced on each layer by the multilayer partitions in Fig. 5.1(e,f,g) are optimal for at least one layer when $\omega = 0$ (see Fig. 5.1(d)). This second example illustrates how inter-layer coupling can identify intra-layer partitions that are not optimal for any individual layer when $\omega = 0$ but which reflect connectivity patterns that are shared across layers.

In Fig. 5.2, we consider an unweighted multilayer network with $|\mathcal{T}| = 3$ layers and $N = 13$ nodes in each layer. Every s^{th} layer contains four 3-node cliques and a node that is connected to each of the three nodes in the s^{th} clique, and to nodes 10_s and 12_s in the 4th clique. We show the layers of the multilayer network in panels (a)–(c). We examine its communities using a U null network with a resolution-parameter value of $\gamma = 1$. The optimal partition in each layer is unique and is $\{\{1_1, 2_1, 3_1, 13_1\}, \{4_1, 5_1, 6_1\}, \{7_1, 8_1, 9_1\}, \{10_1, 11_1, 12_1\}\}$ for layer 1, $\{\{1_2, 2_2, 3_2\}, \{4_2, 5_2, 6_2, 13_2\}, \{7_2, 8_2, 9_2\}, \{10_2, 11_2, 12_2\}\}$ for layer 2, and $\{\{1_3, 2_3, 3_3\}, \{4_3, 5_3, 6_3\}, \{7_3, 8_3, 9_3, 13_3\}, \{10_3, 11_3, 12_3\}\}$ for layer 3. We obtain the multilayer partition \mathcal{C}_1 in panel (d) by combining these sets such that induced intra-layer partitions are optimal for each layer when $\omega = 0$ and persistence is maximized between layers. The multilayer partition \mathcal{C}_2 in panel (e) reflects connectivity patterns that are shared by all layers (i.e., node 13_s is with the fourth 3-node clique instead of the s^{th} 3-node clique); but its intra-layer partitions are not optimal for any layer when $\omega = 0$. By carrying out similar calculations to those in the previous toy example, one can show that when $\omega > 3/2$,¹² the multilayer partition in panel (e) yields a higher modularity value than the multilayer partition in panel (d). We illustrate this result in Fig. 5.2(f) by plotting $Q(\mathcal{C}_2|\mathcal{B}) - Q(\mathcal{C}_1|\mathcal{B})$ against ω . This example is a simple illustration of how inter-layer connections can help identify connectivity patterns that are shared across

¹²i.e., when $4\omega + 6[2(1 - \langle \mathbf{A} \rangle_s) - \langle \mathbf{A} \rangle_s - 3(1 - \langle \mathbf{A} \rangle_s)] > 0$, with $\langle \mathbf{A} \rangle_1 = \langle \mathbf{A} \rangle_2 = \langle \mathbf{A} \rangle_3$ by construction in this example.

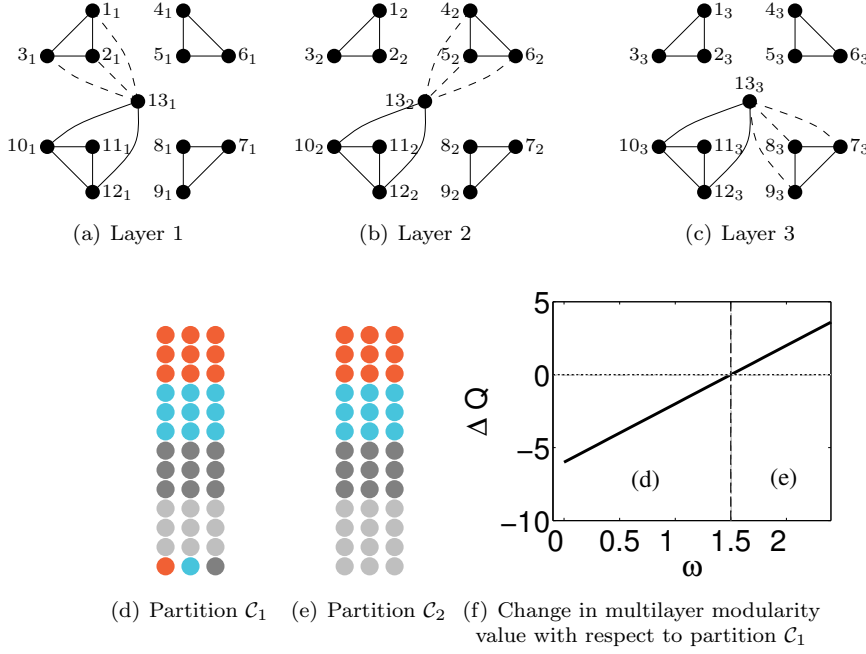


FIG. 5.2. *Toy example illustrating the use of ordinal diagonal and uniform inter-layer coupling for detecting shared connectivity patterns across layers. We consider three layers ($|\mathcal{T}| = 3$) with thirteen nodes ($N = 13$) in each layer. We show the network structures in (a) layer 1, (b) layer 2, and (c) layer 3. Solid lines represent edges present in all three layers and dashed lines represent edges that are only present in one of the layers. Panels (d) and (e) illustrate two different multilayer partitions. In each panel, the s^{th} column of circles represents the nodes in the s^{th} layer, which we order 1 to 13. We show sets of nodes in the same community using colors in panels (d) and (e). In panel (f), we show the difference between the multilayer modularity value between the partition in panel (e) and the partition in panel (d) for different values of ω . We include the horizontal dotted line to show the point at which the line intercepts the horizontal axis. The panel labels in the regions defined by the area between two consecutive vertical lines in panel (f) indicate which of the multilayer partitions in panels (d) and (e) has a higher value of multilayer modularity.*

layers.

5.2. Some properties of multilayer partitions. We now ask how introducing positive ordinal diagonal and uniform coupling (i.e., $\omega > 0$) alters the set of maximum-modularity partitions of static networks (i.e., the case $\omega = 0$). To clearly differentiate between intra-layer and inter-layer modularity contributions, we denote the quality function by

$$Q(C|\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega) := \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \text{Pers}(C)$$

instead of $Q(C|\mathbf{B})$. Let $\mathcal{C}_{\max}(\omega)$ denote the set of optimal partitions for the multilayer modularity-maximization problem (5.3), and let C_{\max}^ω be an arbitrary partition in $\mathcal{C}_{\max}(\omega)$. We prove several propositions that hold for an arbitrary choice of the matrices \mathbf{B}_s (for example, if one uses the modularity quality function with a U null

network and a resolution parameter value of 1, then $\mathbf{B}_s = \mathbf{A}_s - \langle \mathbf{A}_s \rangle \mathbf{1}_N$.

PROPOSITION 5.1. $\text{Pers}(C_{\max}^\omega) > 0 \Leftrightarrow \omega > 0$.

Proposition 5.1 ensures that as soon as (and only when) the value of ω is strictly positive, the value of persistence of an optimal solution is also positive. To prove this, it suffices to observe that if one rearranges sets in a multilayer partition by combining some of the sets into the same set *without* changing the partitions induced on individual layers, then one only changes the value of persistence in the expression of multilayer modularity. For example, this phenomenon occurs in Fig. 5.1 when going from the partition in panel (d) to the partition in panel (e).

Proof.

\Rightarrow : We prove the contrapositive. Assume that $\omega = 0$ and consider a multilayer partition C such that $\text{Pers}(C) > 0$. The partition C contains at least one set with disconnected components (because $\text{Pers}(C) > 0$ and nodes in different layers are not connected), and C is not optimal by our assumption on global optima in Section 3.

