
Efficient Near-Optimal Testing of Community Changes in Balanced Stochastic Block Models

Aditya Gangrade
Boston University
gangrade@bu.edu

Praveen Venkatesh
Carnegie Mellon University
vpraveen@cmu.edu

Bobak Nazer
Boston University
bobak@bu.edu

Venkatesh Saligrama
Boston University
srv@bu.edu

Abstract

We propose and analyze the problems of *community goodness-of-fit* and *two-sample testing* for stochastic block models (SBM), where changes arise due to modification in community memberships of nodes. Motivated by practical applications, we consider the challenging sparse regime, where expected node degrees are constant, and the inter-community mean degree (b) scales proportionally to intra-community mean degree (a). Prior work has sharply characterized partial or full community recovery in terms of a “signal-to-noise ratio” (SNR) based on a and b . For both problems, we propose computationally-efficient tests that can succeed far beyond the regime where recovery of community membership is even possible. Overall, for large changes, $s \gg \sqrt{n}$, we need only $\text{SNR} = O(1)$ whereas a naïve test based on community recovery with $O(s)$ errors requires $\text{SNR} = \Theta(\log n)$. Conversely, in the small change regime, $s \ll \sqrt{n}$, via an information-theoretic lower bound, we show that, surprisingly, no algorithm can do better than the naïve algorithm that first estimates the community up to $O(s)$ errors and then detects changes. We validate these phenomena numerically on SBMs and on real-world datasets as well as Markov Random Fields where we only observe node data rather than the existence of links.

While community detection and recovery for the stochastic block model (SBM) [Abb18] and, more generally, inference of community structures underlying large-scale network data [GN02; New06; For10] has received significant interest across the machine learning, statistics and information theory literatures, there has been limited work on the important problem of testing changes in community structures. The general problem of testing changes in networks naturally arises in a number of applications such as discovering statistically significant topological changes in gene regulatory networks [Zha+08] or differences in brain networks between healthy and diseased individuals [Bas+08]. Building upon this perspective, we propose testing of differences in the underlying community structure of a network, which can encompass scenarios such as detecting structural changes over time in social networks [AG05; For10], determining whether a set of genes belong to different communities in disease and normal states [JTZ04], and deciding whether there are changes in functional modules, which represent communities, in protein-protein networks [CY06].

Testing structural changes in networks is statistically challenging due to the fact that we may have relatively few independent samples to evaluate combinatorially-many potential changes. In this paper, we propose methods for goodness-of-fit (GoF) testing and two-sample testing (TST) for detecting changes in community memberships under the SBM. The SBM naturally captures the community structures commonly observed in large-scale networks, and serves as a baseline model for more complex networks. Specifically, there are n nodes partitioned into two equal-sized communities, and the

network is observed as a random $n \times n$ adjacency matrix, representing the instantaneous pairwise interactions among individuals in the population. Both intra- and inter-community interactions are allowed. Members within the same community interact with uniform probability a/n , while members belonging to different communities with a smaller probability b/n . We restrict attention to the commonly-considered and practically-relevant setting of $a/b = \Theta(1)$.

For our testing problems, we assume that the network samples are aligned on $n \gg 1$ vertices, and that the latent communities are either the same, or they differ in at least some $s \ll n$ nodes. We pose the GoF problem as: *Decide whether or not the observed random incidence matrix is an instantiation of a given community structure.* For the TST problem, we ask: *Given two random incidence matrices, decide whether or not their latent community structure is identical.*

Sparse vs. Dense Graphs. We focus on scenarios where the observed random incidence matrices are sparse with average node degree bounded by a constant independent of the network size. Within this context we develop minimax optimal methods for GoF and TST in this context. We are motivated by both practical and theoretical concerns. Practically, as observed in [Chu10], realistic graphs such as social networks are sparse (friendships do not grow with network size); in temporal settings, at any given time, only a small subset of interactions are observed; and in other cases ascertaining the presence or absence of each edge in the network being observed is an expensive process, and it makes sense to understand the fundamental limits for when testing is even possible.

From a theoretical standpoint, the sparse setting is challenging due to signal-to-noise ratio (SNR) constraints that do not arise in the dense case. Recovery of the latent community with up to s errors is possible iff $\Lambda \gtrsim \log(n/s)$ [CRV15; ZZ16; FC19], where Λ is a SNR parameter that, in the setting $a/b = \Theta(1)$, scales linearly with the mean degree. In particular, for Λ of constant order, recovery with sublinear distortion fails. The question of *whether testing is possible when recovery fails* is mathematically intriguing. Further, this is the *only* theoretically interesting setting. Indeed, if testing for s changes requires a graph dense enough to allow recovery with $\sim s$ errors, then one might as well recover these communities and compare them.

Our Contributions. We show that optimal tests exhibit a surprising two-phase behavior:

1. For $s \gg \sqrt{n}$, or ‘large changes,’ we propose computationally-efficient schema for GoF and TST that succeed with $\Lambda = O(1)$ - far below the SNR threshold for recovery. For GoF, this requirement is even weaker - we only need $\Lambda \gtrsim n/s^2$, which vanishes with n since $s \gg \sqrt{n}$. Further, we match these bounds up to constants with information-theoretic lower bounds.
2. In contrast, we show via an information-theoretic lower bound that for $s \ll \sqrt{n}$, or ‘small changes,’ both testing problems require $\Lambda = \Omega(\log(n))$ for reliable testing. This means that the naïve strategy of recovering communities and comparing them is tight up to constants in this regime.

We complement the above theoretical study by three experiments: the first implements the above tests on synthetic SBMs, and the second on the political blogs dataset - a popular real world dataset for community detection [AG05]. Both of these experiments show excellent agreement with the theoretical predictions. The third experiment casts a wider net, and instead studies the related problem of testing the underlying community structure of a Gaussian Markov Random Field that has precision matrix $I + \gamma G$ for G drawn from an SBM. This experiment explores the more realistic setting where instead of receiving a graph, we obtain observations at each node of a hidden graph, and wish to reason about the underlying structure. Remarkably, a simple adaptation of our procedure for SBMs shows excellent performance for this problem. This indicates that our observations are not restricted to raw SBMs, but may signal a more general phenomenon that merits exploration.

Related Work. For work on recovery communities we refer to the survey [Abb18]. However, we explicitly point out the papers [CRV15; ZZ16; FC19], which provide various schemes and necessary conditions that show that the partial recovery problem with distortion s can be solved with vanishing error probability if and only if $\Lambda \gtrsim \log(n/s)$. We further point out the lower bounds of [MNS15; DAM17], which assert that if $\Lambda < 2$, then asymptotically, the best possible distortion for partial recovery (or weak recovery, as it is referred to in this constant SNR regime) is $n/2 - o(n)$. Note that reporting a uniformly random community achieves distortion of $s = n/2 - O(\sqrt{n})$.

Ours is the first work to study GoF and TST where both hypothesized models are SBMs. Nevertheless, both GoF and TST in the context of network data as well as SBMs have been studied. Below we highlight the key differences in modeling assumptions and the ensuing technical implications, which renders much of the prior work inapplicable to our setting.

With regards to GoF, [AV14; VA15] study the problem of detecting if a graph is an unstructured

Erdős-Rényi (ER) graph, or if it has a planted dense subgraph, providing detailed characterizations of the feasibility regions and statistical phase transitions in this setting. While this work is aligned with ours in the techniques used, the modeled setting and problem there are different (ER vs. planted dense subgraph), and TST is not explored. Particularly, the dense subgraph model and the SBM are qualitatively different, and conclusions from one cannot be transferred to the other directly.

A number of papers, including [Lei16; BS16; Ban+16; GL17] study various techniques and regimes of determining if a graph is a SBM or an unstructured ER graph, and if the former, the number of communities in the model. Of these, [GL17] approach the problem by counting small motifs in the graphs, [Ban+16] propose a simple scan and [Lei16; BS16] propose testing of the number of communities on the basis of the top singular values of the graph.

[Tan+17] study TST of the model parameters in random dot product graphs, and propose the distance between aligned spectral embeddings of the two graphs as a statistic to do so. They use this to test equality against various transformations of the underlying models, and in particular for SBMs, test if the connectivity probabilities $(a/n, b/n)$ are identical or not for two graphs with latent communities that are randomly drawn. [LL18] adapt these tests by considering the same distance, but weighted by the corresponding singular values of one of the graphs, and use this to study two-sample testing of equality of the latent communities in the graphs - as in this paper.

In contrast to the low-rank structure assumptions in the above work, [Gho+17a; Gho+17b; GL18] study two-sample testing of inhomogeneous ER graphs (i.e., ER graphs where each edge may have a distinct probability of existing). Within this setting, they provide a number of statistics based both on estimates of the Frobenius and operator norms of the differences of the expected graph adjacency matrices, as well as those based on motifs such as triangles, and explore the limits of these tests.

A fundamental drawback of these approaches, in our context, is their reliance on singular values, spectral norms and Frobenius norms. Singular embeddings are particularly sensitive to noise, and stable embeddings require significant edge density (particularly when a sublinear number of alterations to the communities are to be tested). Indeed, in this context, we note that, in contrast to our low SNR, sparse setting, [LL18] require both a degree of $n^{1/2-\epsilon}$ and an SNR of $\log(n)$ corresponding to a high SNR, high edge-density regime, where full community recovery is possible.

Similarly, Frobenius and Spectral norms based tests of [GL18; Gho+17a] are not stable enough to test a sublinear number of changes in a low SNR regime. Functionally, this can be seen by the fact that the square-Frobenius norm of the difference of two graphs is equal to the number of edges that appear in one graph but not the other, and for sparse graphs, *most* edges appear in only one of the two graphs. Similarly, arguments about spectral norms rely on concentration of the same for ER graphs, but the best known concentration radius [LLV17] is far too large to allow testing of small differences in sparse graphs. Indeed, for any of the statistics of [GL18] to have power in our setting, the results of the paper require that the expected degree diverges with n , and that $\Lambda \gtrsim n/s$, which is exponentially above the SNR required to recover communities up to distortion $s/2$.

1 Definitions

The Stochastic Block Model. A vector $x \in \{\pm 1\}^n$ is said to be a *balanced community vector* (or partition) if $\sum x_i = 0$. The *stochastic block model* is defined as a random, simple, undirected graph G on n nodes such that all edges are drawn mutually independently given x , and

$$P(\{i, j\} \in G | x) = \frac{a+b}{2n} + \frac{a-b}{2n} x_i x_j.$$

Note that we treat x as a deterministic but unknown quantity, and thus, $P(\cdot | x)$ is a slight abuse of notation. The parameters (a, b) may vary with n , and we focus on the setting $a, b = O(\log n)$, with emphasis on $O(1)$ ¹, and $a/b = \Theta(1)$. For technical convenience, we require that $a + b < n/4$.

The *signal-to-noise ratio* (SNR) of an SBM is the quantity $\Lambda := \frac{(a-b)^2}{a+b}$, which characterises the recovery problem, as described in earlier discussions.

¹While our main interest is in the constant degree regime, we also show that testing for small changes is impossible in this setting (e.g Thm 1), and instead logarithmic degrees are needed. Thus, to present our results completely, we must allow a, b to vary at least in the range $[\Omega(1), O(\log(n))]$, or, more succinctly $O(\log n)$. Large scales are not of interest since exact recovery is possible at the logarithmic scale.

Note that the partitions x and $-x$ induce the same distribution. Accordingly, the *distortion* between partitions x and y is $d(x, y) := \min(d_H(x, y), d_H(x, -y))$, where d_H is the Hamming distance.

Minimax Testing Problems. We formally define two minimax hypothesis testing problems.

Goodness-of-Fit. We are given a balanced partition x_0 and a parameter s . We receive a graph $G \sim P(G|x)$, where x is an unknown balanced partition that is either exactly equal to x_0 or differs in at least s places. Our goal is to solve the hypothesis test:

$$H_0 : d(x, x_0) = 0 \quad \text{vs.} \quad H_1 : d(x, x_0) \geq s$$

We measure the minimax risk of this problem by

$$R_{\text{GoF}}(n, s, a, b) := \inf_{\phi} \sup_{x_0} \left\{ P(\text{FA}) + \sup_x P(\text{MD}(x)) \right\} \quad (1)$$

where $\phi(G)$ outputs either 0 or 1, $P(\text{FA}) := P(\phi(G) = 1 \mid x_0)$, $P(\text{MD}(x)) := P(\phi(G) = 0 \mid x)$, are respectively the false alarm and missed detection probabilities, and the second supremum is over all x such that $d(x, x_0) \geq s$.

Two-Sample Testing. We are given a parameter s and two independent graphs $G \sim P(G|x)$, $H \sim P(H|y)$, where x and y are unknown balanced communities satisfying $d(x, y) \in \{0\} \cup [s : n/2]$. The goal is to solve the following (composite null) testing problem:

$$H_0 : d(x, y) = 0 \quad \text{vs.} \quad H_1 : d(x, y) \geq s,$$

with the measure of risk

$$R_{\text{TST}}(n, s, a, b) := \inf_{\phi} \sup_{x, y} P(\phi(G, H) \neq \mathbf{1}\{x = y\} \mid x, y), \quad (2)$$

where $\phi(G, H)$ outputs either 0 or 1 and the supremum is over balanced x, y such that $d(x, y) \in \{0\} \cup [s : n/2]$.

As we vary n and (s, a, b) with n as some functions (s_n, a_n, b_n) , the above define a sequence of hypothesis tests. We say that the GoF problem can be solved *reliably* for such a sequence if $R_{\text{GoF}}(n, s_n, a_n, b_n) \rightarrow 0$ as $n \nearrow \infty$, and similarly for TST. Below, we will target $O(1/n)$ bounds. For conciseness, we will suppress the dependence of risks on (n, s, a, b) , writing just $R_{\text{GoF}}/R_{\text{TST}}$.

On balance: The strict balance assumption above can be relaxed to only requiring that both communities are of size linear in n , at the cost of weakening some of the constants left implicit in the theorem statements. While the majority of the analysis in the paper will assume exact balance, we briefly discuss unbalanced but linearly sized communities whilst detailing the proofs. Note that since the communities are no longer balanced, the differences between x and y can be ‘one-sided’ i.e., more nodes can move from, say, $+$ to $-$, than in the other direction. We do not require any control on these other than the total number of changes is at least s .

On constants: We use C and c , and their modifications, as unspecified constants that may change from line to line. While these can be explicitly bounded, we do not expect them to be tight.

2 Community Goodness-of-Fit

We begin by stating our main results regarding the *community goodness-of-fit problem*.

Theorem 1. *Community goodness-of-fit testing is possible with risk $R_{\text{GoF}} \leq \delta$ if $s\Lambda \geq C \log(2/\delta)$ and $\Lambda \geq C \frac{n}{s^2} \log(2/\delta)$ for some constant $C > 0$.*

Conversely, in order to attain $R_{\text{GoF}} \leq \delta \leq 0.25$, we must have that $s\Lambda \geq C' \log(1/\delta)$ and $\Lambda \geq C' \log\left(1 + \frac{n}{s^2}\right)$ for some constant $C' > 0$.

These bounds reveal the following behavior in terms of large and small changes:

- For large changes ($s \geq n^{1/2+c}$ for some $c > 0$), since $n/s^2 \leq 1$ and $\log(1+x) \geq x/2$ for $x \leq 1$, the second converse bound behaves as $\Lambda \geq Cn/s^2$, matching the sufficient condition up to a constant.
- For small changes ($s \leq n^{1/2-c}$ for some $c > 0$), since $n/s^2 \sim n^{2c}$, the second converse bound instead behaves as $\Lambda \gtrsim \log n$. In this regime, community recovery up to $s/2$ errors requires

$\Lambda \geq C \log 2n/s = \tilde{C} \log n$. Thus, estimating x from G and comparing it to x_0 is optimal up to constants.

- The above indicate a phase transition in the GoF testing problem at $\sigma := \log_n(s) = 1/2$. Consider the thermodynamic limit of $n \nearrow \infty$. For $\sigma < 1/2$, the problem is ‘hard’ in that the SNR Λ is required to diverge to ∞ , while for $\sigma > 1/2$, the SNR can tend to zero.

