Query-Limited Community Recovery in Stochastic Block Models*

Sabyasachi Basu
Microsoft Research, India
firstname.lastname@microsoft.com

Manuj Mukherjee
Indraprastha Institute of Information Technology, India
firstname@iiitd.ac.in

Lutz Oettershagen
University of Liverpool, United Kingdom
firstname.lastname@liverpool.ac.uk

Suhas Thejaswi
Aalto University, Finland
firstname.lastname@aalto.fi

Abstract

We study exact community recovery in the two-community stochastic block model on n vertices under limited and noisy access to network data. The learner may query a noisy neighborhood oracle that reveals each true neighbor of a queried vertex independently with fixed probability and never returns non-neighbors, subject to a finite query budget. We consider both oracle-only access and a combined model where the learner also observes a single subsampled copy of the underlying graph. For oracle-only access, balanced uniform querying gives a sharp non-adaptive benchmark: when each vertex is queried the same number of times, the observations reduce to an SBM with attenuated edge probabilities and the Abbe-Bandeira-Hall exact recovery threshold applies. We show that this benchmark is not adaptively optimal: a two-stage adaptive strategy succeeds with $n + o(n)$ queries in a regime where balanced uniform querying requires mn queries for some $m > 1$. With an additional subsampled graph, we prove a sublinear-query adaptivity gap: balanced data-independent uniform querying with a sublinear budget does not improve over the subsampled graph alone, whereas adaptive querying can target a small set of uncertain vertices and achieve exact recovery. Thus adaptive data acquisition can strictly improve the information-theoretic limits of exact recovery.

1 Introduction

Recovering latent community structure from network data is a fundamental problem in machine learning, statistics, and network science, with applications ranging from social and communication networks to biological and information systems [23]. The stochastic block model (SBM) is a canonical random-graph model for this task: in its simplest two-community form each vertex (among n vertices) belongs to one of two communities, an edge between two vertices is present with probability p if they are in the same community and with probability q otherwise. This model, introduced by Holland, Laskey, and Leinhardt [22], and widely studied since, has become a standard framework for understanding the statistical and algorithmic limits of community detection. A central task here is *exact recovery*: reconstructing all community labels, up to a global label flip, from the observed graph. In the sparse logarithmic-degree regime, where $p = \alpha\frac{\log n}{n}$ and $q = \beta\frac{\log n}{n}$, exact recovery exhibits a sharp phase transition governed by the quantity $(\sqrt{\alpha} - \sqrt{\beta})^2$: in particular, recovery is information-theoretically possible precisely above the classical Abbe–Bandeira–Hall threshold [4].

*Authors are listed in alphabetical order.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

In many applications, however, the full graph is not directly observable. Edges may be missing because of subsampling, privacy restrictions, incomplete crawls, or noisy measurements, and additional information may be available only through local probes such as API queries, experiments, or surveys [21, 28, 35]. Network sampling can also distort observed structural properties [37, 38]. Motivated by such settings, prior work has investigated how additional imperfect observations—such as multiple correlated or subsampled graphs—can change the information-theoretic limits of exact recovery [20, 32]. In these models, the observation mechanism is exogenous: it is specified ex ante by a fixed probabilistic model, independent of the procedure for exact recovery [1, 20, 32, 33]. In this paper, we study a complementary source of information. Concretely, we introduce a query-limited model with access to a noisy neighborhood oracle: when queried at a vertex, the oracle reveals each true neighbor independently with probability w and never returns non-neighbors. Subject to a total budget of k queries, the learner must decide which vertices to query and how to allocate its queries. This model is closely related to edge subsampling, since both mechanisms introduce one-sided observation noise. But the key difference is operational: in a subsampled graph, the observation pattern is fixed independently of the learner, whereas in the oracle model the learner chooses where information is acquired, possibly non-uniformly and adaptively. As a consequence, exact recovery depends not only on how much information is available, but also on how it is acquired.