\Leftarrow : Assume that $\omega > 0$ and consider a multilayer partition C such that $\text{Pers}(C) = 0$. We will show that C is not optimal. Let $C' = \bigcup_{s=1}^{|\mathcal{T}|} C|_s$. It then follows that

$$Q(C'|B_1, \dots, B_{|\mathcal{T}|}; \omega) = Q(C|B_1, \dots, B_{|\mathcal{T}|}; \omega).$$

Choose a node i_r at random and let C_{i_r} denote the set in C' that contains i_r . Now let C'' be the partition obtained from C' by combining all sets that contain i_s , for some s , into one set:

$$C'' = \left(C' \setminus \bigcup_{s=1}^{|\mathcal{T}|} \{C_{i_s}\} \right) \cup \left\{ \bigcup_{s=1}^{|\mathcal{T}|} C_{i_s} \right\}.$$

Consequently,

$$Q(C''|B_1, \dots, B_{|\mathcal{T}|}; \omega) \geq Q(C|B_1, \dots, B_{|\mathcal{T}|}; \omega) + 2\omega(|\mathcal{T}| - 1),$$

so C is not optimal. (Note that $|\mathcal{T}| \geq 2$ for a network with more than a single layer, so $2\omega(|\mathcal{T}| - 1)$ is strictly positive for $\omega > 0$.) \square

PROPOSITION 5.2.

If $C_l|_r = \emptyset$ for some $r \in \{1, \dots, |\mathcal{T}| - 1\}$, then $C_l|_s = \emptyset$ for all $s > r$,

where $C_l \in C_{\max}^\omega$ and $C_{\max}^\omega \in \mathcal{C}_{\max}(\omega)$.

Proposition 5.2 ensures that if a community becomes empty in a given layer, then it remains empty in all subsequent layers. We omit the proof as this result follows directly from the sparsity pattern of \mathbf{B} and our assumption that optimal solutions do not contain disconnected components.

PROPOSITION 5.3. $C_{\max}^\omega|_s = C_{\max}^\omega|_{s+1} \Leftrightarrow \text{Pers}(C_{\max}^\omega)|_s = N$.

Proposition 5.3 connects the notion of persistence between a pair of layers to the notion of change in community structure within layers. Various numerical experiments that have been performed with ordinal diagonal and uniform inter-layer coupling

consist of varying the value of ω and using information about when nodes change communities between layers as an indication of change in community structure within these layers [6, 7, 44]. The equivalence relation in Proposition 5.3 motivates the use of $\text{Pers}(C)|_s$ (or a variant thereof) as an indication of intra-layer change in community structure.

Proof.

\Leftarrow : This follows straightforwardly by transitivity of community assignments: if $\delta(c_{j_s}, c_{j_{s+1}}) = \delta(c_{i_s}, c_{i_{s+1}}) = 1$ for all i, j , then $\delta(c_{i_s}, c_{j_s}) = 1$ if and only if $\delta(c_{i_{s+1}}, c_{j_{s+1}}) = 1$ for all i, j . (This direction holds for *any* multilayer partition; it need not be optimal.)

\Rightarrow : Let C be a multilayer partition that does not contain disconnected components such that $C|_s = C|_{s+1}$ and $\text{Pers}(C)|_s < N$ for some $s \in \{1, \dots, |\mathcal{T}|\}$. We show that C is not optimal.¹³ Consider a set $C_l \in C$ such that $C_l|_s \neq \emptyset$. If $\delta(c_{i_s}, c_{i_{s+1}}) = 1$ (respectively, $\delta(c_{i_s}, c_{i_{s+1}}) = 0$) for some $i_s \in C_l|_s$, then $\delta(c_{j_s}, c_{j_{s+1}}) = 1$ (respectively, $\delta(c_{j_s}, c_{j_{s+1}}) = 0$) for all $j_s \in C_l|_s$ by transitivity of community assignments and because $C|_s = C|_{s+1}$ by hypothesis. Because $\text{Pers}(C)|_s < N$ by hypothesis, there exists at least one set of nodes $C_l|_s$, with $C_l \in C$, such that $\delta(c_{i_s}, c_{i_{s+1}}) = 0$ for all $i_s \in C_l|_s$. Let $C_k|_{s+1}$ denote the set of nodes in layer $s+1$ that contains i_{s+1} for all $i_s \in C_l|_s$. Consider the set of nodes $\cup_{r \leq s} C_l|_r$ in C_l that are in layers $\{1, \dots, s\}$ and the set of nodes $\cup_{r > s} C_k|_r$ in C_k that are in layers $\{s+1, \dots, |\mathcal{T}|\}$. Because $\delta(c_{i_s}, c_{i_{s+1}}) = 0$ for all $i_s \in C_l|_s$ and by Proposition 5.2, it follows that $C_l = \cup_{r \leq s} C_l|_r$ and $C_k = \cup_{r > s} C_k|_r$. Define the partition C' by

$$C' = \left(C \setminus (\{C_l\} \cup \{C_k\}) \right) \cup \left(\{C_l \cup C_k\} \right).$$

This partition satisfies $C'|_r = C|_r$ for all $r \in \{1, \dots, |\mathcal{T}|\}$, $\text{Pers}(C')|_r = \text{Pers}(C)|_r$ for all $r \neq s$, and $\text{Pers}(C')|_s > \text{Pers}(C)|_s$. It follows that $Q(C'|_{\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}}; \omega) > Q(C|_{\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}}; \omega)$ and C is not optimal. \square

Propositions 5.1, 5.2, and 5.3 apply to an optimal partition obtained with any positive value of ω . The next two propositions concern the existence of “boundary” values for ω .

PROPOSITION 5.4. *There exists $\omega_0 > 0$ such that*

$$\text{if } \omega < \omega_0, \text{ then } \bigcup_{s=1}^{|\mathcal{T}|} C_{\max}^\omega|_s \in \mathcal{C}_{\max}(0).$$

Proposition 5.4 reinforces the idea of thinking of ω as the cost of breaking static community structure within layers in favor of larger values of persistence across layers. It demonstrates that there is a positive value of inter-layer coupling such that for any smaller coupling, multilayer modularity maximization only gives more information than single-layer modularity maximization in that it identifies the set of partitions in $\mathcal{C}_{\max}(0)$ with highest persistence. The proof of this property relies on the fact that the set of possible modularity values for a given modularity matrix is finite.

Proof.