Proof Sketch for the Achievability. Let us begin with an intuitive development of the test. Since we start with a partition x_0 in hand to test, it is natural to look at the edges across and within the cut defined by x_0 . We thus define the number of edges *across* and *within* this cut:

$$\begin{aligned} N_a^{x_0}(G) &:= |\{(i, j) \in G : x_{0,i} \neq x_{0,j}\}| = \frac{1}{4} x_0^\top (D(G) - G) x_0 \\ N_w^{x_0}(G) &:= |\{(i, j) \in G : x_{0,i} = x_{0,j}\}| = \frac{1}{4} x_0^\top (D(G) + G) x_0 \end{aligned} \quad (3)$$

where the latter expressions treat G as an adjacency matrix and $D(G) = \text{diag}(\text{degree}(i))_{i \in [1:n]}$.² In the null case, these are respectively $\text{Bin}(n^2/4, b/n)$ and $\text{Bin}(2\binom{n}{2}, a/n)$ random variables, while in the alternate case some $s/2 \cdot (n-s)/2$ of each behave like edges of the opposite polarity (i.e. as b/n instead of a/n and vice versa), leading to a excess/deficit of edges of this type. Note that while the ‘average signal strength’, i.e., the amount by which edges are over- or underrepresented is the same in both cases ($\sim s|a-b|$), the group with the larger null parameter suffers greater fluctuations. Thus, we base our test only on edges of smaller bias. This reduces the SNR by at most a factor of 4.

We now define the test. C_1 below is the constant implicit in Lemma 3 in Appendix A.1.

- If $a > b$, we use the test $N_a^{x_0}(G) \underset{H_0}{\overset{H_1}{\gtrless}} \frac{bn}{4} + C_1 \max\left(\sqrt{nb \log(1/\delta)}, \log(1/\delta)\right)$.
- If $b > a$, we use the test $N_w^{x_0}(G) \underset{H_0}{\overset{H_1}{\gtrless}} \frac{an}{4} - \frac{a}{2} + C_1 \max\left(\sqrt{na \log(1/\delta)}, \log(1/\delta)\right)$.

The risks of these tests can be controlled by separating the null and alternate ranges using Bernstein’s inequality. Indeed, the threshold above is just the the expectation plus the concentration radius of the statistic under the null distribution. Let us briefly develop the statistic’s behaviour in the alternate - considering only the case $a > b$, we find that under the alternate, $\binom{n-s}{2} + \binom{s}{2}$ of the edges in $N_a^{x_0}$ continue to behave like $\text{Bern}(b/n)$ bits, while the remaining $s(n-s)/2$ edges behave as $\text{Bern}(a/n)$ bits. Thus, the expectation of $N_a^{x_0}$ is increased by an amount greater than $s(n-s)\frac{a-b}{2n} \geq s(a-b)/4$. Next, Bernstein’s inequality controls the fluctuations at scale $\sqrt{\max(nb, s(a-b)) \log(2/\delta)}$. The conclusion is straightforward to draw from here, and the proof is carried out in Appendix A.1³.

Proof Sketch for the Converse. The proof is relegated to Appendix A.2, and we discuss the strategy here. The converse proof follows Le Cam’s method, which lower bounds the minimax risk by the Bayes risk for conveniently chosen priors - which can be expressed using the TV distance.

To show $\Lambda \gtrsim \log(1 + n/s^2)$, we pick the null x_0 to be any balanced community, and choose the uniform prior on communities that are exactly s -far from x_0 (in fact, we only use a subset of these in order to facilitate easier computations). This is an obvious choice for this setting - we are interested in balanced communities that are at least s far, and choosing a large number of them allows for a greater ‘confusion’ in the testing problem due to a richer alternate hypothesis. The bound follows by invoking inequalities between TV and χ^2 divergences and a lengthy calculation due to the combinatorial objects involved.

To show $s\Lambda \gtrsim -\log(\delta)$, we again pick the null to be any balanced community, and pick the alternate to be an s -far singleton. We then proceed to control d_{TV} by the Hellinger divergence.

3 Two-Sample Testing

We again begin with the main results on *community two-sample testing problem*.

²Note that $D(G) - G$ is the Laplacian of the graph.

³The same also describes the extension of the claims to linearly sized communities

Theorem 2. Assume, for some $\gamma > 0$, $s \geq n^{\frac{1}{2}+\gamma}$. There exist constants C, C' such that if $C' \leq a, b \leq (n/2)^{1/3}$, then two-sample testing of s changes with $R_{\text{TST}} \leq 4/n$ is possible if the SNR satisfies $\Lambda \geq C$.

Conversely, for $n \geq 200$, there exist constants c, c' such that if $s < (\frac{1}{2} - c')n$, then two-sample testing of s changes cannot be carried out with $R_{\text{TST}} \leq 1/4$ unless $\Lambda \geq c$.

Large Changes. The above theorem makes an achievability claim for the setting of large changes. Notice that in this regime the stated upper and lower bounds match up to constants. Specifically, if $n^{\frac{1}{2}+\gamma} < s < (\frac{1}{2} - c')n$, two-sample testing can be solved iff $\Lambda \gtrsim 1$. Further, the condition $a, b \gtrsim 1$ is also tight, as it follows from $a/b = \Theta(1)$, and the necessary condition $\Lambda \gtrsim 1$, since $\Lambda \leq a + b$.

This leaves the condition $\max(a, b) \leq (n/2)^{1/3}$, which we suspect is an artifact of the proof technique and conjecture that, even for our proposed test, it can be removed. In any case, observe that this condition is irrelevant in the setting $a, b = O(\log n)$ considered in this paper. Further, if a/b is bounded away from 1, then TST is directly possible when $a, b = \Omega(\log n)$ by recovering the communities and comparing them, demonstrating that this condition is not present in general.

Small Changes. We claim that for small changes - $s < n^{\frac{1}{2}-\gamma}$ for some $\gamma > 0$ - the naïve scheme of recovering the communities and comparing them is minimax. To see this, note that that GoF testing is reducible to TST - given a TST scheme of a known risk, one may construct a GoF tester of that risk by feeding the TST algorithm the observed graph and a graph drawn from $P(\cdot|x_0)$. Thus, the lower bounds of Theorem 1 apply to TST, and for $a/b = \Theta(1)$, we find that it is necessary that $s\Lambda = \omega(1)$ and that $\Lambda \gtrsim \log(1 + n/s^2)$ to attain vanishing R_{TST} . For small s , the latter lower bound is $\Omega(\log n)$, the claim follows since recovery with up to s errors is possible if $\Lambda \gtrsim \log n$.

Efficiency. Finally, we point out that the above bounds can be attained with computationally efficient tests. Further, for large changes, the test can be made agnostic to knowledge of (a, b) . Instead, it only requires one to be able to estimate $n(a+b)$ to within an additive error of $\tilde{O}(\sqrt{n(a+b)})$, which can be done by simply counting the number of edges in the graphs.

Proof Sketch of the Achievability. We describe the proposed test, and sketch its risk analysis below, completing the same in Appendix B.1. Recall the definition of N_w^z, N_a^z from (3) in §2, and let

$$T^{\hat{x}}(G) := N_w^{\hat{x}}(G) - N_a^{\hat{x}}(G). \quad (4)$$

We show that the routine ‘TwoSampleTester’ below attains a risk smaller than $4/n$. In words, the test computes a partition \hat{x} for the graph G by using about half the edges in the graph. This is represented in the ‘PartialRecovery’ step below, for which any such method may be used - concretely, that of [CRV15]. Next, we compute the statistic $T^{\hat{x}}$ above for both the remaining part of the first graph, and for the second graph. Notice that unlike the GoF statistic, which was only N_a , $T^{\hat{x}}$ takes the difference of N_a and N_w . This is necessary because the partition \hat{x} derived from partial recovery cannot be very well correlated with the true partition x . This means the reduced fluctuations from only considering one part does not apply, and we instead use the whole cut.

Since the edges within communities, and across communities in the graph are (separately) exchangeable, the errors made in \hat{x} distribute uniformly over the two communities⁴. This allows us to explicitly control the behaviour of T as defined in the test *provided \hat{x} is non-trivially correlated with x* - i.e., given that it makes $< (1/2 - c)n$ errors for some $c > 0$. The condition $\Lambda \gtrsim 1$ in the theorem arises from this.

A complication in this strategy is that the remaining graph \tilde{G} in the scheme is not independent of the recovered community \hat{x} . This is handled in the analysis by introducing an independent copy of G , called G' , and arguing that $T^{\hat{x}}(\tilde{G}) \approx 1/2 T^{\hat{x}}(G')$. This step is the origin of the nuisance condition $\max(a, b) \lesssim n^{1/3}$ in the theorem.

Algorithm 1: TwoSampleTester(G, H, δ)

- 1: $G_1 \leftarrow$ subsampling of edges of G at rate $1/2$ uniformly at random.
 - 2: $\tilde{G} \leftarrow G - G_1$.
 - 3: $\hat{x} \leftarrow \text{PartialRecovery}(G_1)$.
 - 4: Compute $T^{\hat{x}}(\tilde{G}), T^{\hat{x}}(H)$.
 - 5: $T \leftarrow |2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(H)|$.
 - 6: Return $T \underset{H_0}{\overset{H_1}{\gtrless}} \sqrt{Cn(a+b)\log(6n)}$.
-

⁴For a proof: since $x, -x$ induce the same law, and since the communities are balanced, for every realization of G such that \hat{x} makes e_+, e_- errors in the community $+, -$ respectively, there is a realization of equal probability where it makes e_-, e_+ errors. Further, within community exchangeability implies that errors distribute uniformly.

Lastly, we point out that while the above exploits the exact balance by using the description of the error distribution it enables, one can derive the same results (but with weakened constants) even without this assumption, so long as both communities are at of size linear in n . In this case, one cannot rely on the errors distributing uniformly over the nodes, but the within-community uniformity of errors, which follows due to within community exchangeability, can be exploited in a similar way. We describe this extension in Appendix B.1.1.

Proof Sketch of the Converse. The necessary condition is shown via Le Cam’s method, but with the twist that the null model is chosen to be a two-step procedure - one that draws a balanced community uniformly at random, and then generates a graph according to it, while the alternate models are drawn uniformly from the balanced communities that are at least s -far from the chosen null. This allows a comparison to the unstructured Erdős-Rényi graph on n vertices with mean degree $(a + b)/2$. Bounds can then be drawn in from the study of the so-called distinguishability problem [Ban+16], and we invoke results from [WX18] to show that total variation distance between the null and alternate distributions is small when Λ is a small enough constant, allowing us to conclude using Neyman-Pearson. See Appendix B.3 for a detailed argument.

4 Experiments

We perform three different sets of numerical experiments. We first run our tests on SBMs with 1000 nodes. Next, we demonstrate that our tests perform similarly for a real dataset, specifically the Political Blogs dataset [AG05]. Finally, we examine SBM-supported Gaussian Markov Random Fields (GMRFs) as an example of a “node observation” model, where the SBM-generated edges form the precision matrix for the Gaussian vector consisting of the random variables assigned to each node. In particular, we need to determine if the underlying community of the graph has changed without explicitly observing (or recovering) the edges of the graph. For the sake of brevity, precise details of the experiments are moved to Appendix C.

4.1 SBM Experiments

We perform experiments implementing our GoF and TST strategies as well as the naïve scheme of reconstructing communities and comparing. Recovery is performed by regularised spectral clustering, for which a detailed description is given in Appendix C.1. The graphs are drawn on $n = 1000$ nodes for a range of (s, Λ) pairs and the high and low risk regimes are plotted in Figure 1. First, note that for ‘large changes,’ $s \geq \sqrt{n \log(10)} \approx 50$, our GoF and TST tests can succeed for lower SNR values. In contrast, for ‘small changes,’ $s < \sqrt{n} \approx 30$, the naïve test is more powerful in the high SNR regime. Additionally, both tests fail for TST unless the SNR is larger than a constant, as predicted by our lower bound in Theorem 2.

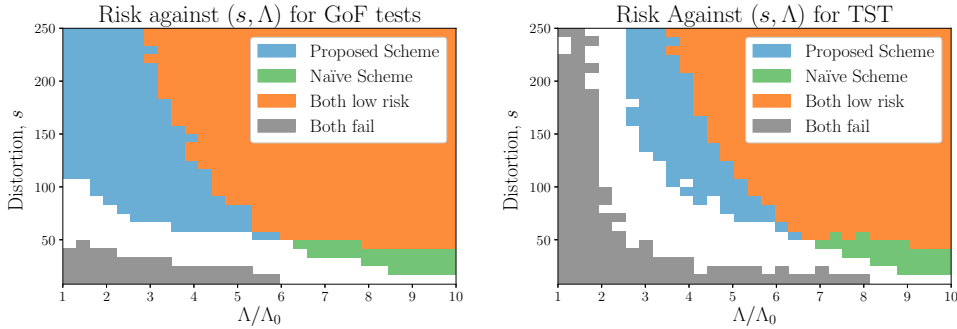


Figure 1: Risks of the proposed tests from Sections 2 and 3 for GoF and TST respectively, and the performance of the naïve scheme, on synthetic SBMs with $n = 1000$, $a/b = 3$. Both schemes attain high risk ($> 1 - \delta$) in the grey region, intermediate risk in the white, and the colours indicate which of the schema attain low risk ($< \delta$), where $\delta = 0.01$ for GoF and $\delta = 0.1$ for TST.

4.2 Political Blogs Dataset [AG05]

The political blogs dataset [AG05] is canonical in the study of community detection, and consists of $n = 1222$ nodes. Here, we vary the effective SNR by randomly subsampling the edges of the graphs at rate ρ . See Appendix C.2 for further details. In this dataset, the ground truth partition x_{True} is available, which in turn yields accurate estimates of the connectivity probabilities (a, b) . For this graph $a/b \approx 10$. Further, spectral clustering alone incurs ≈ 50 errors in this graph, which is larger than $\sqrt{1222} \approx 35$. As a consequence, the behaviour in the ‘small changes’ regime where the test relies on recovery - is not well illustrated in the following.




Figure 2: Risks of the tests applied to the Political Blogs graphs - colour scheme is retained from Fig. 1. The X-axis plots the sparsification factor, which serves as a proxy for SNR. Features similar to Fig. 1 can be seen. The GoF plot improves since a/b is bigger, while the TST plot suffers since the political blogs graph is not completely described as a 2-community SBM [Lei16].

Goodness-of-Fit. We determine the size of the test by running the GoF procedures against x_{True} . To determine power, we construct a partition y by relabelling a random set of nodes of size s , and running the GoF procedures against y with the same graph.

Two-Sample Testing. We compare the political blogs graph G against two other graphs drawn from SBMs. Size is determined by drawing G' according to an SBM of community x_{True} and running the TST procedure, and power is determined by drawing a y as above, generating H according to an SBM of community y , and running the TST procedure. Note that this experiment is thus semi-synthetic.

4.3 Gaussian Markov Random Fields (GMRFs)

Frequently instead of simply receiving a graph, one receives i.i.d. samples from a graph-structured distribution, and it is of interest to be able to cluster nodes with respect to the latent graph. For example, in large-scale calcium imaging, it is possible to simultaneously record the activity pattern of thousands of neurons, but not their underlying synaptic connectivity [Pne+16]. Here, we explore the behavior of our tests for GMRFs where the underlying graph structure is randomly drawn from an SBM and we only observe the nodes.

A heuristic reason for why our methods might succeed in such a situation arises from the local tree-like property of sparse random graphs (see, e.g. [DM10]). For graphs with mean degree $d \ll n$, typical nodes do not lie in cycles shorter than $\sim \frac{\log n}{2 \log d}$. In MRFs, this tree-like property induces correlation decay: the correlation between two nodes decays geometrically up to graph-distance $\sim \frac{\log n}{2 \log d}$. Thus, the covariance matrix closely approximates $\sigma_1 G + \sum_{i=2}^k (\sigma_1 G)^i + \sigma_0 \mathbf{1}\mathbf{1}^T$ for some $\sigma_0 \ll \sigma_1$, small k , and G , the adjacency matrix of the graph. Since the local structure of the graph is so expressed, both clustering and testing applied directly to the covariance matrix should be viable.

We report experimentation on the GMRF (see, e.g. [WJ08, Ch. 3]), which comprises random vectors $\zeta \sim \mathcal{N}(0, \Theta^{-1})$, where the non-zero entries of the precision matrix Θ encode the conditional dependence structure of ζ . Following standard parametrisations [WWR10], we set $\Theta = I + \gamma G$, where $G \sim P(G|x)$ is an adjacency matrix from an SBM with latent parameter x , and γ is a scalar. Below, we fix the SBM parameters a, b and the level γ , and explore risks against s and sample size t .

Following the above heuristic, we naïvely adapt community recovery and testing to this setting, by replacing all instances of the graph adjacency matrix in previous settings with the sample covariance matrix. Figure 3 presents our simulations of the risk of this test when $n = 1000$, and $(a, b) \approx (12.3 \log n, 1.23 \log n)$, at $\Lambda \approx 9 \log(n)$ (for details see Appx. C.3). This large SNR is chosen so that community recovery would be easy if the graph was recovered;⁵ this emphasizes the role of the sample size, t . Importantly, in this implementation, the threshold for rejecting the null has been fit using data (unlike in the previous sections). This is since we lack a rigorous theoretical understanding of this problem, and have not analytically derived expressions for the thresholds. As a result, these plots should be treated as speculative research intended to underscore the presence of interesting testing effects in this scenario, and to encourage future work along these lines.