This leads to a natural question: when can adaptive oracle querying improve exact recovery over uniform querying, and what initial information is required for such an improvement? We investigate this question in two models. In the *oracle-only* model, the learner has no initial graph observation and only has access to a noisy-neighborhood oracle. In the *subsampled-graph-plus-oracle* model, the learner first observes a subsampled graph G_{sub} and may then use oracle queries to refine the recovery[2]. This two-part formulation separates the case in which adaptivity can only improve constants from the case in which even coarse side information enables a sublinear-query adaptivity gap. This setting also captures incremental network acquisition, where an incomplete graph is available and additional local neighborhood information is acquired selectively by probing vertices [36].

Our work connects exact recovery to adaptive data acquisition for network inference. When local measurements are expensive, the learner must decide not only how to infer communities from partial observations, but also how to allocate a limited budget to enable exact recovery. Such measurements may come from API calls, experiments, surveys, or other procedures that reveal incomplete local neighborhood information. In this case, one can start with uniform querying as a natural baseline, but it need not be the most efficient strategy once a coarse global estimate can identify vertices that remain *difficult* to classify. Our results identify regimes in which adaptive querying can selectively target these "difficult" vertices and thereby improve the information-theoretic limits of exact recovery. In this sense, our work extends beyond exogenous observation settings studied in the literature on exact recovery in the SBM with correlated side information, i.e., where the observation mechanism is fixed in advance, by making (adaptive) data acquisition itself part of the inference problem.

Our contributions are as follows:[3]

(i) **Oracle-only access: a linear-regime adaptivity gap.** We first analyze the oracle-only setting. For balanced uniform querying, where each vertex is queried the same number of times, the observed graph is again an SBM with attenuated (i.e., reduced) edge probabilities. This gives a sharp threshold for uniform sampling via the Abbe-Bandeira-Hall criterion, with effective edge-retention factor $\theta_{\mathrm{unif}} = 1 - (1 - w)^{2k/n}$. We then show that uniform querying is not necessarily optimal. For an explicit choice of parameters, a two-stage adaptive strategy first constructs a rough global sketch, identifies a sublinear set of uncertain vertices, and then spends additional queries only on that set. This strategy succeeds with $n + o(n)$ queries, whereas balanced uniform querying requires $x_\star n$ queries for some $x_\star > 1$. We also show that every successful oracle-only strategy requires $\Omega(n)$ queries, so the adaptive construction matches the optimal query order while improving the uniform constant.

(ii) **Subsampled graph plus oracle access: a sublinear adaptivity gap.** We next consider the setting in which the learner is also given a subsampled copy of the graph. In this case, the subsampled graph provides an initial global sketch before any oracle queries are made. We show that, with a sublinear budget $k = o(n)$, uniform oracle querying does not improve the exact recovery threshold beyond that of the subsampled graph alone. In contrast, a two-stage adaptive strategy can use G_{sub} to

[2] Unlike models with multiple passively observed correlated graphs [20, 32], our setting combines a single subsampled graph with adaptive neighborhood queries.

[3] Omitted details and full proofs are available in the appendix.

identify a small set of uncertain vertices and spend the sublinear query budget exclusively on them. For suitable parameters, this achieves exact recovery in a regime where uniform querying with the same budget fails, yielding a genuine adaptivity gap.

It is worth highlighting that our results establish different kinds of adaptivity gains in the two models. In the oracle-only setting, exact recovery requires $\Omega(n)$ queries even for adaptive strategies; our improvement therefore concerns the leading constant of a linear query budget. With a supplied subsampled graph, the learner already has global information before making any oracle queries. In this setting, we identify a parameter regime in which adaptive querying achieves exact recovery with $o(n)$ additional queries, whereas data-independent uniform querying with the same budget cannot. Thus, the second result establishes a separation between the possible and the impossible at sublinear budgets, rather than merely a constant-factor improvement.

Table 1: Summary of passive baselines and our query-limited results. The *Acquisition* column distinguishes externally generated observations without learner control (Exogenous), learner-controlled but data-independent balanced querying (Uniform), and observation-dependent querying (Adaptive). "Any" denotes a lower bound over all querying strategies, including adaptive ones. Here, $\Delta^2 = (\sqrt{\alpha} - \sqrt{\beta})^2$, x is the number of queries per vertex, w is the oracle reliability, and $\theta(x, w) = 1 - (1 - w)^{2x}$ is the induced edge-retention probability.