Let C be an arbitrary partition such that $\cup_{s=1}^{|\mathcal{T}|} C|_s \notin \mathcal{C}_{\max}(0)$. We will show that there

¹³Imposing $\text{Pers}(C)|_s = N$ is not sufficient because changing $\text{Pers}(C)|_s$ can change partitions induced on individual layers.

exists a value ω_0 of the inter-layer coupling parameter ω such that C is never optimal for any inter-layer coupling below ω_0 . Given a sequence of single-layer modularity matrices $\{\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}\}$, the set of possible multilayer modularity values for a fixed value of $\omega > 0$ is finite and is given by

$$\mathcal{Q}_\omega = \{Q(C|\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega), C \in \mathcal{C}\}.$$

Let $Q_0^1 = \max \mathcal{Q}_0$, $Q_0^2 = \max \mathcal{Q}_0 \setminus \{Q_0^1\}$, and $\Delta Q = Q_0^1 - Q_0^2 > 0$. By hypothesis,

$$Q(C|\mathbf{B}_1, \dots, \mathbf{B}_1; 0) < Q(C_{\max}^0|\mathbf{B}_1, \dots, \mathbf{B}_1; 0),$$

where $C_{\max}^0 \in \mathcal{C}_{\max}(0)$. Furthermore, by definition of persistence, it follows that

$$Q(C|\mathbf{B}_1, \dots, \mathbf{B}_1; \omega) \leq Q_0^2 + 2\omega N(|\mathcal{T}| - 1)$$

for all values of ω . By choosing $\omega < \omega_0$, with $\omega_0 = \Delta Q / 2N(|\mathcal{T}| - 1)$, we obtain

$$Q(C|\mathbf{B}_1, \dots, \mathbf{B}_1; \omega) \leq Q_0^2 + 2\omega N(|\mathcal{T}| - 1) < Q_0^2 + \Delta Q = Q_1^0,$$

so C is not optimal for any inter-layer coupling below ω_0 . \square

Clearly, $\omega_0 = \Delta Q / 2N(|\mathcal{T}| - 1)$ is not an upper bound for the set $\{\omega \in \mathbb{R}^+ : \bigcup_{s=1}^{|\mathcal{T}|} C_{\max|s}^\omega \in \mathcal{C}_{\max}(0)\}$,¹⁴ but our main concern is that the smallest upper bound of this set is *not* zero (in fact, we have shown that it must be larger than $\Delta Q / 2N(|\mathcal{T}| - 1) > 0$.)

PROPOSITION 5.5. *There exists $\omega_\infty > 0$ such that*

$$\text{if } \omega > \omega_\infty, \text{ then } \text{Pers}(C_{\max}^\omega)|_s = N \text{ for all } s \in \{1, \dots, |\mathcal{T}|\}.$$

Proposition 5.5 implies that a sufficiently large value of inter-layer coupling ω guarantees that $C_{\max|s}^\omega$ remains the same across layers (by Proposition 5.3). The proof of this proposition is similar to the proof of Proposition 5.4.

Proof.

Let C be an arbitrary partition of a multilayer network such that $\text{Pers}(C)|_s < N$ for some $s \in \{1, \dots, |\mathcal{T}|\}$. We show that there exists a value $\omega_\infty > 0$ of the inter-layer coupling parameter ω such that C is never optimal for $\omega > \omega_\infty$. We first rewrite the quality function as

$$Q(C|\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega) = \beta_1 + 2\omega(N(|\mathcal{T}| - 1) - A),$$

where $\beta_1 = \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s})$ and $A \geq 1$ because $\text{Pers}(C) < N(|\mathcal{T}| - 1)$ by assumption. Now consider the set of values on the diagonal blocks of the multilayer modularity matrix \mathcal{B} :

$$\mathcal{B}_{\text{diag}} = \{B_{ijs}|i, j \in \{1, \dots, N\}, s \in \{1, \dots, |\mathcal{T}|\}\}, \quad (5.4)$$

¹⁴For example, one could replace $N(T - 1)$ in (5.4) by $N(|\mathcal{T}| - 1) - \text{Pers}(C(0))$, where $\text{Pers}(C(0))$ denotes the maximum value of persistence that one can obtain by combining sets in each partition of $\mathcal{C}_{\max}(0)$ without changing the partitions induced on individual layers. Taking $\omega_0 = \Delta Q / [2N(|\mathcal{T}| - 1) - \text{Pers}(C(0))]$ satisfies proposition 5.4, and $\Delta Q / [2N(|\mathcal{T}| - 1) - \text{Pers}(C(0))] > \Delta Q / 2N(|\mathcal{T}| - 1)$.

and let $\text{Max}(\mathcal{B}_{\text{diag}})$ and $\text{Min}(\mathcal{B}_{\text{diag}})$, respectively, denote the maximum and minimum values of the set $\mathcal{B}_{\text{diag}}$. Let C' be any multilayer partition with a maximal value of persistence. It then follows that

$$Q(C'|B_1, \dots, B_{|\mathcal{T}|}; \omega) = \beta_2 + 2\omega N(|\mathcal{T}| - 1)$$

for some $\beta_2 \in \mathbb{R}$. Because $A \geq 1$, choosing

$$2\omega > N^2[\text{Max}(\mathcal{B}_{\text{diag}}) - \text{Min}(\mathcal{B}_{\text{diag}})] \geq \beta_1 - \beta_2$$

ensures that C' yields a higher value of multilayer modularity than C for any β_1 and for all $A \in \{1, \dots, N(|\mathcal{T}| - 1)\}$. \square

The following proposition follows directly from proposition 5.5.

PROPOSITION 5.6. *There exists $\omega_\infty > 0$ such that*

$$C_{\max|s}^\omega \text{ is a solution of } \max_{C \in \mathcal{C}} Q\left(C \middle| \sum_{s=1}^{|\mathcal{T}|} B_s\right)$$

for all $\omega > \omega_\infty$.

Propositions 5.5 and 5.6 imply the existence of a “boundary value” for ω above which single-layer partitions induced by optimal multilayer partitions (1) are the same on all layers and (2) are optimal solutions for the single-layer modularity-maximization problem defined on the mean modularity matrix.

Proof.

Suppose that $\omega > \omega_\infty$, where ω_∞ is as defined in Proposition 5.5, and let $C_{\max}^\omega \in \mathcal{C}_{\max}(\omega)$. By Proposition 5.5, it then follows that $\text{Pers}(C_{\max}^\omega) = N(|\mathcal{T}| - 1)$ and community assignments in C_{\max}^ω are the same across layers. Consequently, for $\omega > \omega_\infty$

$$\begin{aligned} & \max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \text{Pers}(C) \\ & \Leftrightarrow \max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_i, c_j) + 2\omega N(|\mathcal{T}| - 1) \\ & \Leftrightarrow \max_{C \in \mathcal{C}} \sum_{i,j}^N \left(\sum_{s=1}^{|\mathcal{T}|} B_{ijs} \right) \delta(c_i, c_j), \end{aligned}$$

where c_i denotes the community assignment of node i in all layers. \square

The next two propositions formalize the intuition that an optimal multilayer partition measures a trade-off between static community structure within layers and persistence of community structure across layers.

PROPOSITION 5.7. *Let $\omega_1 > \omega_2 > 0$. For all $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_2)$, one of the following two conditions must hold:*

- (1) $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_1)$,
- or (2) $\text{Pers}(C_{\max}^{\omega_2}) < \text{Pers}(C_{\max}^{\omega_1})$ for all $C_{\max}^{\omega_1} \in \mathcal{C}_{\max}(\omega_1)$.

Proof.

Let $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_2)$. If $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_1)$, then condition (1) is satisfied. Suppose that $C_{\max}^{\omega_2} \notin \mathcal{C}_{\max}(\omega_1)$, and assume that $\text{Pers}(C_{\max}^{\omega_2}) \geq \text{Pers}(C_{\max}^{\omega_1})$ for some $C_{\max}^{\omega_1} \in \mathcal{C}_{\max}(\omega_1)$. By definition of optimality, $C_{\max}^{\omega_2} \notin \mathcal{C}_{\max}(\omega_1)$ implies that

$$Q(C_{\max}^{\omega_2} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega_1) < Q(C_{\max}^{\omega_1} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega_1), \quad (5.5)$$

where $\omega_1 > \omega_2$ by hypothesis. By writing

$$Q(C_{\max}^{\omega_k} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega_{k'}) = \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}^{\omega_k}, c_{j_s}^{\omega_{k'}}) + 2\omega_{k'} \text{Pers}(C_{\max}^{\omega_k}),$$

where $c_{i_s}^{\omega_k}$ is the community assignment of node i_s in $C_{\max}^{\omega_k}$ and $k, k' \in \{1, 2\}$; and by substituting ω_1 by $\omega_2 + \Delta$ for some $\Delta > 0$, one can show that the inequality (5.5) implies

$$Q(C_{\max}^{\omega_2} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega_2) < Q(C_{\max}^{\omega_1} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega_2),$$

which contradicts the optimality of $C_{\max}^{\omega_2}$. \square

One can similarly prove the following proposition.