Figure 3: Risks for adaptation of our tests to GMRFs - colour scheme is retained from Fig. 1. The plots show structural similarity to Fig. 1, but with two differences - In GoF, we don't find a high risk region at the sample sizes considered, and the proposed scheme always outperforms the Naïve scheme based on spectral clustering.

5 Directions for Future Work

The development of the recovery problem for SBMs suggests a number of directions for further work on the testing problems considered above. For instance, one may investigate the exact constants in the testing threshold that the above work suggests, or one may study the testing problem for SBMs with $k > 2$ communities, which is a practically relevant setting since many real-world networks are significantly better described as k -SBMs than as 2-SBMs. In the latter vein, testing problems such as the above may be studied in richer random graph models, such as degree corrected SBMs, or geometric block models. Additionally, testing of strongly imbalanced communities, where one of the communities has size sublinear in n is conceptually unexplored and of interest.

One open problem that draws from the above exposition is if there exists an algorithm for TST in the 2 community setting that does not pass through a partial recovery step and yet works for sparse graphs. We expect that such a method would be necessary for determining exact testing thresholds (for large changes), since the recovery step necessarily requires some subsampling, which reduces the effective SNR available for testing. In addition, this would be conceptually pleasant, and would eliminate the dissonance in the above work where showing testing guarantees requires passing through recovery guarantees. Such a scheme would also more generally allow study of the testing problem for situations where partial recovery is ill understood.

Finally, we mention that more work is needed on the practical investigation of the effectiveness of the above methods - while the experiments we have run validate the theory, the real-world applicability of the methods above require deeper experimentation. A significant lacuna for this line is the lack of a good real-world dataset for the testing of community changes.

⁵Note, however, we expect graph recovery to be impossible at these sample sizes. Lower bounds from [WWR10] indicate this would require > 3300 samples theoretically.

Acknowledgement

We thank the anonymous reviewers of this and of previous versions of this paper, particularly one for suggesting the proof of the converse in Theorem 2.

References

- [Abb18] Emmanuel Abbe. “Community Detection and Stochastic Block Models: Recent Developments”. In: *Journal of Machine Learning Research* 18.177 (2018), pp. 1–86.
- [AG05] Lada A Adamic and Natalie Glance. “The political blogosphere and the 2004 US election: divided they blog”. In: *Proceedings of the 3rd international workshop on Link discovery*. ACM. 2005, pp. 36–43.
- [AV14] Ery Arias-Castro and Nicolas Verzelen. “Community detection in dense random networks”. In: *The Annals of Statistics* 42.3 (2014), pp. 940–969.
- [Ban+16] Jess Banks, Cristopher Moore, Joe Neeman, and Praneeth Netrapalli. “Information-theoretic thresholds for community detection in sparse networks”. In: *Conference on Learning Theory*. 2016, pp. 383–416.
- [Bas+08] Danielle S Bassett, Edward Bullmore, Beth A Verchinski, Venkata S Mattay, Daniel R Weinberger, and Andreas Meyer-Lindenberg. “Hierarchical organization of human cortical networks in health and schizophrenia”. In: *Journal of Neuroscience* 28.37 (2008), pp. 9239–9248.
- [BS16] Peter J Bickel and Purnamrita Sarkar. “Hypothesis testing for automated community detection in networks”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78.1 (2016), pp. 253–273.
- [Chu10] Fan Chung. “Graph theory in the information age”. In: *Notices of the AMS* 57.6 (2010), pp. 726–732.
- [CL06] Fan Chung and Linyuan Lu. *Complex Graphs and Networks (Cbms Regional Conference Series in Mathematics)*. Boston, MA, USA: American Mathematical Society, 2006. ISBN: 0821836579.
- [CRV15] Peter Chin, Anup Rao, and Van Vu. “Stochastic block model and community detection in sparse graphs: A spectral algorithm with optimal rate of recovery”. In: *Conference on Learning Theory*. 2015, pp. 391–423.
- [CY06] Jingchun Chen and Bo Yuan. “Detecting functional modules in the yeast protein–protein interaction network”. In: *Bioinformatics* 22.18 (2006), pp. 2283–2290. DOI: [10.1093/bioinformatics/btl1370](https://doi.org/10.1093/bioinformatics/btl1370).
- [DAM17] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. “Asymptotic mutual information for the balanced binary stochastic block model”. In: *Information and Inference: A Journal of the IMA* 6.2 (2017), pp. 125–170. DOI: [10.1093/imaiai/iaw017](https://doi.org/10.1093/imaiai/iaw017).
- [DM10] Amir Dembo and Andrea Montanari. “Gibbs measures and phase transitions on sparse random graphs”. In: *Brazilian Journal of Probability and Statistics* 24.2 (2010), pp. 137–211.
- [FC19] Yingjie Fei and Yudong Chen. “Exponential error rates of SDP for block models: Beyond Grothendieck’s inequality”. In: *IEEE Transactions on Information Theory* 65.1 (2019), pp. 551–571.
- [For10] Santo Fortunato. “Community detection in graphs”. In: *Physics reports* 486.3-5 (2010), pp. 75–174.
- [Gal68] Robert G Gallager. *Information theory and reliable communication*. Vol. 588. Springer, 1968.
- [Gho+17a] Debarghya Ghoshdastidar, Maurilio Gutzeit, Alexandra Carpentier, and Ulrike von Luxburg. “Two-sample Hypothesis Testing for Inhomogeneous Random Graphs”. In: *arXiv preprint arXiv:1707.00833* (2017).
- [Gho+17b] Debarghya Ghoshdastidar, Maurilio Gutzeit, Alexandra Carpentier, and Ulrike von Luxburg. “Two-Sample Tests for Large Random Graphs Using Network Statistics”. In: *Proceedings of the 2017 Conference on Learning Theory*. 2017, pp. 954–977.
- [GL17] Chao Gao and John Lafferty. “Testing network structure using relations between small subgraph probabilities”. In: *arXiv preprint arXiv:1704.06742* (2017).

- [GL18] Debarghya Ghoshdastidar and Ulrike von Luxburg. “Practical methods for graph two-sample testing”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 3019–3028.
- [GN02] Michelle Girvan and Mark EJ Newman. “Community structure in social and biological networks”. In: *Proceedings of the national academy of sciences* 99.12 (2002), pp. 7821–7826.
- [Hoe63] Wassily Hoeffding. “Probability Inequalities for Sums of Bounded Random Variables”. In: *Journal of the American Statistical Association* 58.301 (1963), pp. 13–30. DOI: [10.1080/01621459.1963.10500830](https://doi.org/10.1080/01621459.1963.10500830).
- [Hun07] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing In Science & Engineering* 9.3 (2007), pp. 90–95. DOI: [10.5281/zenodo.1098480](https://doi.org/10.5281/zenodo.1098480).
- [JOP01] Eric Jones, Travis Oliphant, and Pearu Peterson. *SciPy: Open source scientific tools for Python*. [Online]. 2001. URL: <http://www.scipy.org/>.
- [JTZ04] Daxin Jiang, Chun Tang, and Aidong Zhang. “Cluster analysis for gene expression data: A survey”. In: *IEEE Transactions on Knowledge & Data Engineering* 11 (2004), pp. 1370–1386.
- [JY16] Antony Joseph and Bin Yu. “Impact of regularization on spectral clustering”. In: *Ann. Statist.* 44.4 (2016), pp. 1765–1791. DOI: [10.1214/16-AOS1447](https://doi.org/10.1214/16-AOS1447).
- [Lei16] Jing Lei. “A goodness-of-fit test for stochastic block models”. In: *The Annals of Statistics* 44.1 (2016), pp. 401–424.
- [LL18] Yezheng Li and Hongzhe Li. “Two-sample Test of Community Memberships of Weighted Stochastic Block Models”. In: *arXiv preprint arXiv:1811.12593* (2018).
- [LLV17] Can M Le, Elizaveta Levina, and Roman Vershynin. “Concentration and regularization of random graphs”. In: *Random Structures & Algorithms* 51.3 (2017), pp. 538–561.
- [LR06] Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer Science & Business Media, 2006.
- [Lux07] Ulrike von Luxburg. “A tutorial on spectral clustering”. In: *Statistics and Computing* 17.4 (2007), pp. 395–416. ISSN: 1573-1375. DOI: [10.1007/s11222-007-9033-z](https://doi.org/10.1007/s11222-007-9033-z).
- [MNS15] Elchanan Mossel, Joe Neeman, and Allan Sly. “Reconstruction and estimation in the planted partition model”. In: *Probability Theory and Related Fields* 162.3 (Aug. 2015), pp. 431–461. ISSN: 1432-2064. DOI: [10.1007/s00440-014-0576-6](https://doi.org/10.1007/s00440-014-0576-6).
- [New06] Mark EJ Newman. “Modularity and community structure in networks”. In: *Proceedings of the national academy of sciences* 103.23 (2006), pp. 8577–8582.
- [Oli06] Travis E Oliphant. *A guide to NumPy*. 2006.
- [Ped+11] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [Pne+16] Eftychios A Pnevmatikakis, Daniel Soudry, Yuanjun Gao, Timothy A Machado, Josh Merel, David Pfau, Thomas Reardon, Yu Mu, Clay Lacefield, Weijian Yang, et al. “Simultaneous denoising, deconvolution, and demixing of calcium imaging data”. In: *Neuron* 89.2 (2016), pp. 285–299.
- [Tan+17] Minh Tang, Avanti Athreya, Daniel L. Sussman, Vince Lyzinski, Youngser Park, and Carey E. Priebe. “A Semiparametric Two-Sample Hypothesis Testing Problem for Random Graphs”. In: *Journal of Computational and Graphical Statistics* 26.2 (2017), pp. 344–354. DOI: [10.1080/10618600.2016.1193505](https://doi.org/10.1080/10618600.2016.1193505).
- [VA15] Nicolas Verzelen and Ery Arias-Castro. “Community detection in sparse random networks”. In: *The Annals of Applied Probability* 25.6 (2015), pp. 3465–3510.
- [WJ08] Martin J. Wainwright and Michael I. Jordan. “Graphical Models, Exponential Families, and Variational Inference”. In: *Foundations and Trends® in Machine Learning* 1.1–2 (2008), pp. 1–305. ISSN: 1935-8237. DOI: [10.1561/2200000001](https://doi.org/10.1561/2200000001).
- [WWR10] Wei Wang, Martin J Wainwright, and Kannan Ramchandran. “Information-theoretic bounds on model selection for Gaussian Markov random fields”. In: *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*. IEEE. 2010, pp. 1373–1377.

- [WX18] Yihong Wu and Jiaming Xu. “Statistical Problems with Planted Structures: Information-Theoretical and Computational Limits”. In: *arXiv preprint arXiv:1806.00118* (2018).
- [Zha+08] Bai Zhang, Huai Li, Rebecca B. Riggins, Ming Zhan, Jianhua Xuan, Zhen Zhang, Eric P. Hoffman, Robert Clarke, and Yue Wang. “Differential dependency network analysis to identify condition-specific topological changes in biological networks”. In: *Bioinformatics* 25.4 (2008), pp. 526–532. DOI: [10.1093/bioinformatics/btn660](https://doi.org/10.1093/bioinformatics/btn660).
- [ZZ16] Anderson Y Zhang and Harrison H Zhou. “Minimax rates of community detection in stochastic block models”. In: *The Annals of Statistics* 44.5 (2016), pp. 2252–2280.

Appendix

A Proofs omitted in Section 2

A.1 Proof of Achievability in Theorem 1

We will restrict attention to the case $a > b$ below. The $b > a$ case follows identically. Recall the test in this setting:

$$N_a^{x_0}(G) \underset{H_0}{\overset{H_1}{\geq}} \frac{bn}{4} + C_1 \max(\sqrt{nb \log(2/\delta)}, \log(2/\delta)),$$

where C_1 is the constant implicit in Lemma 3 below.

Under the null distribution, $N_a^{x_0}(G)$ is distributed as $\text{Bin}(n^2/4, b/n)$, while under the alternate, it is distributed as $\text{Bin}((n-s)^2/4 + s^2/4, b/n) * \text{Bin}(s(n-s)/2, a/n)$. These distributions can be separated by Bernstein concentration bounds [CL06, Ch. 2], as summarised by the following Lemma, which is proved in subsequent sections.

Lemma 3. *There exist constants $C_0, C_1 > 1$ such that, if $nb + s(a-b) > C_0 \log(1/\delta)$, then with probability at least $1 - \delta/2$*

$$\begin{aligned} (\alpha) \text{ Under } H_0: N_a^{x_0}(G) &\leq \frac{bn}{4} + C_1 \max\left(\sqrt{nb \log(2/\delta)}, \log(2/\delta)\right). \\ (\beta) \text{ Under } H_1: N_a^{x_0}(G) &\geq \frac{bn}{4} + \frac{s(a-b)}{4} - C_1 \sqrt{(nb + s(a-b)) \log(2/\delta)}. \end{aligned}$$

As the proof of the above lemma discusses, results of the above type hold in the more generic situation where both the communities and the changes can be unbalanced, so long as each community is of at least linear in n size. This allows one to extend the entirety of this theorem to the setting $n^+ n^- = \Omega(n^2)$ on replacing $bn/4$ above with $\mathbb{E}_{\text{Null}}[N_a^{x_0}(G)]$, where n^+ and n^- are the sizes of the two communities, i.e., the number of i such that $x_i = +1$ and $x_i = -1$ respectively.

Since $s|a-b| \geq s\Lambda \geq C \log(2/\delta)$, the lemma above holds in our setting on picking C large enough. (α) in Lemma 3 indicates that the false alarm error of test is $\leq \delta/2$. Further, since $(nb + s(a-b)) \log(2/\delta) > \log^2(2/\delta)$, part (β) shows that missed detection error is $\leq \delta/2$ if

$$\frac{1}{4} s(a-b) > 2C_1 \sqrt{(nb + s(a-b)) \log(2/\delta)} \iff \frac{(a-b)^2}{nb + s(a-b)} > C \frac{\log(2/\delta)}{s^2}.$$

The argument is concluded by some casework:

- (i) If $nb \leq s(a-b)$, then the left hand side of the condition above can be bounded from below by $s(a-b)/2$, and thus $s(a-b) \geq 2C_1 \log(2/\delta)$ is sufficient. But $s(a-b) \geq s(a-b)^2/(a+b) = s\Lambda$ is larger than $C \log(1/\delta)$, and choosing C large enough is sufficient.
- (ii) On the other hand, if $nb > s(a-b)$, the left hand side is instead lower bounded by $s^2(a-b)^2/2nb \geq s^2\Lambda/2n$, and thus $s^2\Lambda \gtrsim n \log(2/\delta)$ is sufficient to satisfy the same. \square

A.1.1 Proof of Lemma 3

The proof proceeds by establishing the centres of the statistic $N_a^{x_0}$ under the null and alternate distributions, and then invoking Bernstein-type bounds [CL06, Ch 2] to show the claimed statements separately.

- (α) For the null, $N_a^{x_0}(G)$ is distributed as $\text{Bin}(n^2/4, b/n)$. Thus, clearly $\mathbb{E}_{\text{Null}}[N] = bn/4$. Further, by Bernstein's inequality for the upper tail,

$$\begin{aligned} P_{\text{Null}}(N_a^{x_0}(G) > \mathbb{E}_{\text{Null}}[N_a^{x_0}(G)] + nt) &\leq \exp\left(-\frac{n^2/4 \times t^2/2}{n^2/4 \times (b/n) + nt/3}\right) \\ &\leq \exp\left(-\frac{3}{2} \frac{nt^2}{b + 4t}\right) \leq \exp\left(-\frac{3}{8} \frac{nt^2}{t + b}\right). \end{aligned}$$

Thus, if

$$\frac{nt^2}{b+t} \geq \frac{8}{3} \log(2/\delta),$$

then this tail has mass at most $\delta/2$. We may now consider the two cases

- (i) If $nb \leq 16/3 \log(2/\delta)$, then plugging in $t = 16/3 \frac{\log(2/\delta)}{n}$ above yields that the condition is satisfied, since then

$$\frac{nt^2}{b+t} \geq \frac{nt^2}{2t} = \frac{nt}{2} = \frac{8}{3} \log(2/\delta).$$

- (ii) If $nb \geq 16/3 \log(2/\delta)$, then setting $t = \sqrt{(16/3) \frac{b}{n} \log 2/\delta}$ we can bound

$$\frac{nt^2}{b+t} = \frac{16/3 \log(2/\delta)}{1 + \sqrt{(16/3) \log(2/\delta)/nb}} \geq \frac{16/3 \log(2/\delta)}{2}.$$

As a consequence, picking $nt = \max(\sqrt{(16/3)nb \log(2/\delta)}, 16/3 \log(2/\delta))$ implies that the probability in question is at most $\delta/2$.