Setting	**Acquisition**	**#Queries**	**Exact recovery condition**	**Ref./Result**
Standard SBM	Exogenous	–	$\Delta^2 > 2$	[4]
Correlated graph(s)	Exogenous	–	Modified threshold	[20, 32, 33]
Oracle only	Uniform	xn	$\theta(x, w)\,\Delta^2 > 2$	Prop. 1
Oracle only	Adaptive	$n + o(n)$	Below uniform threshold in Ex. 1	Thm. 3 (Ex.1)
Oracle only	Any	$o(n)$	Impossible	Thm. 5
Subsampled graph + oracle	Uniform	$o(n)$	Fails in the gap regime	Thm. 8
Subsampled graph + oracle	Adaptive	$o(n)$	Succeeds in the gap regime	Thm. 8

Table 1 summarizes the passive baselines and the query-limited results proved in this paper.

Technical challenges. The main technical difficulty is that adaptive querying breaks the homogeneous observation structure used in standard exact recovery analyses for the SBM. In the adaptive setting, the probability that an edge is observed may depend on non-uniform, random, and transcript-dependent query counts. Moreover, targeted querying creates a selection effect: a vertex is selected for additional queries precisely because the observations already available around that vertex look ambiguous. Those same observations therefore cannot be reused as unbiased evidence in the final local test.

Leave-one-out-screening. We overcome these challenges through a leave-one-out screening framework, which plays a fundamentally different role from similar arguments used in prior SBM and random-matrix analyses [5, 14, 19, 42]. In those works, leave-one-out is primarily a device for controlling dependencies in passive observation models. In our setting, it is the mechanism that makes adaptive data acquisition analyzable. For each vertex v, we compute a reference labeling $\hat{\sigma}^{(-v)}$ using observations that exclude edges incident to v, use one comparison set to decide whether v should be selected for refinement, and reserve a disjoint comparison set for the final likelihood test with fresh oracle observations. A key step in our analysis is to prove that the resulting screened set is sublinear, so that only a vanishing fraction of vertices require refinement. This separation decouples adaptive selection from local refinement and makes it possible to concentrate the oracle budget on the sublinear vertex subset responsible for the remaining recovery errors and turns leave-one-out from a technical dependence-control tool into a structural ingredient of the adaptive recovery procedure.

Related Work. SBMs are a central model for studying community recovery, following the formulation of Holland et al. [22] and the exact recovery theory developed in subsequent work [1, 4, 6, 13, 26]. The SBM literature spans sharp recovery thresholds, optimal algorithms, and numerous model extensions, including general multi-community limits, spectral and belief-propagation methods, degree-corrected and hierarchical models, attributed graphs, and active querying, e.g., [3, 8, 12, 17, 25, 29, 31]. We focus here on the lines of work most closely related to partial observation and adaptive data acquisition. Recovery from incomplete network data has been studied in settings with sampled vertices, missing edges, and non-random missingness [21, 35]. These models motivate partial observability, but the sampling mechanism is typically fixed independently of the

learner. Our setting differs by making observation itself a decision: the learner allocates a finite budget of noisy neighborhood queries.

Adaptive information acquisition for community detection has been studied through random and adaptive edge sampling [41] and active label queries [16]. Related semi-supervised SBM models investigate recovery with revealed labels or side information [24, 34]. Unlike these models, our oracle reveals noisy neighborhood information rather than community labels or individual edge observations. Related active-learning approaches also emphasize query design [7, 27, 39]. Two-stage recovery and local refinement are established techniques in SBM inference [2, 18]. In particular, Abbe and Sandon [2] use graph splitting to preserve independence during recovery. Our contribution differs by using preliminary observations to adaptively select vertices for additional queries, while separating screening and refinement observations to control selection bias. Our model is also related to correlated SBMs and multiple graph observations, where additional correlated networks can alter exact recovery thresholds [20, 32, 33]. Recent computational work further shows that correlated observations introduce new algorithmic phenomena [10, 11, 15]. In these settings, however, the observations are fixed in advance; in ours, the learner chooses where to acquire additional local information. Other SBM extensions, including bipartite, attributed, multiplex, and testing problems, show how richer observation structures affect inference [30, 40]. None address the query-limited noisy-oracle setting studied here. Our contribution is to distinguish regimes where adaptive querying gives no information-theoretic benefit from regimes where coarse side information enables a genuine adaptivity gap.