PROPOSITION 5.8. *Let $\omega_1 > \omega_2 > 0$. For all $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_2)$, one of the following two conditions must hold:*

- (1) $C_{\max}^{\omega_2} \in \mathcal{C}_{\max}(\omega_1)$,
- or (2) $Q(C_{\max}^{\omega_2} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0) > Q(C_{\max}^{\omega_1} | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0)$ for all $C_{\max}^{\omega_1} \in \mathcal{C}_{\max}(\omega_1)$.

(To visualize Propositions 5.7 and 5.8 graphically, it is helpful to think of a multi-layer quality function $Q(C | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; \omega)$ for a given partition C as a linear function of ω with slope $\text{Pers}(C)$ that crosses the vertical axis at $Q(C | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0)$ [see, for e.g., the last panel of Fig. 5.1 and Fig. 5.2].) The next three corollaries follow straightforwardly from Propositions 5.7 and 5.8. The first states that the highest achievable value of persistence for an optimal partition obtained with a given value of inter-layer coupling is a non-decreasing function in ω . The second states that the highest achievable value of intra-layer modularity contributions for an optimal partition obtained with a given value of inter-layer coupling is a non-increasing function in ω . The third property states that if two distinct values of ω have the same set of optimal partitions, then this set is also optimal for all intermediate values.

COROLLARY 5.9. *Let $\omega_1 > \omega_2$. Then*

$$\text{Pers}(\mathcal{C}_{\max}(\omega_1)) \geq \text{Pers}(\mathcal{C}_{\max}(\omega_2)),$$

where $\text{Pers}(\mathcal{C}_{\max}(\omega)) := \max \{ \text{Pers}(C_{\max}^{\omega}), C_{\max}^{\omega} \in \mathcal{C}_{\max}(\omega) \}$.

COROLLARY 5.10. *Let $\omega_1 > \omega_2$. Then*

$$Q(\mathcal{C}_{\max}(\omega_1) | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0) \leq Q(\mathcal{C}_{\max}(\omega_2) | \mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0),$$




FIG. 5.3. Toy example illustrating the effect of post-processing on a multilayer partition by increasing multilayer modularity via community-assignment swaps that increase the value of persistence but do not change intra-layer partitions. The colors in panels (a)–(c) scale with the entries of the adjacency matrix. Panel (d) (respectively, panel (e)) represents the output multilayer partition obtained with Louvain before (respectively, after) post-processing. The horizontal axis represents the layers, and the vertical axis represents the nodes. The shading in panels (d,e) represents the community assignments of nodes in each layer.

where $Q(\mathcal{C}_{max}(\omega)|\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0) := \max \{Q(C_{max}^\omega|\mathbf{B}_1, \dots, \mathbf{B}_{|\mathcal{T}|}; 0), C_{max}^\omega \in \mathcal{C}_{max}(\omega)\}$.

COROLLARY 5.11. Assume that $\mathcal{C}_{max}(\omega_1) = \mathcal{C}_{max}(\omega_2)$ for $\omega_1 > \omega_2$. Then

$$\mathcal{C}_{max}(\omega_1) = \mathcal{C}_{max}(\omega) = \mathcal{C}_{max}(\omega_2) \text{ for all } \omega \in (\omega_2, \omega_1).$$

One can extend the proofs of Propositions 5.1–5.7 so that they apply for inter-layer coupling that is uniform between each pair of contiguous layers but may differ from pair to pair.

5.3. Implementation issues. We now examine issues that can arise when using the Louvain heuristic (see Section 2.2) to maximize multilayer modularity (3.3).

5.3.1. Underemphasis of persistence. Consider the example network in Fig. 5.3, which is a 3-layer network that has 5 nodes in each layer. Suppose that all nodes are strongly connected to each other in layers 1 and 3, and that the edge weight between node 1_2 and nodes $\{2_2, 3_2, 4_2, 5_2\}$ is smaller in layer 2 than the edge weight between node 1_s and nodes $\{2_s, 3_s, 4_s, 5_s\}$ when $s = 1, 3$. We use the uniform null network with $\gamma = 0.5$ and set $\omega = 0.1$. This produces a multilayer modularity matrix in which all the single-layer modularity entries B_{ijs} except those of node 1_2 are positive and exceed the value of inter-layer coupling. Suppose that one loops over the nodes ordered from 1 to $N|\mathcal{T}|$ in the first phase of the Louvain heuristic. The initial partition consists of $N|\mathcal{T}|$ singletons, and each node is then moved to the set that maximally increases modularity. The partition at the end of phase 1 is

$\{\{1_1, 2_1, 3_1, 4_1, 5_1, 1_2\}, \{2_2, 3_2, 4_2, 5_2\}, \{1_3, 2_3, 3_3, 4_3, 5_3\}\}$. In phase 2, the second and third sets merge to form a single set,¹⁵ and the Louvain heuristic gets trapped in a local optimum in which the smaller set of nodes (i.e., $\{1_1\}$) remains in the same community across layers 1 and 2 and the larger set of nodes (i.e., $\{2_1, 3_1, 4_1, 5_1\}$) changes community. We show this multilayer partition in Fig. 5.3(d). Repeating this experiment 1000 times using a randomized node order at the start of each iteration of phase 1 of the Louvain heuristic yields the same multilayer partition. One can modify this multilayer partition to obtain a new partition with a higher value of multilayer modularity by increasing the value of persistence across layers without changing intra-layer partitions (we use this idea in the proof of proposition 5.1). We show an example of this situation in Fig. 5.3(e).

In Fig. 5.3(d) we illustrate the above issue visually via abrupt changes in colors between layers (these are more noticeable in larger networks). Such changes are misleading because they imply a strong decrease in persistence that might not be accompanied by a strong change in intra-layer partitions. In Fig. 5.3(d), for example, the intra-layer partitions differ in the community assignment of only a single node. To mitigate this problem, we apply a post-processing function to all output partitions that increases persistence between layers without changing the partitions that are induced on each layer (thereby producing a partition with a higher value of multilayer modularity). We do this by relabeling the community assignments of nodes in each layer such that 1) the number of nodes that remain in the same community between consecutive layers is increased and 2) the partition induced on each layer by the original multilayer partition is unchanged.