We note that this calculation can be made more robust, in that if the communities are unbalanced but linearly sized with n , then the number of edges crossing is $n^+(n - n^+) = \Omega(n^2)$ in the above, and essentially the same goes through with $n^2/4$ replaced by n^2/C for some constant C .

- (β) This proof proceeds in much the same way as the above. With the modification that the distribution of $N_a^{x_0}(G)$ is now $\text{Bin}(n^2/4 - s(n-s)/2, b/n) * \text{Bin}(s(n-s)/2, a/n)$, since $2 \times s(n-s)/4$ of the edges are now between nodes of the same communities. The centre of this is easily seen to be $\frac{nb}{4} + \frac{s(n-s)}{2} \frac{a-b}{n}$. Further invoking the Bernstein lower tail, we find that

$$\begin{aligned} P_{\text{Alt}}(N_a^{x_0}(G) \leq \mathbb{E}_{\text{Alt}}[N_a^{x_0}(G)] - nt) &\leq \exp\left(-\frac{1}{2} \frac{n^2 t^2}{\frac{s(n-s)}{2} \cdot \frac{a}{n} + \frac{n^2 - 2s(n-s)}{4} \cdot \frac{b}{n}}\right) \\ &\leq \exp\left(-\frac{n^2 t^2}{nb + s(a-b)}\right) \end{aligned}$$

The required claim now follows directly by setting $t = \sqrt{\frac{(nb + s(a-b)) \log(2/\delta)}{n}}$.

Again, the above can also be rendered more robust to imbalance. Suppose that the communities and the changes are both imbalanced, and let n^+, n^- be the sizes of the communities in x_0 , and s^+, s^- be the number of nodes that are moved from $+$ to $-$ and vice-versa according to the alternate x . Then the number of edges which behave according to a/n in the alternate is $\tau = s^+(n^- - s^-) + s^-(n^+ - s^+)$. But $\tau \leq s(n^+ + n^- - s^+ - s^-) = sn$, so the concentration results go through with a weakening of a factor of 2. Further, assume wlog that $s^+ \geq s^-$. since $s^+ + s^- = s$, and $n^+ + n^- = n$, we have that

$$\tau = s^+(n - s) + (s - 2s^+)(n^+ - s^+).$$

Minimising the above subject to $s^+ \in [s/2 : s]$, we find that the minima can be uniformly lower bounded by $\min(s \min(n^+, n^-), s(n-s)/2)$. So long as each community is of linear size, this is $\Omega(sn)$, and thus the centre of the statistic moves by $\Omega(s(a-b))$ with respect to the null statistic.

Putting the two effects above together, we can write that under the alternate distribution, with probability $\geq 1 - \delta/2$,

$$N_a^{x_0}(G) \geq \mathbb{E}_{\text{Null}}[N_a^{x_0}(G)] + \frac{1}{C_1} s(a-b) - C_2 \sqrt{(nb + s(a-b)) \log(2/\delta)}.$$

In conjunction with the discussion for unbalanced but linearly sized communities in case (α), the above allows the claims of the achievability part of Theorem 1 to hold for the case where both communities are of linear size and changes are not constrained to be balanced without any change other than a weakening of the constants implicit in the same. The only modification required for this is to update the tests to threshold at $\mathbb{E}_{\text{Null}}[N_a^{x_0}(G)] + (\text{fluctuation term})$ instead of at $bn/4$ as presented in the main text.

A.2 Proofs of converse bounds from Theorem 1

This section begins with an exposition of Le Cam’s method, which is the general proof strategy we employ to show both these converse bounds. This is followed by separate subsections devoted to each converse bound claimed in Theorem 1.

A.2.1 Le Cam’s method.

The generic lower bound strategy is constructed by noting that the minimax risk of the goodness-of-fit problem is lower bounded by the risk of the same with any given prior on the alternate communities, i.e. the risk of the problem

$$H_0 : x = x_0 \quad \text{vs} \quad H_1 : x \sim \pi$$

for a π supported on $\{x : d(x, x_0) \geq s\}$ (or some restriction of the same, as in the following sections), and the Bayes risk

$$R_\pi := \inf_{\varphi: G \rightarrow \{H_0, H_1\}} P(\varphi = H_1 | x_0) + \sum_{x: d(x, x_0) \geq s} \pi(x) P(\varphi = H_0 | x).$$

By classical Neyman-Pearson theory [see, e.g., [LR06](#)], the likelihood ratio test is optimal under the above risk, and

$$R_\pi = 1 - d_{\text{TV}}(P(G|x_0), \langle P(G|x) \rangle_\pi),$$

where $\langle P(G|x) \rangle_\pi := \sum_x \pi(x) P(G|x)$, and d_{TV} is the total variation distance

$$d_{\text{TV}}(P, Q) := \frac{1}{2} \|P - Q\|_1.$$

We proceed by bounding d_{TV} by an f -divergence more conducive to tensorisation in order to exploit the (conditional) independence of the edges in an SBM, and then by choosing an appropriate π . The f -divergence inequalities we use are

1. χ^2 bound: Recall that

$$D_{\chi^2}(Q||P) = \sum_x P(x) \left(\frac{Q(x) - P(x)}{P(x)} \right)^2 = \mathbb{E}_P[L^2(X)] - 1,$$

where $L(x) := Q(x)/P(x)$ is the likelihood ratio. It holds that

$$d_{\text{TV}}(P, Q) \leq \sqrt{\frac{1}{2} \log(1 + D_{\chi^2}(Q||P))},$$

which follows from Pinsker’s inequality and the fact that

$$D_{\text{KL}}(Q||P) \leq \log(1 + D_{\chi^2}(Q||P)),$$

which is a consequence of Jensen’s inequality applied to the log (or, equivalently, the monotonicity of Rényi divergences).

Invoking the above inequality and Le Cam’s method, we find that for any choice of π , and for $L(G) := \frac{\langle P(G|x) \rangle_\pi}{P(G|x_0)}$, the following is necessary for the minimax risk of the GoF problem to be bounded above by δ :

$$\mathbb{E}_{x_0}[L^2(G)] \geq \exp(2(1 - \delta)^2).$$

For $\delta \leq 1/4$, this yields a necessary lower bound of $\mathbb{E}_{x_0}[L^2] > 3.08$.

2. Hellinger bound: The Bhattacharya coefficient of P, Q is defined as

$$\text{BC}(P, Q) := \sum_x \sqrt{P(x)Q(x)},$$

and the Hellinger divergence as

$$D_{\text{H}}(P, Q) := \sqrt{1 - \text{BC}(P, Q)} = \frac{1}{\sqrt{2}} \|\sqrt{P} - \sqrt{Q}\|_2.$$

We exploit the relation

$$d_{\text{TV}}(P, Q) \leq \sqrt{D_{\text{H}}^2(P, Q)(2 - D_{\text{H}}^2(P, Q))} = \sqrt{1 - \text{BC}^2(P, Q)},$$

which is a consequence of the Cauchy-Schwarz inequality.

Again plugging this in with $Q = \langle P(G|x) \rangle_{\pi}$, we find that in order for the risk to be smaller than δ , we must have that

$$\delta \geq 1 - \sqrt{1 - \text{BC}^2} \geq \frac{\text{BC}^2}{2} \implies \text{BC} \leq \sqrt{2\delta},$$

where $\text{BC} = \text{BC}(\langle P(G|x) \rangle_{\pi}, P(G|x_0))$.

We now proceed to show the claimed bounds. Recall that we are required to show that if $R_{\text{GoF}} < \delta \leq 1/4$, then

$$\Lambda \gtrsim \log(1 + n/s^2) \tag{5}$$

$$s\Lambda \gtrsim \log(1/\delta). \tag{6}$$

A.2.2 Proof of the converse bound (5)

For convenience, we let

$$\nu := (a - b)^2 \left(\frac{1}{a(1 - a/n)} + \frac{1}{b(1 - b/n)} \right). \tag{7}$$

Since $a, b \leq n/2$, and since $a/b = \Theta(1)$, we have $\Lambda \asymp \nu$, and it suffices to show the same bound on the latter.

We invoke Le Cam's method with a χ^2 -bound. Let $m := n/2, t := s/2$ and let x_0 be the partition $([1 : m], [m + 1 : 2m])$.

The alternate prior is chosen to be the uniform prior on the set of alternate partitions constructed as follows. For each $T \subset [1 : m]$, we define the partition

$$\begin{aligned} y_T(+) &= [1 : m] \cup (m + T) \sim T \\ y_T(-) &= [m + 1 : 2m] \cup T \sim (m + T), \end{aligned}$$

where $(m + T) = \{i + m : i \in T\}$. Let $\mathcal{Y}_t := \{y_T : T \subset [1 : m], |T| = t\}$. For conciseness, we define the measures on \mathcal{G} :

$$P_{y_T}(\cdot) := P_T(\cdot) := P(\cdot \mid y_T),$$

and set $P_0 = P(\cdot \mid x_0)$. Further, for convenience, we set $p = a/n$ and $q = b/n$.

For a graph G , we find that $L(G) := \frac{1}{|\mathcal{Y}_t|} \sum_{x \in \mathcal{Y}_t} \frac{P_x(G)}{P_0(G)}$. To invoke Le Cam's method (§A.2.1), we need to upper bound $\mathbb{E}_{P_0}[L^2(G)]$.

To this end, we will define for an edge $e = (u, v)$, and a graph G (which is implicit in the notation)

$$f_e(q, p) := (q/p)^e ((1 - q)/(1 - p))^{1-e}. \tag{8}$$

Above, $f_e(q, p)$ arises as a ratio of the probabilities of a $\text{Bern}(q)$ and a $\text{Bern}(p)$ random variable. Thus, it is the likelihood ratio of an edge being between nodes in the different and in the same community.

First observe that

$$\frac{P_T}{P_0} = \left(\prod_{\substack{i \in [1:m] \sim T, \\ j \in m+T}} f_{ij}(p, q) \right) \left(\prod_{\substack{i \in [m+1:2m] \sim m+T, \\ j \in T}} f_{ij}(p, q) \right) \left(\prod_{\substack{i \in [1:m] \sim T, \\ j \in T}} f_{ij}(q, p) \right) \left(\prod_{\substack{i \in [m+1:2m] \sim m+T, \\ j \in m+T}} f_{ij}(q, p) \right) \tag{9}$$

An important feature of the setup above is that every term in the above product is independently distributed, and wherever $f_{ij}(p, q)$ appears, the corresponding e_{ij} is $\text{Bern}(q)$, and similarly with $f_{ij}(q, p)$ and $\text{Bern}(p)$.

Note that

$$\mathbb{E}_{P_0}[L^2(G)] = \sum_{T_1, T_2 \subseteq [1:m] \text{ of size } t} \mathbb{E}_{P_0} \left[\frac{P_{T_1}(G)P_{T_2}(G)}{P_0^2(G)} \right],$$

and so we must control expectations of this form in order to apply Le Cam's method. Let us fix T_1 and T_2 for now, and partition the nodes into groups as described by the Figure 4⁶.




Figure 4: A schematic of the nodes, partitioned according to their labellings in x_0, y_{T_1}, y_{T_2} . The two ovals denote the partition induced by x_0 into groups marked $+$ and $-$. The section $1F2F^+$ denotes the nodes in the $+$ group whose labels remain *fixed* to $+$ in both y_{T_1}, y_{T_2} . The section marked $1S2F^+$ denotes the nodes in the $+$ group whose labels are *switched* to $-$ in y_{T_1} but remain *fixed* to $+$ in y_{T_2} . Other labels are analogously defined.

Note that in the figure, $1F2F^+ = [1 : m] \sim (T_1 \cup T_2)$, $1S2S^- = (m + T_1) \cap (m + T_2)$ and so on. Also, importantly, the size of groups with the same number of S s and F s in the above representation is identical (i.e., $|1F2S^+| = |1F2S^-| = |1S2F^+| = |1S2F^-|$ and so on.)

We consider how the terms relating to the edge (u, v) for any $u, v \in [1 : 2m]$ appear in the product $\frac{P_{T_1}P_{T_2}}{P_0^2}$. Below,

- Clearly, if u and v are both in the same group in both settings, the behaviour of the edge (u, v) under the alternate distributions and the null distribution is identical, and these terms will not appear in the product.
- If both $(u, v) \in 1F2F^+ \times 1F2F^- \cup 1S2S^+ \times 1S2S^-$, then again, the edge (u, v) has identical distribution under both alternates and the null, and these terms do not appear in the product.
- If $(u, v) \in 1F2F^+ \times 1F2S^+$, then the (u, v) term does not appear in P_{T_1}/P_0 , but appears once in P_{T_2}/P_0 . Since likelihoods must average to 1, and since the distributions of the edges are independent, any term which appears just once is averaged out when we take expectations with respect to P_0 . Thus, even though these terms appear in the product, we may ignore them due to our eventual use of the expectation operator. A quick check will show that the same effect happens for $(u, v) \in \Gamma_1 \times \Gamma_2$, where Γ_1 can be obtained by inverting one instance of an F to a S or vice versa, and possibly changing the sign (e.g. $1F2S^- \times 1S2S^+$.) Thus, all such pairs can be safely ignored.

⁶The argument, while simple, gets a little notationally hairy at this point. We recommend that the reader consults Figure 4 frequently, preferably a printed copy that allows one to sketch the various types of connections on it.

- This leaves us with edges of the form $\{1F2F^\pm \times 1S2S^\pm\} \cup \{1F2S^\pm \times 1S2F^\pm\}$. In these cases, if the signs of the two choices match - i.e.

$$(u, v) \in \Gamma^+ \times \tilde{\Gamma}^+ \text{ for } (\Gamma, \tilde{\Gamma}) \in \{(1F2F, 1S2S), (1S2F, 1F2S)\},$$

then we will obtain a contribution of $f_{uv}(q, p)^2$ to the product. On the other hand, if they differ, then we will obtain a contribution of $f_{uv}(p, q)^2$

Accounting for the above, and taking expectation, we have that

$$\mathbb{E} \left[\frac{P_{T_1} P_{T_2}}{P_0^2} \right] = (\Psi)^{|1F2F^+| \cdot |1S2S^+| + |1F2F^-| \cdot |1S2S^-| + |1S2F^+| \cdot |1F1S^-| + |1F2S^+| \cdot |1S2F^-|}, \quad (10)$$

where

$$\Psi := \mathbb{E}_{e \sim \text{Bern}(p)} [f_e(q, p)^2] \mathbb{E}_{e \sim \text{Bern}(q)} [f_e(p, q)^2] \quad (11)$$

Further, since in our choice of the alternate communities the groups with the same number of S s and F s have identical size, and thus we may rewrite (10) above as

$$\mathbb{E} \left[\frac{P_{T_1} P_{T_2}}{P_0^2} \right] = \Psi^{2(|1F2F^+| + |1S2S^+| + |1S2F^+|^2)}.$$

For convenience, let $|1S2S^+| = |T_1 \cap T_2| = k$. We then have that $|1S2F^+| = t - k$ and $|1F2F^+| = m + k - 2t$.

We thus have that

$$\mathbb{E}_{P_0} \frac{P_{T_1} P_{T_2}}{P_0^2} = \exp((\log \Psi)(2k(m + k - 2t) + 2(t - k)^2)) \quad (12)$$

$$= \exp((\log \Psi)(2mk + 2k^2 - 4kt + 2k^2 + 2t^2 - 4kt)) \quad (13)$$

$$= \exp((\log \Psi)(2mk + 4k^2 - 8kt + 2t^2)) \quad (14)$$

$$\leq \exp((\log \Psi)((2m - 4t)k + 2t^2)), \quad (15)$$

where we have used that $k \leq t$.