2 Preliminaries

We consider a stochastic block model (SBM) on n vertices with two communities. Throughout the paper, we assume that n is even, as is standard practice in the SBM literature (see *e.g.*, [4, 29]). A random graph $G = (V, E)$ with $|V| = n$ is said to be an SBM(n, p, q), for some $p, q \in [0, 1]$, if there exists a partition $V = V^+ \cup V^-$, with $|V^+| = |V^-| = n/2$, such that, conditional on the partition, edges are generated independently and (i, j) is present with probability p if i and j belong to the same community, and with probability q if i and j belong to different communities. The partition is generated uniformly at random among all balanced bipartitions of V. We denote the ground-truth community labeling by $\sigma^* \in \{\pm 1\}^n$, where $\sigma_i^* = +1$ if $i \in V^+$ and $\sigma_i^* = -1$ if $i \in V^-$. Thus $\sum_{i \in V} \sigma_i^* = 0$. Exact recovery is always understood up to a global sign flip.

Let A denote the adjacency matrix of G, where $A_{ij} = 1$ ($A_{ij} = 0$ resp.) indicate that $(i, j) \in E$ ($(i, j) \notin E$ resp.). Following the notation, we say $G \sim \text{SBM}(n, p, q)$ if, for all $i \neq j$, $\Pr(A_{ij} = 1 | \sigma_i^* = \sigma_j^*) = p$ and $\Pr(A_{ij} = 1 | \sigma_i^* \neq \sigma_j^*) = q$. For a positive integer $k \in \mathbb{N}$, we denote $[k] := \{1, \ldots, k\}$, and for convenience, we consider the vertex set to be $V = [n]$. For $\alpha, \beta \geq 0$, we use $\Delta^2(\alpha, \beta) := \left(\sqrt{\alpha} - \sqrt{\beta}\right)^2$, and $D_+(\alpha, \beta) := \frac{1}{2}\Delta^2(\alpha, \beta)$. The ABH exact recovery threshold for the balanced two-community SBM with edge probabilities $\alpha \frac{\log n}{n}$ and $\beta \frac{\log n}{n}$ is $D_+(\alpha, \beta) > 1$, or equivalently $\Delta^2(\alpha, \beta) > 2$. When the ambient SBM parameters are α, β, we use $\Delta^2 := \Delta^2(\alpha, \beta)$.

A community recovery algorithm takes a graph G as input and outputs an estimate $\tilde{\sigma} \in \{\pm 1\}^n$ of the community labels, without the knowledge of the ground truth σ^*. The success of the algorithm is measured by the *overlap* between $\tilde{\sigma}$ and σ^*. And it is the normalized difference between the number of correctly and incorrectly labeled vertices. Formally, the overlap is expressed as

$$\text{ov}(\sigma^*, \tilde{\sigma}) := \frac{1}{n} \left| \sum_{i \in V} \sigma_i^* \, \tilde{\sigma}_i \right|. \tag{1}$$

Note that, $\text{ov}(\sigma^*, \tilde{\sigma}) \in [0, 1]$ and larger values of $\text{ov}(\sigma^*, \tilde{\sigma})$ indicate better estimate. The *exact recovery* occurs when $\tilde{\sigma} = \sigma^*$ which is equivalent to $\text{ov}(\sigma^*, \tilde{\sigma}) = 1$.

Observation models: We study exact recovery under two related observation models. Both are based on the same underlying graph $G \sim \text{SBM}(n, \alpha \frac{\log n}{n}, \beta \frac{\log n}{n})$, but they differ in what information is available before the learner starts querying the oracle.

Model 1 (Oracle-only access). Consider a two-community stochastic block model $\text{SBM}(n, \alpha \frac{\log n}{n}, \beta \frac{\log n}{n})$. The learner does not observe the graph directly, but may make at most k queries to a noisy neighborhood oracle. When queried at a vertex, the oracle reveals each true neighbor independently with probability w and never reports non-neighbors. The learner may choose

queries non-adaptively or adaptively as a function of previous oracle responses. The goal is exact recovery of the community labels from the oracle transcript.