5.3.2. Abrupt drop in the number of intra-layer merges. The Louvain heuristic faces a second problem in multilayer networks. When the value of inter-layer coupling satisfies

$$\omega > \text{Max}(\mathcal{B}_{\text{diag}}), \quad (5.6)$$

where $\mathcal{B}_{\text{diag}}$ is the set of values on the diagonal blocks of \mathcal{B} defined in equation (5.4), the inter-layer contributions to multilayer modularity are larger than the intra-layer contributions for all pairs of nodes. Consequently, only inter-layer merges occur during the convergence of phase 1 in the first iteration of the Louvain heuristic. In Fig. 5.4(a), we illustrate this phenomenon using data set 1. The mean number of intra-layer merges drops from roughly $N = 98$ (almost every node contains at least one other node from the same layer in its community) to 0. For ω values larger than $\text{Max}(\mathcal{B}_{\text{diag}})$, every set at the end of the first iteration of phase 1 only contains copies of each node in different layers and, in particular, does not contain nodes from the same layer. This can yield abrupt changes in the partitions induced on individual layers of the output multilayer partition

In Fig. 5.4(c), we show an example using data set 1 of how the above issue can lead to an abrupt change in a quantitative measure computed from a multilayer output partition obtained with the Louvain heuristic. Note that the mean size of the sets after convergence of phase 1 in Fig. 5.4(a) is relatively small for data set 1. (The mean is 3 nodes per set, and the maximum possible number of nodes per set is $|\mathcal{T}| = 238$.) Nevertheless, there is a sudden drop in the value of $(1 - \text{Pers}(C)|_s/N)$

¹⁵Note that combining the first and second set into a single set decreases modularity because the value of inter-layer coupling is too small to compensate for the decrease in intra-layer contributions to modularity.

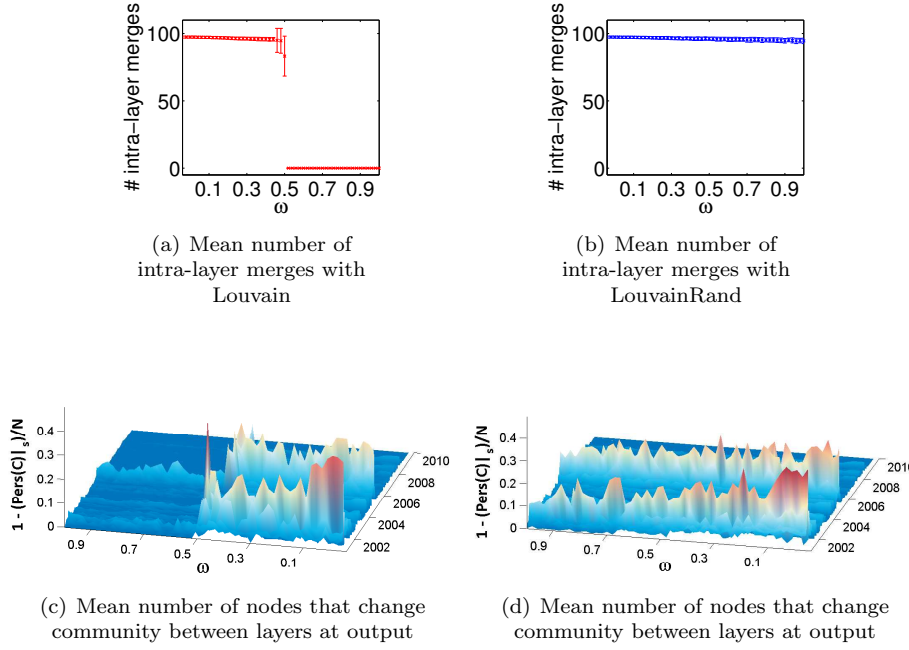


FIG. 5.4. Comparison between the Louvain and LouvainRand algorithms. The sample of inter-layer coupling values is the set $\{0, 0.02, \dots, 0.98, 1\}$ with a discretization step of 0.02 between each pair of consecutive values. (a,b) The number of nodes that have been merged with at least one node from the same layer after convergence of the first phase of (a) the Louvain heuristic and (b) the LouvainRand heuristic. For each heuristic, we average this value over $|\mathcal{T}| = 238$ layers and 100 iterations. The error bars in panels (a,b) indicate standard deviations. (c,d) The value of $1 - \text{Pers}(C)|_s/N$ averaged over 100 runs of (c) the Louvain heuristic and (d) the LouvainRand heuristic after convergence of the algorithms to a local optimum.

between consecutive layers at $\omega = \text{Max}(\mathcal{B}_{\text{diag}})$ [see Fig. 5.4(c)]. Nonzero values of $(1 - \text{Pers}(C)|_s)/N$ indicate that community assignments have changed between layers s and $s + 1$ (by proposition 5.3).

This problem manifests itself when the values of inter-layer coupling are large relative to the entries of $\mathcal{B}_{\text{diag}}$. In the correlation multilayer networks that we consider (or in unweighted multilayer networks), entries of the adjacency matrix satisfy $|A_{ijs}| \leq 1$. Assuming that one uses the modularity quality function on each layer and that $P_{ijs} \geq 0$ (e.g., $P_{ijs} = \langle \mathbf{A}_s \rangle$), this implies that

$$\text{Max}(\mathcal{B}_{\text{diag}}) \leq 1 \text{ for all } \gamma \in [\gamma^-, \gamma^+].$$

For networks in which the modularity cost of changing intra-layer partitions in favor of persistence is large in comparison to the values of $\text{Max}(\mathcal{B}_{\text{diag}})$, it might be desirable to use $\omega > 1$ to gain insight into a network's multilayer community structure (e.g., this occurs in both toy examples of Section 5.1).

To mitigate this problem, we change the condition for merging nodes in the Louvain heuristic. Instead of moving a node to a community that maximally increases modularity, we move a node to a community chosen uniformly at random from those that increase modularity. We call this heuristic *LouvainRand* [1], and we illustrate the results of using it in Figs. 5.4(b,d). Although LouvainRand increases the output variability (by increasing the search space of the optimization process), it seems