Now, for $(p, q) = (a/n, b/n)$,

$$\Psi = \left(\frac{q^2}{p} + \frac{(1-q)^2}{(1-p)} \right) \left(\frac{p^2}{q} + \frac{(1-p)^2}{(1-q)} \right) \quad (16)$$

$$= \left(1 + \frac{(p-q)^2}{p(1-p)} \right) \left(1 + \frac{(p-q)^2}{q(1-q)} \right) \quad (17)$$

$$= \left(1 + \frac{(a-b)^2}{na(1-a/n)} \right) \left(1 + \frac{(a-b)^2}{nb(1-b/n)} \right) \quad (18)$$

$$= 1 + \frac{\nu}{n} + O(n^{-2}) \leq 1 + 2\frac{\nu}{n}. \quad (19)$$

As a consequence, using $2m = n$, and the development above,

$$\mathbb{E}_{P_0} \frac{P_{T_1} P_{T_2}}{P_0^2} \leq \exp\left(\frac{4t^2}{n}\nu\right) \exp(2k\nu(1 - 4t/n)). \quad (20)$$

The above is insular to the precise identities of T_1, T_2 . Further, for a given T_1 , the number of partitions T_2 such that $|T_1 \cap T_2| = t$ is $\binom{t}{k} \binom{m-t}{t-k}$. Feeding this into the expression for $\mathbb{E}[L^2(G)]$ and some simple manipulations yield that

$$\mathbb{E}_{P_0}[L^2(G)] \leq \frac{e^{\frac{4t^2}{n}\nu}}{\binom{m}{t}} \sum_{k=0}^t \binom{t}{k} \binom{m-t}{t-k} \exp(2k\nu(1 - 4t/n)), \quad (21)$$

where we remind the reader that $t = s/2, m = n/2$.

Recall from §A.2.1 that if $\mathbb{E}_{P_0}[L^2] < 3$, then the risk exceeds 0.25. Thus, we will aim to upper bound (21) by 3.

We begin by rewriting

$$\mathbb{E}_{P_0}[L^2(G)] \leq \frac{e^{\frac{4t^2}{n}\nu}}{\binom{m}{t}} \sum_{k=0}^t \binom{t}{k} \binom{m-t}{t-k} \exp(2k\nu(1-4t/n)), \quad (22)$$

$$= e^{\frac{4t^2}{n}\nu} \mathbb{E}[\xi^Z], \quad (23)$$

where $\xi := \exp(2\nu(1-4t/n)) > 1$ and $Z = \sum_{i=1}^t Z_i$, where Z_i are sampled without replacement from the collection of t (+1)s and $m-t$ (0)s. Note that $z \mapsto \xi^z$ is continuous and convex for $\xi \geq 1$. By Theorem 4 of [Hoe63],

$$\mathbb{E}[\xi^Z] \leq \mathbb{E}[\xi^{\tilde{Z}}],$$

for $\tilde{Z} = \sum_{i=1}^t \tilde{Z}_i$, where \tilde{Z}_i are drawn by sampling with replacement from the same collection. But \tilde{Z} is just a Binomial random variable with parameters $(t, t/m)$. Thus, we have that

$$\mathbb{E}_{P_0}[L^2(G)] \leq e^{\frac{2t^2}{m}\nu} \left(1 + \frac{t}{m} (\exp(2\nu(1-2t/m)) - 1)\right)^t \quad (24)$$

$$\leq \exp\left(2\frac{t^2}{m}\nu + \frac{t^2}{m} (\exp(2\nu(1-2t/m)) - 1)\right) \quad (25)$$

$$\leq \exp\left(\frac{t^2}{m} (2\nu + e^{2\nu} - 1)\right) \quad (26)$$

$$\leq \exp\left(2\frac{t^2}{m} (e^{2\nu} - 1)\right), \quad (27)$$

where the final inequality uses $u < e^u - 1$. Using the above, and noting that $m/2t^2 = n/s^2$, we find that

$$\nu \leq \frac{1}{2} \log\left(1 + \frac{\log(3)n}{s^2}\right) \implies \mathbb{E}_{P_0}[L^2(G)] \leq 3,$$

finishing the argument. \square

A.2.3 Proof of the converse bound (6)

Recall that this part of the theorem claims that if $R_{\text{Gof}} \leq \delta \leq 1/4$, then $s\Lambda \geq C \log(1/\delta)$.

We will again use Le Cam's method (§A.2.1), this time controlling the total variation distance by a Hellinger bound.

Let $x_0 = ([1 : n/2], [n/2+1 : n])$ be the null partition, and $\mathcal{Y} := \{y\}$, with $y := ([1 : n/2-s/2] \cup [n/2+1 : n/2+s/2], [n/2-s/2+1 : n/2] \cup [n/2+s/2+1 : n])$. We let $P_{x_0}(G) := P(G|x_0)$, and similarly P_y . Recall from the section on Le Cam's method that the following is a necessary condition for the risk to be smaller than δ

$$\text{BC}(P_{x_0}, P_y) \leq \sqrt{2\delta}.$$

The Bhattacharya Coefficient can be estimated directly in this setting. (We omit the derivation below)

$$\text{BC}(P_y, P_{x_0}) = \left(\sqrt{\frac{ab}{n^2}} + \sqrt{\left(1 - \frac{a}{n}\right) \left(1 - \frac{b}{n}\right)} \right)^{s(n-s)} \quad (28)$$

For $u, v < 3/4$,

$$\sqrt{(1-u)(1-v)} \geq 1 - (u+v)/2 - 2(u-v)^2.$$

Thus

$$\text{BC}(P_y, P_{x_0}) \geq \left(1 - \frac{a+b}{2n} + \frac{ab}{n} - 2\frac{(a-b)^2}{n^2}\right)^{s(n-s)} \quad (29)$$

$$= \left(1 - \frac{(\sqrt{a} - \sqrt{b})^2}{2n} - 2\frac{(a-b)^2}{n^2}\right)^{s(n-s)} \quad (30)$$

$$\geq \left(1 - \frac{(\sqrt{a} - \sqrt{b})^2}{n}\right)^{s(n-s)} \quad (31)$$

$$\geq \exp\left(-2s(\sqrt{a} - \sqrt{b})^2\right), \quad (32)$$

where the third inequality uses $(a+b) < n/4$, and the final uses $1 - u \geq e^{-2u}$ for $0 < u \leq 0.75$ —which applies since $0 < (\sqrt{a} - \sqrt{b})^2 < \max(a, b) < n/4$ —and $n - s \leq n$.

Now note that

$$(\sqrt{a} - \sqrt{b})^2 = \frac{(a-b)^2}{(\sqrt{a} + \sqrt{b})^2} \leq \frac{(a-b)^2}{a+b} = \Lambda,$$

and thus,

$$\text{BC}(P_y, P_{x_0}) \geq \exp(-2s\Lambda).$$

Invoking the condition for $R_{\text{GoF}} \leq \delta$ above, we have

$$\begin{aligned} \exp(-2s\Lambda) &\leq \sqrt{2\delta} \\ \iff s\Lambda &\geq \frac{1}{4} \log \frac{1}{2\delta}. \end{aligned}$$

For $\delta \leq 1/4$, we may further lower bound the above by $(\log(1/\delta))/8$. \square

A.3 A comment on the role of Λ when $a/b \neq \Theta(1)$

The main text concentrates on the setting where a/b is a constant. Here, we briefly comment on the setting where the ratio $\rho := \frac{\max(a,b)}{\min(a,b)}$ is *diverging* with n . In the setting of balanced communities and divergent ρ , the behaviour of the goodness-of-fit problem is no longer described by the quantity $\Lambda = \frac{(a-b)^2}{a+b}$, but instead depends on

$$\mu := \frac{(a-b)^2}{\min(a,b)}.$$

Specifically, our proofs can, with minimal changes, be adapted to say that for balanced GoF, R_{GoF} can be solved with vanishing risk if the following hold:

$$\begin{aligned} s\Lambda &= \omega(1) \\ \mu &= \omega(n/s^2), \end{aligned}$$

and further, to attain the same, it is necessary to have

$$\begin{aligned} s\Lambda &= \omega(1) \\ \mu &\gtrsim \log(1 + n/s^2). \end{aligned}$$

Indeed, for the lower bounds, $\mu \leq \nu \leq 4\mu$ uniformly, where ν is the SNR quantity in the previous section, and the upper bounds naturally feature μ .

Together, the above offer a tight characterisation of the GoF problem in the setting of balanced communities and *large* s . Note that $\mu/\Lambda = 1 + \rho$ diverges with ρ , and thus the above indicate that GoF testing becomes much easier as this ratio blows up - something to be expected.

Despite the above developments, we concentrated on the setting $\rho = \Theta(1)$ in the main text. This is largely because the majority of the literature on the SBM focuses on this regime, as this is the

hardest setting for inference about the planted structure. Thus, in order to compare to existing work, we examined the $a \asymp b$ setting.

As an aside, we note that unlike the above GoF results, the TST results do not alter in the setting of divergent ρ . Theorem 2, and in particular the converse bound $\Lambda \gtrsim 1$, continues to hold for this setting.

On the whole, this line of work is still under investigation, particularly whether the behaviour of GoF for large ρ continues to be driven by μ in the setting of small changes. We plan to explore this question in later work.

B Proofs omitted from section 3

B.1 Proof of Achievability in Theorem 2

We carry out the analysis for the case $a > b$. The reverse can be handled similarly. Note that this assumption will implicitly be made in all the lemmata that follow.

Recall that the scheme in Algorithm 1 utilises a partial recovery routine. For the purposes of the following argument, we invoke the method of [CRV15], which provides a procedure that, under the conditions of the theorem, that attains with probability at least $1 - 1/n$ recovery with at most $\varepsilon_{\max}n$ errors, where $\varepsilon_{\max} = \min(1/2, 2e^{-C\Lambda})$ for an explicit constant C . We choose Λ large enough so that ε_{\max} is bounded strictly below $1/2$ - for convenience, say by $1/3$.

Let $G' \sim P(\cdot|x)$ be an independent copy of G , useful in the analysis, and recall the definition of \tilde{G}, G_1 from Algorithm 1. We define the following events that we will condition on in the sequel:

$$\mathcal{E}(G_1) = \{\text{Number of edges in } G_1 \leq an/2\} \quad \mathcal{E}(\hat{x}) = \{d(\hat{x}, x) \leq \varepsilon_{\max}n\}$$

For succinctness, we let $\mathcal{E} := \mathcal{E}(G_1) \cap \mathcal{E}(\hat{x})$. The analysis proceeds in four steps:

(L1) **Lemma 4.** $P(\mathcal{E}) \geq 1 - 4/3n$.

(L2) **Lemma 5.** $\left| \mathbb{E}[2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') \mid \mathcal{E}] \right| \leq a^2$.

(L3) **Lemma 6.** If $d(x, y) \geq s$, then for $\kappa := (1 - 2\varepsilon_{\max})^2 - 1/(n - 1)$,

$$\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \mathcal{E}] \geq \kappa \frac{(a - b)}{n} (n - s)s.$$

(L4) **Lemma 7.** Let $\xi := a^2 + 5\sqrt{2na \log(6n)}$. Then

$$\begin{aligned} P_{\text{Null}}(T \geq \xi \mid \mathcal{E}) &\leq 2/3n \\ P_{\text{Alt.}}(T \leq \kappa(a - b)s/2 - \xi \mid \mathcal{E}) &\leq 4/3n \end{aligned}$$

We briefly describe the functional roles of the above, and relegate their proofs to the following sections.

(L1) allows us to make use of the typicality of G_1 and the recovery guarantees of \hat{x} . The former is primarily useful for (L2), while the latter induces (L3).

(L2) lets us avoid the technical issues arising from the fact \tilde{G} and G_1, \hat{x} are correlated, and allows us to work with the simpler G' . It also shows that under the null, the mean of T is small. This lemma is likely loose, and introduces the nuisance condition $a \leq n^{1/3}$.

(L3) shows that under the alternate, the centre of T linearly grows with s despite the weak recovery procedure's errors.

(L4) serves to control the fluctuations in T . The \sqrt{n} -level term arises from the randomness in \tilde{G}, H, G' , and the a^2 term from our use of G' and (L2).

Putting the above together, we find that the risk is bounded by $4/3n + 2/3n + 4/3n \leq 4/n$ if

$$\kappa(a - b)s \geq 4(a^2 + 5\sqrt{2na \log(6n)}).$$

Since $a^2 = a^{3/2}\sqrt{a} \leq \sqrt{na}$, and for Λ a large enough constant, $\varepsilon_{\max} \leq 1/3 \implies \kappa \geq (1/3 - 1/(n-1))^2 \geq 1/36$ for $n \geq 7$, the above condition is equivalent to

$$(a-b)s \geq C' \sqrt{na \log(6n)}$$

for a large enough C' . Rearranging and squaring, this is equivalent to

$$\frac{(a-b)^2}{a} \gtrsim \frac{n \log(6n)}{s^2}.$$

For $s \geq n^{1/2+c}$ as in the statement, the quantity on the right hand side is decaying with n . Further, Λ is smaller than the left hand side, so it being bigger than a constant forces the above to hold.

Note that the threshold in Algorithm 1 alters the fluctuation range above from \sqrt{na} to $\sqrt{n(a+b)}$. The reason for this is that this relaxation allows Algorithm 1 to be agnostic to the knowledge of (a, b) - generic spectral clustering schemes do not require this knowledge, and the threshold of our scheme depends only on $n(a+b)$, which can be robustly estimated in our setting since the number of edges in the graph is proportional to this. In addition, invoking the bounds of [CRV15] allows explicit control on κ above, and thus provides an explicit value of the constant C in Algorithm 1. \square

B.1.1 Relaxing Exact Balance for TST

We briefly discuss the modifications required to the above analysis in order to extend the same to unbalanced but linearly sized communities. Of the four lemmata used in the proof described above, the proof of Lemma 5 is completely agnostic to the sizes of the communities. In addition, while the proof presented in [CRV15] concentrates on the case of exactly balanced communities, as noted in their Section 1, it can be extended to unbalanced but linearly sized communities with minimal changes, although with a corresponding weakening of the constants in the rate with which error decays with increasing Λ . This extends Lemma 4 to linearly sized communities.

In contrast, Lemma 6 does rely on the assumption of balance. To sidestep this, we show the following version for use in this setting:

Lemma 8. *Let the communities be of sizes n_+, n_- . If $d(x, y) \geq s$, then*

$$\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \mathcal{E}] \geq \frac{(a-b)s}{2} \left(1 - 2\varepsilon \left(\frac{n}{\min(n_+, n_-)} + 2 \right) \right).$$

When $\min(n_+, n_-) = cn$, for some $c > 0$, then by enforcing that ε is smaller than, say, $(3(2 + 1/c))^{-1}$, which may be done by choosing a large enough, but $O(1)$, value of Λ , the above can be expressed as $\Omega(s(a-b))$.

Lastly, the alternate case in Lemma 7 must be adjusted. However, this concentration result is actually proved by arguing that the event $\{T \leq \mathbb{E}_{\text{Alt}}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \mathcal{E}] - \xi\}$ has low probability given \mathcal{E} , and the above lemma implies that this expectation is at least $\Omega(s(a-b))$, so the corresponding tail bound goes through to the required form trivially.

At this point, the concluding remarks of the above proof apply to finish the argument.

B.2 Proofs of Lemmata used in B.1

B.2.1 Proof of Lemma 4

We first note that by the work of [CRV15], or [FC19], under the conditions of the theorem, $\mathcal{E}(\hat{x})$ holds with probability at least $1 - 1/n$. By a union bound, it suffices to show that $P(\mathcal{E}(G_1)) \geq 1 - \frac{1}{3n}$. Recall that

$$P((e, v) \in G_1 \mid x) = \begin{cases} \frac{a}{2n} & x_u = x_v \\ \frac{b}{2n} & x_u \neq x_v, \end{cases} \quad (33)$$

and that edges are independent. Thus the number of edges in G_1 is a sum of Bernoulli random variables of parameter $\leq a/2n$. The factor of 2 arises since G_1 is sub-sampled at rate $1/2$. Let $\#G_1$ be the number of edges in G_1 . We have

$$\mathbb{E}[\#G_1] \leq \binom{n}{2} \frac{a}{2n} \leq \frac{na}{4} \quad (34)$$

$$P(\#G_1 \geq \mathbb{E}[\#G_1] + \sqrt{na \log(3n)}) \leq 1/3n, \quad (35)$$

where the first bound follows from inspection, and the second follows from the Bernstein upper tail bound of [CL06, Ch. 2] and the condition $a \geq 16 \log(6n)/n$. Further invoking this condition we find that $\sqrt{na \log(3n)} \leq na/4$, and thus

$$P(\mathcal{E}(G_1)) = P(\#G_1 \leq na/2) \geq 1 - \frac{1}{3n}. \quad \square$$

B.2.2 Proof of Lemma 5

Let

$$c_{uv} := \frac{(a+b) + (a-b)x_u x_v}{2} \leq a.$$

Recall that c_{uv}/n is the probability under x of the edge (u, v) existing.

Also note that for a graph Γ and a partition z ,

$$T^z(\Gamma) = \sum_{1 \leq u < v \leq n} z_u z_v \Gamma_{uv},$$

where $\Gamma_{uv} := \mathbf{1}\{(u, v) \in \Gamma\}$.