Model 2 (Subsampled graph plus oracle access). Consider the same underlying SBM, but suppose the learner first observes a subsampled graph G_{sub}, obtained by retaining each edge independently with probability $s \in [0,1]$. After observing G_{sub}, the learner may make at most k queries to the noisy neighborhood oracle that reveals each true neighbor from the underlying graph G with probability w. The goal is exact recovery from the combined observation $(G_{\mathrm{sub}}, \text{oracle transcript})$.

The main question is how the query allocation affects the exact recovery threshold. We compare balanced uniform querying, in which the budget is spread evenly over all vertices, with adaptive or non-uniform querying, in which the allocation may depend on the information observed so far.

3 Community Recovery with Oracle Only Information

In this section, we investigate Model 1, where the learner has no initial graph observation and must rely entirely on the noisy neighborhood oracle. In Section 3.1, we start with balanced uniform querying, which provides the natural (non-adaptive) benchmark and yields a sharp exact recovery threshold. In Section 3.2, we turn to strategies where the learner may choose to adapt the querying strategy based on the information revealed during the querying process. This comparison isolates the role of adaptivity in the oracle-only model. Our results show that a linear query budget is necessary for exact recovery even under adaptive querying. At the same time, adaptivity can still be strictly beneficial at the level of constants: for suitable SBM parameters, a two-stage adaptive strategy succeeds with fewer linear-order queries than balanced uniform querying.

3.1 Uniform Oracle Access

For a sharp-homogeneous SBM benchmark, we impose a divisibility convention $k/n \in \mathbb{N}_0$ and each vertex is queried exactly $m := k/n$ times. Under this policy, the revealed graph is again an SBM with attenuated edge probabilities, so the classical Abbe-Bandeira-Hall threshold for exact recovery applies directly. More precisely, an existing edge $\{i,j\}$ is missed by all oracle queries at both endpoints with probability $(1-w)^{2m} = (1-w)^{2k/n}$, and is therefore revealed with probability $\theta_{\mathrm{unif}}(k,w) := 1-(1-w)^{2k/n}$. Consequently, the observed graph is exactly an SBM with effective parameters $\tilde{\alpha} = \alpha \cdot \theta_{\mathrm{unif}}(k,w)$ and $\tilde{\beta} = \beta \cdot \theta_{\mathrm{unif}}(k,w)$. Applying the classical exact recovery threshold gives the following benchmark.

Proposition 1 (Uniform oracle access benchmark)**.** *Let $G \sim SBM(n, \alpha \log n/n, \beta \log n/n)$. Suppose $k/n \in \mathbb{N}_0$ and the learner uses the uniform non-adaptive oracle policy with reliability $w \in [0,1]$. Then exact recovery from the observed graph is possible w.h.p. if and only if*

$$\theta_{\mathrm{unif}}(k,w)(\sqrt{\alpha}-\sqrt{\beta})^2 > 2, \qquad \theta_{\mathrm{unif}}(k,w) := 1-(1-w)^{2k/n}.$$

Thus balanced uniform querying succeeds exactly when the attenuated graph satisfies the threshold $\theta_{\mathrm{unif}}(k,w)(\sqrt{\alpha}-\sqrt{\beta})^2 > 2$. The proposition immediately yields the corresponding requirement for the query budget. Next, in Corollary 2, we show that whenever balanced uniform querying succeeds, it requires a *linear number of queries* for exact recovery.

Corollary 2. *Assume $w \in (0,1)$ and $\Delta^2 > 2$. Under the divisibility convention $k = mn$, exact recovery under uniform oracle access is possible if and only if $1-(1-w)^{2m} > \frac{2}{\Delta^2}$. Equivalently, the minimum number of queries per vertex is*

$$m_{\min} := \min\left\{ m \in \mathbb{N}_0 : 1-(1-w)^{2m} > \frac{2}{\Delta^2} \right\},$$

and the corresponding minimum budget is $k_{\min} = n m_{\min}$. In the continuous relaxation, the threshold corresponds to $m > \frac{\log\left(1-\frac{2}{\Delta^2}\right)}{2\log(1-w)}$. In particular, $k_{\min} = \Omega(n)$.