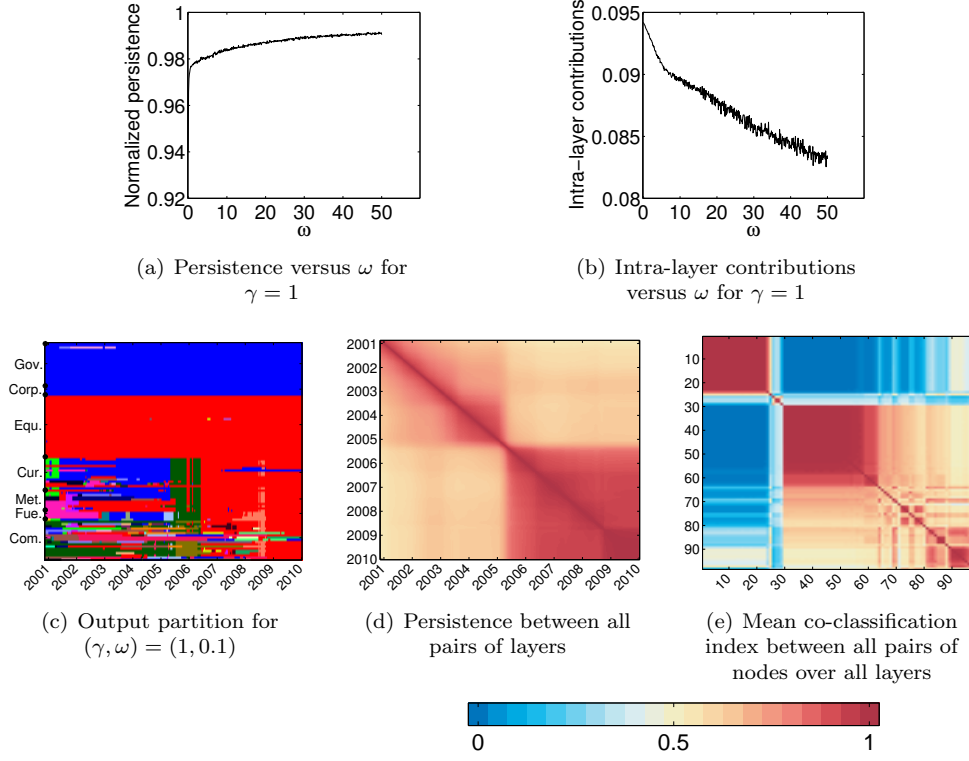


FIG. 5.5. Numerical experiments with data set 1. We sample the set of inter-layer edge weights uniformly from the interval $[0, 50]$ with a discretization step of 0.1 (so there are 501 values of ω in total), and we use the uniform null network (i.e., $P_{ijs} = \langle \mathbf{A}_s \rangle$) with $\gamma = 1$. (a) The persistence normalized by $N(|\mathcal{T}| - 1)$ for each value of ω averaged over 20 runs of LouvainRand. (b) The intra-layer modularity contributions $\sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{is}, c_{js})$ normalized by $\sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N A_{ijs}$ for each value of ω averaged over 20 runs of LouvainRand. (d) Sample output multilayer partition. Each point on the horizontal axis represents a single time window, and each position on the vertical axis is an asset. We order the assets by asset class, and the colors represent communities. (e) Association matrix of normalized persistence values between all pairs of layers averaged over all values of $\omega \in [0, 50]$ in our sample and 20 runs for each value. The normalized persistence between a pair of layers $\{s, r\}$ is $\sum_{i=1}^N \delta(c_{is}, c_{ir}) / N$. (f) Association matrix indicating the co-classification of nodes averaged over the set of partitions induced on each layer for each value of ω and 20 runs of LouvainRand.

to mitigate the problem for multilayer networks with ordinal diagonal and uniform coupling.

5.4. Multilayer community structure in asset correlation networks. In this section, we show the results of computational experiments in which we fix the value of the resolution parameter γ and vary the value of inter-layer coupling ω . We use the uniform null network (i.e., $P_{ijs} = \langle \mathbf{A} \rangle_s$) and set $\gamma = 1$. We use the LouvainRand heuristic to identify multilayer partitions and apply our post-processing procedure that increases persistence without changing partitions induced on individual layer to all output multilayer partitions. We showed in proposition 5.5 that for $2\omega > \omega_\infty = N^2[\text{Max}(\mathcal{B}_{\text{diag}}) - \text{Min}(\mathcal{B}_{\text{diag}})]$, the set $\mathcal{C}_{\text{max}}(\omega)$ of global optima no longer

changes and every optimal partition in this set has maximal persistence.¹⁶ In our example, $N^2[\text{Max}(\mathcal{B}_{\text{diag}}) - \text{Min}(\mathcal{B}_{\text{diag}})] \leq 2N^2$, with $N = 98$. However, for the purposes of the present paper, we take the set $\{0, 0.1, \dots, 0.49, 50\}$ with a discretization step of 0.1 between consecutive values (giving 501 values in total) as our sample of ω values.

In agreement with the properties derived in propositions 5.7 and 5.8, we observe in Fig. 5.5(a) that normalized persistence (given by $\text{Pers}(C)/[N(|\mathcal{T}| - 1)]$) tends to be larger for larger values of inter-layer coupling, and in Fig. 5.5(b) that intra-layer modularity contributions (which we normalize by $\sum_{s=1}^{|\mathcal{T}|} (\mathbf{1}^T \mathbf{A}_s \mathbf{1})$) tend to be smaller for larger values of inter-layer coupling. The increase of persistence and the decrease of intra-layer modularity contributions need not be monotonic, because we are finding a set of local optima for each value of ω rather than the set of global optima.

In Fig. 5.5(c), we show a sample output of the multilayer partition (which contains 35 communities). Some of the changes in community structure correspond to known events (e.g., Lehman Bankruptcy in September 2008 [marked by an increase in the size of the equity asset class]). Observe that the two largest communities are the ones that contain the government bond assets and the equity assets. In particular, the community that contains equities becomes noticeably larger in 2006 and in 2008. For larger values of the resolution parameter, this community instead becomes noticeably larger only in 2008. (By inspecting the correlation matrices, one can check that the increase in correlation between equities and other assets is greater in 2008 than in 2006.)

In Fig. 5.5(d), we show the matrix of mean values of persistence between all pairs of layers. The $(s, r)^{\text{th}}$ entry is the term $\sum_{i=1}^N \delta(c_{i_s}, c_{i_r})$, where $s, r \in \{1, \dots, |\mathcal{T}|\}$ need not be from consecutive layers, averaged over nodes, all values of $\omega \in [0, 50]$ in our sample, and multiple runs for each value of ω . Instead of only plotting $\text{Pers}(C)|_s$ for consecutive layers, Fig. 5.5 gives some indication as to whether nodes change communities between layers s and $s + 1$ to join a community that contains a copy of some of these nodes from another time layer (i.e., $\sum_{i=1}^N \delta(c_{i_{s+1}}, c_{i_r}) \neq 0$ for some r) or to join a community that does not contain a copy of these nodes in any other time layer (i.e., $\sum_{i=1}^N \delta(c_{i_{s+1}}, c_{i_r}) = 0$ for all r). Figure 5.5 also gives some insight into whether there are sets of consecutive layers across which persistence values remain relatively large. This may shed light on when connectivity patterns change in a multilayer network. As indicated by the values on the color scale, the values of persistence in Fig. 5.5(d) remain relatively high (which can partly be explained by the fact that equities and bonds remain in the same community across almost all layers, and these constitute roughly 50 % of the node sample). The most noticeable diagonal block separation in the middle of Fig. 5.5(d) corresponds to the change in Fig. 5.5(c) between 2005 and 2006, and the smaller diagonal block at the bottom right in Fig. 5.5(d) corresponds to the change in Fig. 5.5(c) after the Lehman Bankruptcy between 2008 and 2009.

In Fig. 5.5(e), we show the co-classification index of nodes in partitions induced on individual layers, which we average over layers, all values of $\omega \in [0, 50]$ in our sample, and multiple runs for each value of ω (we re-order the nodes to emphasize diagonal blocks in the association matrix). This figure yields insight into what sets of nodes belong to the same community across layers for increasing values of ω . This may shed light on connectivity patterns that are shared across layers. Unsurprisingly,

¹⁶Note that there can also be smaller values of ω_∞ for which this is true; in other words, we did not show that ω_∞ is the smallest lower bound of the set $\{\omega : \text{Pers}(C_{\text{max}}^\omega) = N(|\mathcal{T}| - 1) \text{ for all } C_{\text{max}}^\omega \in \mathcal{C}_{\text{max}}(\omega)\}$.

the first diagonal block mainly corresponds to bond assets and the second diagonal block mainly corresponds to equity assets. Figures 5.5(d,e) complement each other: at a given resolution, the latter gives an idea about *when* community structure has changed, and the former gives an idea about *how* it has changed.

6. Conclusions. Modularity maximization in temporal multilayer networks is a clustering technique that produces a time-evolving partition of nodes. We have investigated two questions that arise when using this method: (1) the role of null networks in modularity maximization, and (2) the effect of inter-layer edges on the multilayer modularity-maximization problem. We demonstrated that one must be cautious in interpreting communities obtained with a null network in which the distribution of expected edge weights is sample-dependent. Furthermore, we showed that an optimal partition in multilayer modularity maximization reflects a trade-off between static community structure within layers and persistence of community structure across layers. One can try to exploit this in practice to detect changes in connectivity patterns and shared connectivity in a time-dependent network.