We're interested in controlling

$$T = 2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') = \sum \hat{x}_u \hat{x}_v (2\tilde{G}_{uv} - G'_{uv}).$$

Since \hat{x} is a deterministic function of G_1 , \tilde{G} is independent of \hat{x} given G_1 . Further, G' is independent of (G_1, \tilde{G}) . Lastly observe that

$$P((u, v) \in \tilde{G} \mid G_1) = \frac{c_{uv}/2n}{1 - c_{uv}/2n} (1 - (G_1)_{uv}).$$

As a consequence,

$$\mathbb{E}[T \mid G_1] = \sum \hat{x}_u \hat{x}_v \left(2 \cdot \frac{c_{uv}/2n}{1 - c_{uv}/2n} (1 - (G_1)_{uv}) - \frac{c_{uv}}{n} \right) \quad (36)$$

$$= \sum \hat{x}_u \hat{x}_v \frac{c_{uv}^2/2n^2}{1 - c_{uv}/2n} - \sum \hat{x}_u \hat{x}_v \frac{c_{uv}/n}{1 - c_{uv}/2n} (G_1)_{uv} \quad (37)$$

$$\implies |\mathbb{E}[T \mid G_1]| \leq \sum_{u < v} \frac{c_{uv}^2/2n^2}{1 - c_{uv}/2n} + \sum_{u < v} \frac{c_{uv}/n}{1 - c_{uv}/2n} (G_1)_{uv} \quad (38)$$

$$\leq \frac{a^2/2n^2}{1 - a/2n} \binom{n}{2} + \frac{a/n}{1 - a/2n} \#G_1. \quad (39)$$

where recall that $\#G_1$ is the number of edges in G_1 . Note that we may condition on \mathcal{E} , the occurrence of which is a deterministic function of G_1 . Since under \mathcal{E} we have $\#G_1 \leq an/2$, we find that

$$|\mathbb{E}[T \mid G_1, \mathcal{E}]| \leq \frac{1}{1 - a/2n} \left(\frac{a^2}{2n^2} \frac{n^2}{2} + \frac{a}{n} \frac{an}{2} \right) \leq a^2, \quad (40)$$

where the final inequality uses that $1/(1 - a/2n) \leq 4/3$, which follows from $a \leq (n/2)^{1/3}$, and $n \geq 2$.

Finally observe that the right hand side of the equation above does not depend on G_1 . Thus, we may integrate over $P(G_1 \mid \mathcal{E})$ to find that $|\mathbb{E}[T \mid \mathcal{E}]| \leq a^2$.

Remark This lemma is likely rather weak. In particular, the upper bound on $|\mathbb{E}[T \mid G_1]|$ completely ignores the relationship between \hat{x} & G_1 , and that between G_1 & c_{uv} . Indeed, (36) may also be rewritten as

$$\mathbb{E}[T \mid G_1] = \sum \frac{c_{uv}/n}{1 - c_{uv}/2n} \hat{x}_u \hat{x}_v \left(\frac{c_{uv}}{2n} - (G_1)_{uv} \right).$$

Since $(G_1)_{uv} \sim \text{Bern}(c_{uv}/2n)$, and \hat{x} is a clustering derived from G_1 , it may be possible to control the above to something much smaller than a^2 . This may require nontrivial use of the $\mathcal{E}(\hat{x})$ conditioning here, which is unused in the above argument. Unfortunately it seems that such control would closely depend on the scheme used to obtain \hat{x} , which tend to be complex - most schemes involve non-trivial regularisation of G_1 , as well as some amount of quantisation of the solution to an optimisation problem to produce \hat{x} , due to which the covariance of G_1 and \hat{x} is difficult to understand. For completeness' sake we point out that an upper bound on the same of $O(a^2/n)$ would remove the nuisance condition of $a \leq n^{1/3}$ present in Theorem 2. \square

B.2.3 Setting up the Proof of Lemma 6

We proceed by first developing some intuition behind the proof of Lemma 6 instead of launching straight into the same. Further, we assume throughout that $d(x, y) \geq s$.

Let

$$\begin{aligned} \text{Incorrect} &:= \{u \in [1 : n] : x(u) \neq \hat{x}(u)\} \\ \text{Unchanged} &:= \{u \in [1 : n] : x(u) = y(u)\}. \end{aligned}$$

and the sets 'Correct' and 'Changed' be their respective complements. We show in Appendix B.2.4 the following lemma

Lemma 9.

$$\begin{aligned} \mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \hat{x}] &= \frac{(a-b)}{n} \left(n(\text{Unchanged}) - 2n(\text{Incorrect}, \text{Unchanged}) \right) \\ &\quad \times \left(n(\text{Changed}) - 2n(\text{Incorrect}, \text{Changed}) \right), \end{aligned} \quad (41)$$

where

$$\begin{aligned} n(\text{Unchanged}) &= |\text{Unchanged}| \\ n(\text{Incorrect}, \text{Unchanged}) &= |\text{Incorrect} \cap \text{Unchanged}|, \end{aligned}$$

and the other terms are defined analogously.

We note that the above lemma holds irrespective of the balance assumptions in the theorem.

Suppose $n(\text{Incorrect}) = k$. Due to the exchangeability of the nodes when $|\{u : x(u) = +\}| = |\{u : x(u) = -\}|$, the incorrectly labelled nodes in \hat{x} correspond to a choice of $k \in [0 : n/2]$ nodes picked without replacement from $[1 : n]$ uniformly at random. Further, since the changes made in y are chosen independently of the graphs, they are independent of \hat{x} . Thus, the number of correct and incorrect nodes changed forms hypergeometric distribution. The expected number of Incorrect nodes changed is precisely $\frac{s}{n} \cdot k$, where s is the number of changes made, and similarly for Incorrect nodes unchanged.

Further invoking the results of [FC19], if $\Lambda \geq C \log(1/\varepsilon_{\max})$, then $k \leq \varepsilon_{\max} n$ with probability at least $1 - 1/n$. As a consequence, the bound in Lemma 9 remains large in magnitude even on integrating over the randomness in \hat{x} . This was the subject of Lemma 6 from the text, reproduced below for convenience.

Lemma 6

$$\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \mathcal{E}] \geq \left((1 - 2\varepsilon_{\max})^2 - \frac{1}{n-1} \right) \frac{(a-b)}{n} (n-s)s,$$

the proof of which is the subject of Appendix B.2.5.

B.2.4 Proof of Lemma 9

We will require explicit counting of a number of groups of nodes. Let us first define them:

Let

$$\begin{aligned} S^{++} &:= \{u \in [1 : n] : \hat{x}(u) = +1, x(u) = +1\}, & n^{++} &:= |S^{++}|, \\ S^{+-} &:= \{u \in [1 : n] : \hat{x}(u) = +1, x(u) = -1\}, & n^{+-} &:= |S^{+-}|, \\ S^{-+} &:= \{u \in [1 : n] : \hat{x}(u) = -1, x(u) = +1\}, & n^{-+} &:= |S^{-+}|, \\ S^{--} &:= \{u \in [1 : n] : \hat{x}(u) = -1, x(u) = -1\}, & n^{--} &:= |S^{--}|, \end{aligned}$$

The sets above encode the partitions induced by \hat{x} and x , with the first symbol in the superscript denoting the label given by \hat{x} . Observe that S^{+-}, S^{-+} are the sets of nodes mislabelled in \hat{x} .

Lastly, for $(i, j) \in \{+, -\}^2$, let

$$\begin{aligned} C^{i,j} &:= S^{i,j} \cap \{u \in [1 : n] : x(u) \neq y(u)\} \\ n_C^{i,j} &:= |C^{i,j}| \end{aligned}$$

These are the nodes that change their labels in y . Note that the values of each of the above objects is a function of \hat{x} . For now we will fix \hat{x} , and compute expectations over the randomness in G', H alone.

We first study N_w : $N_w^{\hat{x}}(G) = N_w^{\hat{x}}(G[+]) + N_w^{\hat{x}}(G[-])$, where $G[+]$ is the induced subgraph on the nodes $\{u \in [1 : n] : \hat{x}(u) = +\}$ and similarly $G[-]$.

By simple counting arguments,

$$\mathbb{E}[N_w^{\hat{x}}(G'[+]) \mid \hat{x}] = \binom{n^{++} + n^{+-}}{2} \frac{a}{n} - \frac{(a-b)}{n} n^{++} n^{+-}. \quad (42)$$

Under H , the nodes in C^{++} behave as if they were in S^{+-} and those in C^{+-} as if they were in S^{++} . Computations analogous to before lead to

$$\mathbb{E}[N_w^{\hat{x}}(G'[+]) - N_w^{\hat{x}}(H[+]) \mid \hat{x}] = \frac{a-b}{n} ((n^{++} - n_C^{++}) - (n^{+-} - n_C^{+-})) (n_C^{++} - n_C^{+-}) \quad (43)$$

By symmetry, we can obtain the above for $G[-]$ s by toggling the group labels above. Thus, conditioned on a fixed \hat{x} , we have

$$\begin{aligned} \mathbb{E}[N_w^{\hat{x}}(G') - N_w^{\hat{x}}(H) \mid \hat{x}] &= \frac{(a-b)}{n} \left(((n^{++} - n_C^{++}) - (n^{+-} - n_C^{+-})) (n_C^{++} - n_C^{+-}) \right. \\ &\quad \left. + ((n^{--} - n_C^{--}) - (n^{-+} - n_C^{-+})) (n_C^{--} - n_C^{-+}) \right). \quad (44) \end{aligned}$$

Similar calculations can be performed for N_a . Since in edges across the true partitions, the edges in the same group appear with probability a/n and in different groups with b/n , the roles of a and b will be exchanged in this case, leading to a factor of $+(a-b)$ instead of $-(a-b)$. We will suppress the tedious computations, and simply state that

$$\begin{aligned} \mathbb{E}[N_a^{\hat{x}}(G') - N_a^{\hat{x}}(H) \mid \hat{x}] &= \frac{(a-b)}{n} \left(((n^{++} - n_C^{++}) - (n^{+-} - n_C^{+-})) (n_C^{--} - n_C^{-+}) \right. \\ &\quad \left. + ((n^{--} - n_C^{--}) - (n^{-+} - n_C^{-+})) (n_C^{++} - n_C^{+-}) \right). \quad (45) \end{aligned}$$

For convenience, we define

$$\begin{aligned} n(\text{Correct, Unchanged}) &:= (n^{++} + n^{--}) - (n_C^{++} + n_C^{--}) \\ n(\text{Correct, Changed}) &:= (n_C^{++} + n_C^{--}) \\ n(\text{Incorrect, Unchanged}) &:= (n^{+-} + n^{-+}) - (n_C^{+-} + n_C^{-+}) \\ n(\text{Incorrect, Changed}) &:= (n_C^{+-} + n_C^{-+}) \end{aligned}$$

where ‘correctness’ corresponds to the nodes u such that $\hat{x}(u) = x(u)$, while ‘unchangedness’ to u such that $x(u) = y(u)$.

Subtracting (45) from (44) then yields that for fixed \hat{x}

$$\begin{aligned} \mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \hat{x}] &= \frac{(a-b)}{n} \left(n(\text{Correct, Unchanged}) - n(\text{Incorrect, Unchanged}) \right) \\ &\quad \times \left(n(\text{Correct, Changed}) - n(\text{Incorrect, Changed}) \right). \end{aligned} \quad (46)$$

The lemma now follows on observing that

$$n(\text{Unchanged}) = n(\text{Correct, Unchanged}) + n(\text{Incorrect, Unchanged}),$$

and similarly $n(\text{Changed})$. \square

B.2.5 Proof of Lemma 6

Below we will simply assume that $d(x, y) = s$. The proof is easily extended to $> s$.

Effectively, we are considering the following process: we have a bag of n balls - corresponding to the nodes - of two colours (types), Changed and Unchanged, and we are picking $k \leq n/2$ of them uniformly at random without replacement. Let

$$\eta_1 := n(\text{Unchanged, Incorrect}) \quad (47)$$

$$\eta_2 := n(\text{Changed, Incorrect}) \quad (48)$$

and

$$\begin{aligned} \zeta &:= (n(\text{Unchanged}) - 2n(\text{Incorrect, Unchanged}))(n(\text{Changed}) - 2n(\text{Incorrect, Changed})) \\ &= (n - s - 2\eta_1)(s - 2\eta_2). \end{aligned} \quad (49)$$

We now condition on the number of errors being k , which imposes the condition that $\eta_1 + \eta_2 = k$. Recall the sampling without replacement distribution, which implies that

$$P(\eta_1 = k - j, \eta_2 = j \mid d(\hat{x}, x) = k) = \frac{\binom{n-s}{k-j} \binom{s}{j}}{\binom{n}{k}}. \quad (50)$$

Thus,

$$\begin{aligned} \mathbb{E}[\eta_1 \mid d(x, \hat{x}) = k] &= \frac{k}{n}(n-s) \\ \mathbb{E}[\eta_2 \mid d(x, \hat{x}) = k] &= \frac{k}{n}(s) \\ \mathbb{E}[\eta_1 \eta_2 \mid d(x, \hat{x}) = k] &= (n-s)(s) \frac{k(k-1)}{n(n-1)} = s(n-s) \left(\frac{k^2}{n^2} - \frac{k(n-k)}{n^2(n-1)} \right). \end{aligned}$$

As a consequence, we obtain that

$$\begin{aligned} \mathbb{E}[\zeta \mid d(x, \hat{x}) = k] &= s(n-s) \left(1 - 4\frac{k}{n} + 4\frac{k^2}{n^2} - 4\frac{k(n-k)}{n^2(n-1)} \right) \\ &= s(n-s) \left(\left(1 - 2\frac{k}{n} \right)^2 - 4\frac{k(n-k)}{n^2(n-1)} \right) \end{aligned} \quad (51)$$

Note that the above is decreasing as k increases for $k \leq n/2$.

Note further that the Markov chain $\zeta - d(\hat{x}, x) - G_1$ holds. Thus the above also holds for $\mathbb{E}[\zeta \mid \mathcal{E}(G_1), d(x, \hat{x}) = k]$.

We now condition on $\mathcal{E}(\hat{x})$ to find that

$$\frac{\mathbb{E}[\zeta \mid \mathcal{E}(\hat{x}), \mathcal{E}(G_1)]}{s(n-s)} \geq \left((1-2\varepsilon_{\max})^2 - 4 \frac{\varepsilon_{\max}(1-\varepsilon_{\max})}{n-1} \right) \quad (52)$$

$$\geq (1-2\varepsilon_{\max})^2 - \frac{1}{n-1} \quad (53)$$

where we have used $\varepsilon_{\max} \leq 1/2$, and the (unstated but obvious) condition that $n \geq 2$.

Applying the above to the result of Lemma 9, we find that

$$\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \mathcal{E}] \geq \left((1-2\varepsilon_{\max})^2 - \frac{1}{n-1} \right) \frac{(a-b)}{n} (n-s)s. \quad \square$$

B.2.6 Proof of Lemma 8

For continuity of exposition, we prove Lemma 8 before Lemma 7.

Below we will simply assume that $d(x, y) = s$. The proof is easily extended to the same being $> s$.

We begin with recalling Lemma 9, and noting that its proof does not utilise the exact balance assumption. We begin as in the proof of Lemma 6, by defining

$$\begin{aligned} \eta_1 &:= n(\text{Unchanged, Incorrect}) \\ \eta_2 &:= n(\text{Changed, Incorrect}) \end{aligned}$$

and noting from Lemma 9 that $\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) \mid \hat{x}] \geq \frac{(a-b)}{n} (n-s-2\eta_1)(s-2\eta_2) =: \frac{(a-b)}{n} \zeta$. Once again, let's fix the errors number of errors made by \hat{x} as some k , and let s_+ , s_- be the number of changes made in communities $+$ and $-$ respectively. Note that $s_+ + s_- = s$.

Using the above definitions,

$$\zeta = (n-s)(s-2\eta_2) - 2s\eta_1 + 4\eta_1\eta_2 \geq \frac{n}{2}(s-2\eta_2) - 2ks,$$

by noting that $\eta_1\eta_2 \geq 0$, $\eta_1 \leq k$ and that $s \leq n/2$. Thus,

$$\mathbb{E}[\zeta \mid d(\hat{x}, x) \leq k] \geq \frac{ns}{2} \left(1 - \frac{4k}{n} - 2 \frac{\mathbb{E}[\eta_2 \mid d(\hat{x}, x) \leq k]}{s} \right).$$

Suppose that the recovery procedure makes (k^+, k^-) -errors in communities $+$ and $-$ respectively, with $k^+ + k^- \leq k$. *Within-community* exchangability of nodes implies that the errors made within a community must be uniformly distributed over the community. Since the changes are made independently of these errors, we must have that the number of changed nodes in community $*$ $\in \{+1, -1\}$ that are incorrectly inferred must be $\text{Hyp}(n_*, s_*, k^*)$ distributed. In particular, this yields that

$$\mathbb{E}[\eta_2 \mid (k^+, k^-)] = \frac{s_+}{n_+} k^+ + \frac{s_-}{n_-} k^- \leq \frac{s_+ k^+ + s_- k^-}{\min(n_+, n_-)} \leq \frac{k^+ + k^-}{\min(n_+, n_-)} s.$$

The above immediately yields that

$$\mathbb{E}[\zeta \mid d(\hat{x}, x) \leq k] \geq \frac{ns}{2} \left(1 - \frac{4k}{n} - 2 \frac{k}{\min(n_+, n_-)} \right).$$

On \mathcal{E} , $k \leq \varepsilon n$, leading to the claimed bound. \square

B.2.7 Proof of Lemma 7

Recall the notation from Appendix B.2.2. Under the null $H \stackrel{\text{law}}{=} G'$. Below, we will use G' as a proxy for H in the null distribution, and use H only in the alternate.