At the heart of modularity maximization is a comparison between what one anticipates and what one observes. The ability to specify what is anticipated is a desirable (albeit under-exploited) feature of modularity maximization, because one can explicitly adapt it for different applications [5, 6, 19, 60]. By defining a null model as a probability distribution over the space of adjacency matrices and a null network as the expected adjacency matrix under the specified distribution, we highlight the important point that the *same* null network can correspond to *different* null models; this is not something that has been appreciated properly in the literature. Moreover, one needs to be very careful with one’s choice of null network because it determines what one regards as densely connected in a network: different choices in general yield different communities. As we illustrated in Section 4 for financial correlation networks, this choice can have a large impact on results, and can lead to misleading conclusions, so one should be cautious when interpreting communities that one obtains with a null network in which the distribution of expected edge weights is sample-dependent.

In Section 5, we proved several properties that describe the effect of ordinal diagonal and uniform inter-layer coupling on multilayer modularity maximization, or more generally, on any maximization problem that can be cast in the form (3.4). Although our theoretical results do not necessarily apply to the local optima that one attains in practice, they do provide useful guidelines for how to interpret the outcome of a computational heuristic for maximizing modularity: if a multilayer partition is inconsistent with one of the proven properties, then it must be an artifact of the heuristic and not a feature of the quality function.

To further examine multilayer modularity maximization, we defined a measure that we called *persistence* to quantify how much community assignments change in time in a multilayer partition. For zero inter-layer coupling, the value of persistence is 0, and it achieves a maximum finite value for sufficiently large inter-layer coupling. We showed that the highest achievable value of persistence for an optimal partition obtained with a given value of inter-layer coupling ω is a non-decreasing function in ω . Similarly, the highest achievable value of intra-layer contributions to the quality function for an optimal partition obtained with a given value of inter-layer coupling ω is a non-increasing function in ω . The notion of persistence makes it possible to measure this trade-off between static community structure within layers and temporal

persistence across layers:

$$\max_{C \in \mathcal{C}} \sum_{s=1}^{|\mathcal{T}|} \sum_{i,j=1}^N B_{ijs} \delta(c_{i_s}, c_{j_s}) + 2\omega \text{Pers}(C).$$

We illustrated this trade-off in our numerical experiments.

Finally, we showed that the Louvain heuristic can pose two issues when applied to multilayer networks with ordinal diagonal and uniform coupling. These can produce misleading values of persistence (or other quantitative measures of a multilayer partition) and can cause one to draw false conclusions about temporal changes in community structure in a network. We proposed ways to mitigate these problems and showed several numerical experiments on real data as illustrations. To further interpret these results, one needs to investigate more closely how the increase in persistence and the decrease in intra-layer contributions to the quality function actually manifest in a multilayer partition between the “boundary cases”. This may help identify an interval of ω values in which the trade-off between these two quantities yields the most insights.

Acknowledgements. MB acknowledges a CASE studentship award from the EPSRC (BK/10/41). MB and MAP were supported by FET-Proactive project PLEX-MATH (FP7-ICT-2011-8; grant #317614) funded by the European Commission. MAP was also supported by the EPSRC (EP/J001759/1). We thank HSBC Bank for providing the time series data of financial assets. We thank Lucas Jeub, Alex Arenas, and Sergio Gómez for useful discussions.

REFERENCES

- [1] Code available at <http://netwiki.amath.unc.edu/GenLouvain/GenLouvain>.
- [2] A. Arenas, A. Fernández, and S. Gómez. Analysis of the structure of complex networks at different resolution levels. *New J. Phys.*, 10:5:053039, 2008.
- [3] T. Aynaud and J.-L. Guillaume. Static community detection algorithms for evolving networks. *WiOpt’10: Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, 84:508–514, 2010.
- [4] I. Barnett and J.-P. Onnela. Change point detection in correlation networks. *arXiv:1410.0761*.
- [5] D. S. Bassett, E. T. Owens, M. A. Porter, M. L. Manning, and K. E. Daniels. Extraction of force-chain network architecture in granular materials using community detection. *arXiv:1408.3841*, 2014.
- [6] D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson, and P. J. Mucha. Robust detection of dynamic community structure in networks. *Chaos*, 23:1:013142, 2013.
- [7] D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton. Dynamic reconfiguration of human brain networks during learning. *Proc. Natl. Acad. Sci. USA*, 108:7641–7646, 2011.
- [8] S. Battiston, J. B. Glattfelder, D. Garlaschelli, F. Lillo, and G. Caldarelli. *The structure of financial networks*. Springer, London, UK, 2010, Chapter 7.
- [9] E. T. Bell. Exponential numbers. *Amer. Math. Monthly*, 41:411–419, 1934.
- [10] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *J. Statist. Mech.*, 10:P10008, 2008.
- [11] S. Boccaletti, G. Bianconi, R. Criado, C.I. Del Genio, J. Gómez-Gardenes, M. Romance, I. Sendina-Nadal, Z. Wang, and M. Zanin. The structure and dynamics of multilayer networks. *Phys. Rep.*, 544:1–122, 2014.
- [12] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. *Random Struc. Alg.*, 31:3–122, 2007.
- [13] B. Bollobás and O. Riordan. Random graphs and branching processes, in: *Handbook of large-scale random networks*. *Bolyai Soc. Math. Stud.*, 18:15–115, 2009.
- [14] U. Brandes, D. Delling, M. Gaertler, R. Görke, M. Hoefer, Z. Nikoloski, and D. Wagner. On modularity clustering. *IEEE Trans. Knowledge Data Engrg.*, 20:172–188, 2008.

- [15] M. Catanzaro and M. Buchanan. Network opportunity. *Nature Physics*, 9:121–123, 2013.
- [16] F. Chung and L. Lu. The average distance in a random graph with given expected degrees. *Proc. Natl. Acad. Sci. USA*, 99:15879–15882, 2002.
- [17] J.-C. Delvenne, S. N. Yaliraki, and M. Barahona. Stability of graph communities across time scales. *PNAS*, 107(29):12755–12760, 2010.
- [18] M. De Domenico, A. Lancichinetti, A. Arenas, and M. Rosvall. Identifying modular flows on multilayer networks reveals highly overlapping organization in interconnected systems. *arXiv:1408.2925*, 2014.
- [19] P. Expert, T. S. Evans, V. D. Blondel, and R. Lambiotte. Uncovering space-independent communities in spatial networks. *Proc. Natl. Acad. Sci. USA*, 108:7663–7668, 2011.
- [20] D. J. Fenn, M. A. Porter, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones. Dynamic communities in multichannel data: An application to the foreign exchange market during the 2007–2008 credit crisis. *Chaos*, 19(3):033119, 2009.
- [21] D. J. Fenn, M. A. Porter, P. J. Mucha, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones. Dynamical clustering of exchange rates. *Quantitative Finance*, 12(10):1493–1520, 2012.
- [22] D. J. Fenn, M. A. Porter, S. Williams, M. McDonald, N. F. Johnson, and N. S. Jones. Temporal evolution of financial market correlations. *Phys. Rev. E*, 84:2:026109, 2011.
- [23] S. Fortunato. Community detection in graphs. *Physics Reports*, 486:75–174, 2010.
- [24] S. Fortunato and M. Barthélemy. Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA*, 104:36–41, 2006.
- [25] N. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proc. Natl. Acad. Sci. USA*, 99:7821–7826, 2002.
- [26] S. Gómez, P. Jensen, and A. Arenas. Analysis of community structure in networks of correlated data. *Phys. Rev. E*, 80:1:016114, 2009.
- [27] S. González-Bailón, J. Borge-Holthoefer, A. Rivero, and Y. Moreno. The dynamics of protest recruitment through an online network. *Sci. Rep.*, 1(197), 2011.
- [28] B. H. Good, Y.-A. de Montjoye, and A. Clauset. Performance of modularity maximization in practical contexts. *Phys. Rev. E*, 81:4:046106, 2010.
- [29] R. Guimerà, M. Sales-Pardo, and L. A. N. Amaral. Functional cartography of complex metabolic networks. *Nature*, 443:895–900, 2005.
- [30] T. Hoffmann, M. A. Porter, and R. Lambiotte. Social structure of Facebook networks. *Phys. Rev. E*, 86(4):046102, 2012.
- [31] P. Holme and J. Saramäki. Temporal networks. *Phys. Rep.*, 519:97–125, 2012.
- [32] J. Hopcroft, O. Khan, B. Kulli, and B. Selman. Tracking evolving communities in large linked networks. *Proc. Natl. Acad. Sci. U.S.A.*, 101:5250–5253, 2004.
- [33] L. G. S. Jeub, P. Balachandran, M. A. Porter, P. J. Mucha, and M. W. Mahoney. Think locally, act globally: The detection of small, medium-sized, and large communities in large networks. *Phys. Rev. E*, in press (*arXiv:1403.3795*), 2014.
- [34] I. S. Jutla, L. G. S. Jeub, and P. J. Mucha. A generalized louvain method for community detection implemented in matlab, version 2.0. *Code available at <http://netwiki.amath.unc.edu/GenLouvain/GenLouvain>*, 2012.
- [35] M. Kivelä, A. Arenas, M. Barthélemy, J. P. Gleeson, Y. Moreno, and M. A. Porter. Multilayer networks. *Journal of Complex Networks*, 2(3):203–271, 2014.
- [36] R. Lambiotte. Multi-scale modularity in complex networks. *Proceedings of the 8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, pages 546–553, 2010.
- [37] R. Lambiotte, J.-C. Delvenne, and M. Barahona. Random walks, markov processes and the multiscale modular organization of complex networks. *TNSE*, 1(2):76–90, 2014 (*arXiv:0812.1770*, 2009).
- [38] A. Lancichinetti and S. Fortunato. Community detection algorithms: A comparative analysis. *Phys. Rev. E*, 80(5):056117, 2009.
- [39] A. Lancichinetti and S. Fortunato. Consensus clustering in complex networks. *Scientific Reports*, 2(336), 2012.
- [40] A. C. F. Lewis, N. S. Jones, M. A. Porter, and C. M. Deane. The function of communities in protein interaction networks at multiple scales. *BMC Syst. Biol.*, 4:100, 2010.
- [41] M. MacMahon and D. Garlaschelli. Unbiased community detection for correlation matrices. *Physical Review X*, in press (*arXiv:1311.1924*), 2014.
- [42] K. T. Macon, P. J. Mucha, and M. A. Porter. Community structure in the United Nations General Assembly. *Physica A*, 391:343–361, 2012.
- [43] R. N. Mantegna. Hierarchical structure in financial markets. *Eur. Phys. J. B*, 11:193–197, 1999.

- [44] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela. Community structure in time-dependent, multiscale and multiplex networks. *Science*, 328:876–878, 2010.
- [45] P.J. Mucha and M. A. Porter. Communities in multislice voting networks. *Chaos*, 20:4:041108, 2010.
- [46] M. E. J. Newman. Analysis of weighted network. *Physical Review E*, 70:6:066146, 2004.
- [47] M. E. J. Newman. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E*, 74(3:036104), 2006.
- [48] M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, Oxford, UK, 2010.
- [49] M. E. J. Newman. Communities, modules and large-scale structure in networks. *Nature Physics*, 8:25–31, 2012.
- [50] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69(2:026113), 2004.
- [51] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertész, and A. Kanto. Dynamics of market correlations: taxonomy and portfolio analysis. *Phys. Rev. E*, 68:5:056110, 2003.
- [52] G. Palla, A.-L. Barabási, , and T. Vicsek. Quantifying social group evolution. *Nature*, 446:664–667, 2007.
- [53] L. Peel and A. Clauset. Detecting change points in the large-scale structure of evolving networks. *To appear in AAAI (arXiv:1403.0989)*, 2015.
- [54] G. Pietri and P. Expert. Temporal stability of network partitions. *arXiv:1404.7170*, 2014.
- [55] M. A. Porter, J.-P. Onnela, and P. J. Mucha. Communities in networks. *Notices of the Amer. Math. Soc.*, 56:1082–1097, 1164–1166, 2009.
- [56] M. Potters, J.-P. Bouchaud, and L. Laloux. Financial applications of random matrix theory: old laces and new pieces. *Acta Physica Polonica B*, 36:2767–2784, 2005.
- [57] F. Radicchi and A. Arenas. Abrupt transition in the structural formation of interconnected networks. *Nature Physics*, 9:717–720, 2013.
- [58] J. Reichardt and S. Bornholdt. Statistical mechanics of community detection. *Phys. Rev. E*, 74:1:016110, 2006.
- [59] M. Sales-Pardo, R. Guimerà, A. A. Moreira, and L. A. N. Amaral. Extracting the hierarchical organization of complex systems. *Proc. Natl. Acad. Sci. USA*, 39:15224–15229, 2007.
- [60] M. Sarzynska, E. L. Leicht, G. Chowell, , and M. A. Porter. Null models for community detection in spatially-embedded, temporal networks. *arXiv:1407.6297*, 2014.
- [61] R. Schafer and T. Guhr. Local normalization: Uncovering correlations in non-stationary financial time series. *Physica A*, 389:3856–3865, 2010.
- [62] H. A. Simon. The architecture of complexity. *Proc. Am. Phil. Soc*, 10:467–482, 2008.
- [63] S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsey, and M. W. Woolrich. Network modelling methods for FMRI. *NeuroImage*, 54(2):875–891, 2011.
- [64] V. A. Traag and J. Bruggeman. Community detection in networks with positive and negative links. *Phys. Rev. E*, 80:3:036115, 2009.
- [65] V. A. Traag and P. Van Dooren. Narrow scope for resolution-limit-free community detection. *Phys. Rev. E*, 84(1:016114), 2011.
- [66] A. L. Traud, E. D. Kelsic, P. J. Mucha, and M. A. Porter. Comparing community structure to characteristics in online collegiate social networks. *SIAM Rev.*, 53:526–543, 2011.
- [67] S. Williams. Risk on risk off will test investors through 2011. *Financial Times*, 2011.
- [68] Y. Wu, C. Zhou, J. Xiao, J. Kurths, and H.J. Schellnhuber. Evidence for a bimodal distribution in human communication. *Proc. Natl. Acad. Sci. USA*, 107:18803–18808, 2010.
- [69] A. Zalesky, A. Fornito, and E. Bullmore. On the use of correlation as a measure of network connectivity. *NeuroImage*, 60:2096–2106, 2012.
- [70] Q. Zhao, Y. Tian, Q. He, N. Oliver, R. Jin, and W.-C. Lee. Communication motifs: A tool to characterize social communications. *In Proceedings of the 19th ACM international conference on Information and knowledge management*, 84:1645–1648, 2010.