To begin with, observe that both G', H are independent of G_1, \tilde{G}, \hat{x} , and that \tilde{G} is independent of \hat{x} given G_1 . Now, $T^{\hat{x}}$ is a signed sum of independent Bernoulli random variables with parameters

smaller than a/n given G_1 . Thus, invoking results from Ch. 2 of [CL06] (and using that for $a \geq C$ for some large enough C implies that $a \geq 16 \log(6n)/n \iff 1/6n \leq \exp(-na/16)$), we find that for $\Gamma \in \{\tilde{G}, G', H\}$,

$$P\left(|T^{\hat{x}}(\Gamma) - \mathbb{E}[T^{\hat{x}}(\Gamma) | G_1, \mathcal{E}]| \geq \sqrt{2na \log(6n)} | G_1, \mathcal{E}\right) \leq \frac{1}{3n},$$

where we have used that \mathcal{E} is determined given G_1 (i.e. \mathcal{E} lies in the sigma-algebra generated by G_1).

We now control the null and alternate fluctuations given \mathcal{E} .

Null: By the union bound, we find that

$$P\left(2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') \geq \mathbb{E}[2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') | G_1, \mathcal{E}] + 3\sqrt{2na \log(6n)} | G_1, \mathcal{E}\right) \leq \frac{2}{3n}$$

Recall from equation (40) from the proof of Lemma 5 that $\mathbb{E}[2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') | G_1, \mathcal{E}] \leq a^2$. Feeding this in, we find that

$$P\left(2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') \geq a^2 + 3\sqrt{2na \log(6n)} | G_1, \mathcal{E}\right) \leq \frac{2}{3n}.$$

The right hand side above does not depend on G_1 , and neither does the fluctuation radius within the probability. Thus integrating over $P(G_1 | \mathcal{E})$, we find that

$$P\left(T \geq a^2 + 3\sqrt{2na \log(6n)} | \mathcal{E}\right) \leq \frac{2}{3n},$$

where we have used that $T = 2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(H) \stackrel{\text{law}}{=} 2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G')$ under the null.

Alt: Following the above development again, this time with lower tails, we find that given G_1 with probability at least $1 - 2/3n$,

$$\begin{aligned} 2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') &\geq -(\mathbb{E}[2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') | G_1, \mathcal{E}]) - 3\sqrt{2na \log(6n)} \\ T^{\hat{x}}(G') - T^{\hat{x}}(H) &\geq +(\mathbb{E}[T^{\hat{x}}(G') - T^{\hat{x}}(H) | G_1, \mathcal{E}]) - 2\sqrt{2na \log(6n)} \end{aligned}$$

Further, given (G_1, \mathcal{E}) , by Lemmas 5, 6 we have

$$\begin{aligned} 2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(G') &\geq -a^2 - 3\sqrt{2na \log(6n)} \\ T^{\hat{x}}(G') - T^{\hat{x}}(H) &\geq +\kappa(a-b)s(1-s/n) - 2\sqrt{2na \log(6n)}, \end{aligned}$$

where $\kappa = (1 - 2\varepsilon_{\max})^2 - 1/(n-1)$. Adding the above, we find by the union bound that

$$P\left(2T^{\hat{x}}(\tilde{G}) - T^{\hat{x}}(H) \geq \kappa(a-b)s(1-s/n) - a^2 - 5\sqrt{2na \log(6n)} | G_1, \mathcal{E}\right) \geq 1 - \frac{4}{3n}.$$

The claim follows on noting that the right hand side and the fluctuation radius do not depend on G_1 , and integrating the inequality over G_1 . \square

B.3 Proof of the converse bound from Theorem 2.

We restate the lower bound below as a proposition:

Proposition *There exists a universal constant C , and another $c < 1$ that depends on C , such that if $\Lambda \leq C$ and $s \leq \frac{n}{2}(1-c)$, then reliable two-sample testing of balanced communities for s changes is impossible for large enough n .*

In particular, for $a+b < n/4$, the statement holds with $C = 1/8, c = 1/6$ for $n \geq 136$, and in this case, $R_{\text{TST}} \geq 0.25$.

Proof. The proof proceeds by using a variation of Le Cam's method, and importing impossibility results for the so-called distinguishability problem [Ban+16]. In particular, suppose that in the null distribution, the communities are drawn according to the uniform prior on balanced communities, denoted by π . Further, assume that if a s -change is made, then the resulting community is chosen uniformly from all communities that are at least s far from the null community. We have the hypothesis test:

$$H_0 : (G, H) \sim \sum_{x \in \mathcal{B}} \pi_x P(G|x) P(H|x) \quad \text{vs} \quad H_1 : (G, H) \sim \sum_{x, y \in \mathcal{B}} \pi_x \pi_{y|x} P(G|x) P(H|y),$$

where we use \mathcal{B} to denote the set of balanced communities, and $\pi_{y|x}$ is the uniform distribution on $\mathcal{B} \cap \{y : d(x, y) \geq s\}$. For succinctness, let us denote the null and alternate distributions above as p_{null} and p_{alt} respectively.

Once again, by Neyman-Pearson theory,

$$R_{\text{TST}} \geq R_\pi \geq 1 - d_{\text{TV}}(p_{\text{null}}, p_{\text{alt}}) \geq 1 - d_{\text{TV}}(p_{\text{null}}, Q) - d_{\text{TV}}(Q, p_{\text{alt}}),$$

where Q is any distribution, and the last inequality is since d_{TV} is metric.

We choose Q to be the unstructured distribution induced by an Erdős-Rényi graph of parameter $(a+b)/2n$. The primary reason for this is that explicit control on the total variation distance between p_{null} and Q is then available - for instance, by [WX18, §3.1.2], we have

$$D_{\chi^2}(p_{\text{null}} \| Q) + 1 \leq \mathbb{E} \left[\exp \left(\tau \left(\frac{4\mathcal{H} - n}{\sqrt{n}} \right)^2 \right) \right],$$

where \mathcal{H} is a Hypergeometric($n, n/2, n/2$) random variable, and

$$\tau = \frac{(a-b)^2}{2(a+b)} + \frac{(a-b)^2}{2(2-a/n-b/n)}.$$

Notice the extra factor of 2 compared to the expressions in [WX18], which arises since we sum over two independent graphs G, H and not one. We observe that

$$\tau = \Lambda \frac{n}{2n - a - b},$$

and explicitly, if $a + b \leq n/4$, then $\tau \leq \frac{4}{7}\Lambda$.

We now consider the alternate term. As a preliminary, let

$$\gamma := \frac{\sum_{k=0}^{s-1} \binom{n/2}{k/2}^2}{\binom{n}{n/2}}.$$

Note that γ is the probability that two balanced communities chosen independently and uniformly, lie within distortion s . Indeed, since communities are formed by identifying antipodal points in the boolean cube, the probability of picking a community at distortion $< s$ coincides with that of picking a balanced vector at Hamming distance $< s$ from a given balanced vector in the cube $\{0, 1\}^n$. The denominator in γ is clearly the number of balanced vectors in the cube, while the numerator is the number of balanced vectors at a distance of $< s$ from any given balanced vector - we choose $k < s$, and choose $k/2$ points marked 1 and $k/2$ marked 0, and flip them all.

As a consequence, we find that for any $x, y \in \mathcal{B}$,

$$\pi_{y|x} \leq \frac{\pi_y}{1 - \gamma}.$$

Thus, in the χ^2 expressions for p_{alt} , we have

$$\begin{aligned} \mathbb{E}_{(G, H) \sim Q^{\otimes 2}}[(p_{\text{alt}}/Q)^2] &= \sum_{x, y, x', y'} \mathbb{E} \left[\frac{P(G|x)P(G|x')}{Q^2(G)} \frac{P(H|y)P(H|y')}{Q^2(H)} \right] \pi_x \pi_{x'} \pi_{y|x} \pi_{y'|x'} \\ &\leq \frac{1}{(1-\gamma)^2} \sum_{x, y, x', y'} \mathbb{E} \left[\frac{P(G|x)P(G|x')}{Q^2(G)} \frac{P(H|y)P(H|y')}{Q^2(H)} \right] \pi_x \pi_{x'} \pi_y \pi_{y'} \\ &= \frac{1}{(1-\gamma)^2} \left(1 + \chi^2 \left(\sum_{x \in \mathcal{B}} \pi_x P(G|x) \| Q(G) \right) \right)^2 \end{aligned}$$

Since the final quantity is explicitly controlled in the cited section, we also have

$$1 + D_{\chi^2}(p_{\text{alt}} \| Q) \leq \frac{1}{(1-\gamma)^2} \mathbb{E} \left[\exp \left(\frac{\tau}{2} \left(\frac{4\mathcal{H} - n}{\sqrt{n}} \right)^2 \right) \right]^2 \leq \frac{1}{(1-\gamma)^2} \mathbb{E} \left[\exp \left(\tau \left(\frac{4\mathcal{H} - n}{\sqrt{n}} \right)^2 \right) \right],$$

the final relation arising from Jensen's inequality.

Since the quantity appears often, we let

$$\beta := \mathbb{E} \left[\exp \left(\tau \left(\frac{4\mathcal{H} - n}{\sqrt{n}} \right)^2 \right) \right].$$

Invoking the inequality $d_{\text{TV}} \leq \sqrt{\log(1 + D_{\chi^2})/2}$, we find that

$$R_{\text{TST}} \geq 1 - \sqrt{\log(\beta)/2} - \sqrt{\log(\beta(1-\gamma)^{-2})/2} = 1 - \sqrt{\log(\beta/(1-\gamma))}.$$

Note that the only s -dependent term in the above bounds is γ . We first offer control on the γ , and claim that for $s/n < 1/2$, $\gamma \rightarrow 0$. Indeed, since $s \leq n/2$, and by standard refinements of Stirling's approximation (for instance, we use [Gal68, Exercise 5.8] below),

$$\begin{aligned} \gamma &\leq s \frac{\binom{n/2}{s/2}}{\binom{n}{n/2}} \leq s \frac{1}{2\pi} \frac{n/2}{s/2(n-s)/2} 2^{nh_2(s/n)} \left(\sqrt{\frac{n}{8(n/2)^2}} 2^n \right)^{-1} \\ &\leq \sqrt{\frac{2n}{\pi^2}} 2^{-n(1-h_2(s/n))}, \end{aligned}$$

where h_2 is binary entropy in bits.

At this point the argument in the limit as $n \rightarrow \infty$ is complete - since $4(\mathcal{H} - n)/\sqrt{2n} \xrightarrow{\text{Law}} \mathcal{N}(0, 1)$, β is bounded as $n \rightarrow \infty$ by $\sqrt{1 - 2\tau}$ if $\tau < 1/2$, and since in this limit $\tau \rightarrow \Lambda/2$ (for $a, b = o(n)$), we obtain that if $\limsup s/n < 1/2$, and $\Lambda < 1$, then $\liminf R_{\text{TST}} > 0$.

Non-asymptotic bounds can be recovered by giving up space on the constants, leading to the statement we have claimed.

Concretely, to attain $R_{\text{TST}} > 1/4$, it suffice to show that $\beta(1-\gamma)^{-1} < e^{9/16}$. Now, for $s < n/3$, we have

$$(1-\gamma)e^{9/16} \geq \left(1 - \sqrt{2n/\pi^2} 2^{-0.08n}\right) e^{9/16} > 1.75$$

for $n \geq 136$.⁷ Thus, it suffices to control β to below 1.75 in this regime. To this end, note that $u \mapsto \exp(\tau((4u - n)/\sqrt{n})^2)$ is a continuous, convex map, and thus, by [Hoe63, Thm. 4],

$$\beta \leq \mathbb{E} \left[\exp \left(\tau \left(\frac{4\mathcal{B} - n}{\sqrt{n}} \right)^2 \right) \right],$$

where $\mathcal{B} \sim \text{Bin}(n/2, 1/2)$.

By Chernoff's bound, $P(|\mathcal{B} - n/4| \geq \sqrt{n}u) \leq 2e^{-4u^2}$, and thus, we have

$$\begin{aligned} \beta &\leq \int_0^\infty P \left(\exp \left(\tau \left(\frac{4\mathcal{B} - n}{\sqrt{n}} \right)^2 \right) \geq u \right) du \\ &\leq \int_0^\infty \min(1, 2u^{-1/4\tau}) d\tau \\ &= \frac{2^{4\tau}}{1 - 4\tau}, \end{aligned}$$

⁷This is calculated using a computer algebra system. Analytically it is still easy to argue something similar - for $n > 10$, $\sqrt{n}2^{-0.08n}$ is decreasing. Thus, for $n > 100$, $\sqrt{2n/\pi^2} 2^{-0.08n} < 10\sqrt{2}/256\pi < 1/50$, and thus the expression is at least $e^{9/16} \cdot 49/50 > 1.6 \cdot 49/50 > 1.5$. By following the next footnote with this number, this leads to an analytic proof of the conclusion holding for $\Lambda < 7/80 \approx 0.087$.

the final equality holding so long as $1/4\tau > 1 \iff \tau < 1/4$. The original claim follows if

$$\frac{2^{4\tau}}{1-4\tau} \leq \frac{7}{4},$$

which is true for $\tau < 0.074$. Since $\tau \leq 4/7\Lambda$, $\Lambda < 1/8$ implies that $\tau < 4/56 < 0.072$.⁸ \square

A couple of quick comments are useful here:

1. Note that the above cannot be applied usefully to GoF. This is because in GoF, the null is explicitly available, and we do not have the benefit of averaging with π in the TV expressions. This causes the equivalent term $\chi^2(P(G|x_0)||Q(G))$ to grow exponentially with $n\Lambda$.
2. The above characterises the tightness of our claimed bounds for TST of large changes - the method works if $\Lambda = \Omega(1)$ and $s \gg \sqrt{n \log n}$, and by the above argument, no test can work if $\Lambda \ll 1$, as long as the change is not extreme ($\limsup s/n < 1/2$).
3. While the above approach is wasteful in how it utilises s , this is actually a non-issue, since the bounds require a separate control on $d_{TV}(p_{\text{null}}||Q)$, which can only be controlled if $\Lambda = O(1)$. In particular, we cannot pull out better bounds for the small s situation from the above.

C Experimental Details

C.1 Experiments on SBMs

The experimnts simulate an ensemble of GoF and TST test and evaluate the performance of the two schemes using the sum of false alarm and missed detection probabilities ($FA + MD$).

While the GoF scheme is implemented precisely as in the main text, the experiments use a slightly modified version of Algorithm 1 for the TST:

- (i) G_1 subsamples G at a rate η , and the test statistic T is appropriately modified: $T := \frac{1}{1-\eta} T^{\hat{x}_1}(\tilde{G}) - T^{\hat{x}_1}(H)$. Intrinsically, the spectral clustering step is the more singal-sensitive part of the scheme 1. While splitting the graphs equally is fine for theoretical results, it is better in practice to devote more SNR to the clustering step, and less to compute the test statistic, which can be done by increasing η . In the following, we set $\eta = 0.85$. Other values of η are explored in Appendix C.1.2.
- (ii) The constant factor in the threshold developed in the test is conservative, and we vary it to adjust for different values of η and to mitigate its suboptimality. In the experiments, we used the threshold $\frac{3}{4} \sqrt{n(a+b) \log(6n)}$.

As noted in the main text, the experiments are performed for various (s, Λ) for a fixed value of $a/b = 3$. Λ is varied between Λ_0 and $10\Lambda_0$ for $\Lambda_0 = 3/4 \log(n/100) \approx 1.7$. This is significantly below the theoretical threshold of 2 necessary for non-trivial recovery. Further, $8\Lambda_0 = 2 \log(n)$, at which point recovery with constant order distortion becomes viable.

C.1.1 Implementation details

The experiment is setup as follows:

1. We fix a value of $\Lambda_0 = 3/4 \log(n/100)$ as above. Then, for some choice of $b/a = r$, we choose (a, b) satisfying $\Lambda = \alpha\Lambda_0$ and $\alpha \in [1, 10]$. r is set to be $1/3$.
2. For a fixed number of nodes, n , and for $s \in [1 : n/2]$, we consider the balanced partition $x = [x_i]_{i=1}^n$ with

$$x_i = \begin{cases} 0, & 0 \leq i \leq n/2 \\ 1, & n/2 < i \leq n \end{cases}$$

⁸The number 0.074 is calculated using a computer algebra system. Purely analytic calculations are straightforward as well - for example by using $2^{4\tau} \leq 1 + 4\tau$ for $\tau < 1/4$, which can be proved by noting that $1 + 4\tau - 2^{4\tau}$ is initially increasing, and then strictly decreasing after a point, and that $1/4$ is a root of this function. This implies that the conclusion holds so long as $\tau < 3/44$, which holds if $\Lambda < 21/176 \approx 0.119$.

for the null distribution, and the shifted balanced partition $y = [y_i]_{i=1}^n$ with

$$y_i = \begin{cases} 0, & s/2 < i \leq n/2 + s/2 \\ 1, & i \in (n/2, n] \cup [0, s/2] \end{cases}$$

for the alternate distribution. This ensures that $d(x, y) = s$. We take $\lfloor \cdot \rfloor$ whenever s or n are odd.

3. We sample $G, G' \sim P(\cdot | x)$ and $H \sim P(\cdot | y)$, where P represents drawing from an SBM with parameters n, a and b , as described in §1

GoF procedure. Recall that we are given a proposed partition x_0 . Here we set $x_0 = x$. The results of running the tests on the graph G then serve to characterise size, and those on H serve to characterise power.

1. For the naïve scheme, we produce partitions \hat{x} and \hat{y} from G and H respectively via spectral clustering (see below for details), and declare for null in either case if $d(x_0, z) < s/2$, where z is respectively \hat{x} and \hat{y} .
2. For the alternate scheme, we instead compute the statistic from §2, and reject on the basis of the threshold developed there.

TST procedure. Similarly to the above, runs on the pair (G, G') serve to characterise size, and on (G, H) serve to characterise power of the test. Precisely:

1. For the naïve two-sample test based on recovery and comparison, we estimate \hat{x}, \hat{x}' and \hat{y} from G, G' and H respectively. The structure is estimated using spectral clustering (see below for implementation details). We declare that a change has occurred if $d(\hat{x}, \hat{x}') \geq s/2$, and no change if $d(\hat{x}, \hat{y}) < s/2$. We get a false alarm every time we declare a change on the pair (G, G') , and we miss a detection whenever we declare no change on the pair (G, H) . The false alarm and missed detection probabilities are estimated as an average over $M = 100$ samples.
2. For the two-sample test based on Algorithm 1, we follow the algorithm as stated, making only the modifications previously described. To be precise, we estimate \hat{x}_1 from G_1 , a subsampling of (the edges of) G at rate η . Then, we compute the test statistics in the null and alternate distributions:

$$T_{\text{Null}} = \frac{1}{1 - \eta} T^{\hat{x}_1}(\tilde{G}) - T^{\hat{x}_1}(G')$$

and

$$T_{\text{Alt.}} = \frac{1}{1 - \eta} T^{\hat{x}_1}(\tilde{G}) - T^{\hat{x}_1}(H),$$

where $\tilde{G} = G - G_1$.

In both the above cases, the simulations are performed over a range of $\Lambda = \alpha \Lambda_0$ and s , where $\alpha \in [1, 10]$ and $s \in (0, 250)$. Performance is indicated using the sum of false alarm and missed detection rates.

Details associated with the implementation of the aforementioned schemes are given below:

1. All experiments were implemented in the Python programming language (v3.5+), using the Numpy (v1.12+) and Scipy (v0.18+) scientific computing packages [Oli06; JOP01].
2. Structure learning was performed using the Spectral Clustering [Lux07] algorithm, as implemented by the Scikit-learn package (v0.19.1+) [Ped+11].
3. Spectral Clustering was regularized in the manner suggested by [JY16]. Effectively, if G was the adjacency matrix to be submitted to the Scikit-learn spectral clustering function, we performed pre-addition, and instead passed $G + \tau \mathbf{1}\mathbf{1}^\top$. We set $\tau = \frac{1}{10n}$, which proved sufficient to run the spectral clustering function with no errors or warnings.
4. All plots were generated using Matplotlib (v2.1+) [Hun07].

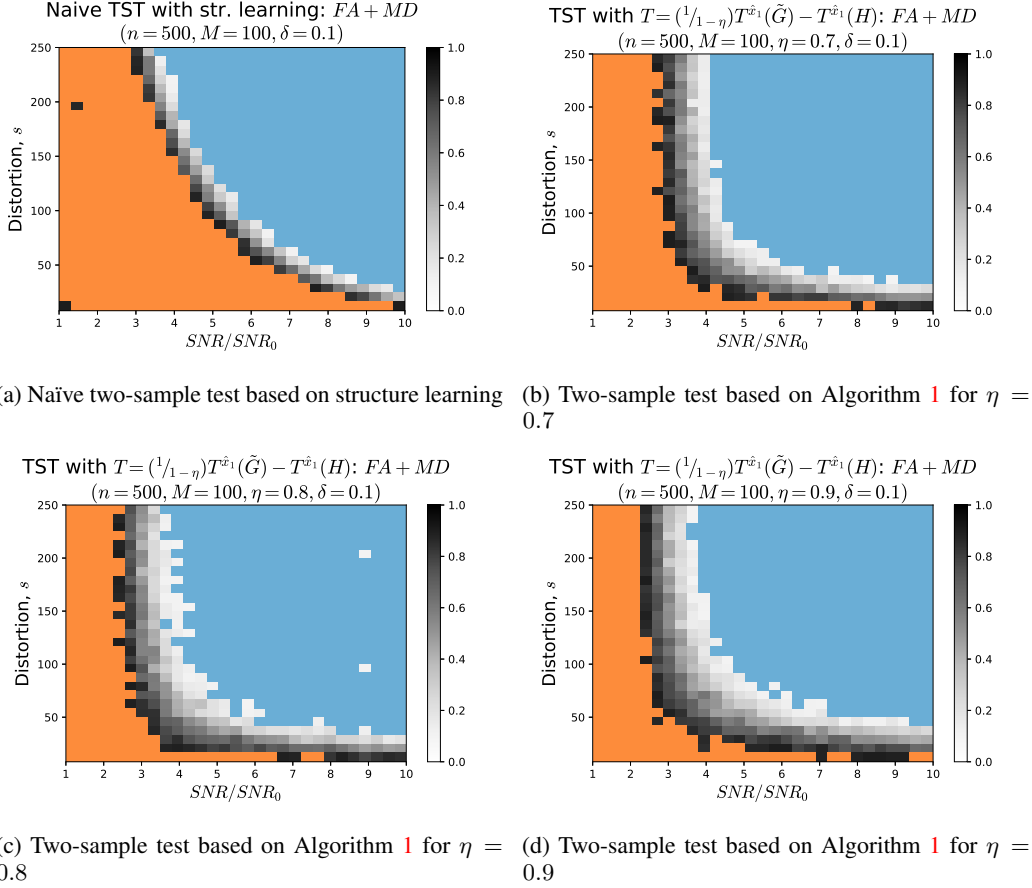


Figure 5: A comparison between the naïve two-sample test based on structure learning, and the two-sample test we propose in Algorithm 1, for $\eta \in \{0.7, 0.8, 0.9\}$. Error rates lower than $\delta = 0.1$ have been shaded blue to represent “success”, while those higher than $1 - \delta = 0.9$ have been shaded orange to represent “failure”.

C.1.2 Modifications to η

For completeness, we demonstrate how the performance of the modified two-sample test based on Algorithm 1 varies as η is changed. Figure 5 compares the naïve two-sample test against the scheme based on Algorithm 1, for three different values of η : 0.7, 0.8 and 0.9.

We use the following parameters: $n = 500$, $\text{SNR}_0 = \frac{3}{8} \log(n/100) = \frac{3}{8} \log 5 \approx 0.5$, $\frac{b}{a} = r = 1/3$. For $\eta = 0.7$ and $\eta = 0.8$, the threshold used is $\sqrt{n(a+b) \log(6n)}$, while for $\eta = 0.9$, we used a higher threshold of $\frac{3}{2} \sqrt{n(a+b) \log(6n)}$.

While differences are rather subtle, a careful examination may reveal that as η increases, the failure region recedes, while the success region advances in the high- s , low-SNR regime. However, the cost of this is an increased threshold to maintain success at $\delta = 0.1$, and a wider transition region, indicating that different η might be optimal at different n .

C.2 Experiments on the Political Blogs dataset

While the original graph has 1490 nodes, we followed standard practice in selecting the largest (weakly) connected component of the graph, which contains 1222 nodes. We denote this graph as G . The true partition of the blogs according to political leaning is available, denoted x_{True} here. This also allows accurate estimates of the graph parameters (a, b) to be made, and we use these estimates

for a, b for GoF, and for the semi-synthetic procedure for TST. We found that $\hat{a} \approx 49.5$, while $\hat{b} \approx 5.2$, giving a ratio $a/b \approx 10$. The communities, according to x_{True} are of sizes 636 and 586.

The regime of low Λ is explored via sparsification. Fixing a $\rho \in (0, 1]$, sparsification is performed by independently flipping coins for each edge in G , and keeping the edge with probability ρ . We refer to ρ as the rate of sparsification.

We lastly note that at no sparsification ($\rho = 1$), spectral clustering produces a partition \hat{x}_1 such that $d(x_{\text{True}}, \hat{x}_1) = 56$.

GoF Procedure.

1. The graph is sparsified at rate ρ . Let the sparsened graph be G_ρ .
2. For the naïve recovery based scheme, spectral clustering is performed on G_ρ as in the previous section to generate \hat{x}_ρ .
3. For the proposed test from §2, the statistic is computed on G_ρ .
4. The size of the test is estimated by running the GoF tests with $x_0 = x_{\text{True}}$. For the naïve scheme, we reject if $d(\hat{x}_\rho, x_0) \geq s/2$; for the proposed scheme, we use the test from §2.
5. To compute the power at distortion s , we first generate y by randomly inverting the community labels of s nodes in x_{True} . We then run the same procedure as in the previous line, but with $x_0 = y$. Note that the graphs are not edited in any way.
6. The precise implementation details are exactly as in Appendix C.1.1, with the minor difference that we use a regularizer of $\tau = 1$ for spectral clustering.

TST Procedure.

1. Recall that TST requires two graphs as input. The experiment compares the political blogs graph against SBMs.
2. To compute the size, we require a graph with the same underlying communities as G . Thus we generate G' , which is drawn as an SBM with the underlying partition x_{True} , and parameters a, b as estimated from the political blogs graph G .
3. To determine the power of the tests we need a graph with an s -far underlying community. For this, we first generate a y such that $d(x_{\text{True}}, y) = s$, as we did in the GoF Procedure. Next, we sample H as an SBM with underlying partition y .
4. The graphs G, G' and H are now all sparsified at rate ρ to get G_ρ, G'_ρ and H_ρ .
5. The size of each test is estimated using the TST procedures, as described in Appendix C.1.1 on the pair (G_ρ, G'_ρ) . Power is similarly estimated using the TST procedures on the pair (G_ρ, H_ρ) .

C.3 Experiments on the GMRFs

Following the heuristic detailed in §4.3, we naïvely generalise community recovery and testing to this setting, by replacing all instances of the graph adjacency matrix in previous settings with the sample covariance matrix.

The Gaussian Markov Random Field is described by its precision matrix Θ (i.e., the inverse covariance matrix of the Gaussian random vector on its nodes). We perform a preliminary examination of the possibility of testing changes in communities for an SBM-structured GMRF even when learning the structure is hard or impossible. As described in Section 4.3, we set

$$\Theta = I + \gamma G,$$

where G is the adjacency matrix of an SBM with known parameters. We generate samples from the GMRF as follows:

1. For a fixed number of nodes n , we fix an SNR for the SBM, Λ , and compute (a, b) satisfying this Λ so that $b/a = r$.

2. Here, we consider $n = 1000$ nodes and take $\Lambda = 30\Lambda_0$, where $\Lambda_0 = \frac{10}{11} \log(n/100) \approx 2.1$, as before, and $r = 1/10$. We find that $(a, b) \approx (12.34 \log n, 1.234 \log n)$. Note that since $\Lambda \approx 63 \approx 10 \log(n)$, recovery of the communities for a raw SBM at this SNR is trivial.
3. We fix a GMRF parameter γ . Here, we take $\gamma = 3/(a + b) \approx 0.032$.
4. We can now construct the precision matrix Θ after sampling G from the SBM. We re-sample to ensure that Θ is positive-definite, but in practice, for the value of γ quoted above, we did not encounter the need to re-sample.
5. To generate i.i.d. samples $\zeta \sim \mathcal{N}(0, \Theta^{-1})$ in a stable manner, we use the following algorithm:
 - (a) Compute the lower-triangular Cholesky factor R of Θ , so that $\Theta = RR^\top$.
 - (b) Sample $\xi \sim \mathcal{N}(0, I)$ from a standard n -dimensional multivariate normal distribution.
 - (c) Solve for ζ in $R^\top \zeta = \xi$.
 This suffices, since, $\zeta = (R^\top)^{-1} \xi$ would then have the covariance matrix $(R^\top)^{-1} R^{-1} = (RR^\top)^{-1} = \Theta^{-1}$.
6. In this manner, we generate samples from the null and alternate distributions: let ζ , ζ' and v respectively denote samples drawn from a GMRF structured using G , G' and H respectively. Here, G , G' and H exactly are as described in Section C.1.1.

Next, we describe how each of the two schemes is evaluated:

1. Assuming we have t i.i.d. samples of ζ , generated as described above, we estimate the covariance matrix $\hat{\Sigma}$ of ζ using the standard estimator:

$$\hat{\Sigma} = \frac{1}{t-1} \sum_{i=1}^t (\zeta_i - \bar{\zeta})(\zeta_i - \bar{\zeta})^\top,$$

where $\bar{\zeta} = \frac{1}{t} \sum_{i=1}^t \zeta_i$. We then compute the correlation matrix,

$$\hat{C} : \hat{C}_{ij} = \frac{\hat{\Sigma}_{ij}}{\sqrt{\hat{\Sigma}_{ii} \hat{\Sigma}_{jj}}},$$

which will be used in place of the adjacency matrix for both two-sample testing schemes.

2. Similarly, we compute \hat{C} , \hat{C}' and \hat{D} from ζ , ζ' and v respectively.
3. The naïve two-sample test based on recovery and comparison is evaluated exactly as described in Section C.1.1, except that \hat{C} , \hat{C}' and \hat{D} are used in place of G , G' and H respectively. False alarm and missed detection rates are also computed in exactly the same way.
4. The two-sample test based on Algorithm 1 has several important variations:
 - (a) We use the test statistics

$$\begin{aligned} T_{\text{Null}} &= T^{\hat{x}}(\hat{C}) - T^{\hat{x}}(\hat{C}') \\ T_{\text{Alt.}} &= T^{\hat{x}}(\hat{C}) - T^{\hat{x}}(\hat{D}), \end{aligned}$$

for the null and alternate distributions respectively. Here, \hat{x} has been estimated from \hat{C} .

- (b) The threshold for the test is estimated from data. That is, we simulate $M = 100$ samples of T_{Null} and $T_{\text{Alt.}}$ each, and fit a classifier to differentiate between the two distributions. The classifier used is a simplistic 1-dimensional Linear Discriminant Analysis.
- (c) We estimate false alarm and missed detection rates by applying the classifier to a hold-out dataset. To use the data as efficiently as possible, we use 10-fold repeated, stratified cross-validation, with 10 repetitions.

Remark on subsampling.

1. Note that in the two-sample test for GMRFs based on Algorithm 1, we do not subsample \hat{C} as we did before in the case of SBMs.
2. While previously, we had subsampled G to create two subgraphs G_1 and \tilde{G} that shared independence properties for ease of theoretical analysis, it should be noted that subsampling results in an effective loss of SNR. This is also the reason why we had to adjust the implementation using a different rate η .
3. However, it emerges empirically that skipping the subsampling entirely, with a completely dependent \hat{x} and G , makes for better separation between the null and alternate distributions, providing a more powerful statistic.
4. Since we could not analytically derive a threshold for this statistic, we presented the subsampled test statistic for the first experiment.
5. Since in the case of GMRFs, we are estimating the threshold from data, we use the more powerful test statistic to show the full extent of possible gains when using a dedicated algorithm for change detection, instead of naïvely looking for changes by learning community structures first.