Remark (non-divisible budgets): If $k/n \notin \mathbb{N}_0$, a balanced allocation queries each vertex either $\lfloor k/n \rfloor$ or $\lceil k/n \rceil$ times. The resulting observation is not exactly homogeneous, but it is bracketed by the two neighboring homogeneous allocations. Thus Proposition 1 remains the correct benchmark up to a one-query-per-vertex rounding window. See Appendix B for precise arguments.

Algorithm 1: Two-stage adaptive screening strategy

Input: Oracle access with reliability w, first-stage budget k_0, refinement parameter L
Output: Community labels $\widehat{\sigma}$

```
// Stage 1
```

Query vertices according to a uniform first-stage allocation using budget k_0.

From first-stage transcript, compute preliminary labeling $\widehat{\sigma}^{(0)}$ and confidence scores $\forall v \in V$.

```
// Screening
```

Let A be the set of vertices whose first-stage local margin is too small.

Keep the labels $\widehat{\sigma}_v^{(0)}$ for vertices $v \notin A$.

```
// Stage 2
```

foreach $v \in A$ **do**

  Query v an additional L times.
  Compare the fresh neighbors of v against a leave-one-out reference labeling.
  Assign v by the resulting local likelihood test.

return the combined labeling: first-stage labels on A^c and refined labels on A.

3.2 Adaptive Oracle Access

We now turn to adaptive querying for exact recovery in Model 1. Here the main idea, summarized in Algorithm 1, is to separate recovery into two stages.[4] In the first stage, the learner uses part of the query budget to construct a coarse global sketch of the communities, and, based on this estimate, produces a preliminary labeling and identifies a small subset of vertices whose labels remain uncertain. In the second stage, the learner spends the remaining queries only on this uncertain set and uses the fresh observations to repair the remaining mistakes. In this section, we formalize when such a two-stage strategy succeeds and show that, for suitable parameters, it can achieve exact recovery with fewer linear-order queries than balanced uniform querying.

A screened-initialization condition. In this first stage, fix $x_0 > 0$ and perform $k_0 = x_0 n + o(n)$ uniform non-adaptive oracle queries. The resulting transcript is said to admit a leave-one-out screened initialization if it produces these global and local per-vertex auxiliary objects as follows. Globally, it outputs a preliminary labeling $\widehat{\sigma}^{(0)}$ and a screened set $\mathcal{A} \subseteq [n]$ whose labels need to be repaired. Locally, for each vertex $v \in [n]$, it produces a leave-one-out reference labeling $\widehat{\sigma}^{(-v)}$ together with two disjoint comparison sets $B_v \subseteq [n] \setminus \{v\}$ and $C_v \subseteq [n] \setminus \{B_v \cup \{v\}\}$. The labeling $\widehat{\sigma}^{(-v)}$ is computed without using the first-stage (edge) observations incident to v, and it is used only when classifying that same vertex v in the second stage. The screening set C_v is used only to decide whether $v \in \mathcal{A}$, while the reference set B_v is reserved for the second-stage likelihood test for v. This separation between B_v and C_v is imposed to avoid reusing the same incident observations from v both for initial screening and for final refinement: the event $\{v \in \mathcal{A}\}$ may depend on first-stage observations incident to v only through pairs $\{u, v\}$ with $u \in C_v$ and not through pairs with $u \in B_v$. Thus the edges from v to B_v are not conditioned on when deciding whether v is screened and remain as uncontaminated evidence for the refinement step. We further require B_v, C_v, and the labels $\widehat{\sigma}^{(-v)}_{B_v \cup C_v}$ (i.e., the collection of leave-one-out labels assigned to vertices in $B_v \cup C_v$) to be functions only of first-stage observations not incident to v. We formalize these aspects in the following.

Definition 1 (Leave-one-out screened initialization)**.** For $\gamma \in (0, 1)$ and $\zeta \in (0, 1]$, we say that a leave-one-out screened initialization is γ-screening with reference parameter ζ if with probability $1 - o(1)$, there exists a common global sign $s \in \{\pm 1\}$ such that

$$|\mathcal{A}| \le n^{\gamma}, \qquad \widehat{\sigma}_u^{(0)} = s\sigma_u^* \quad \text{for all } u \notin \mathcal{A}, \tag{2}$$

and, for every $v \in \mathcal{A}$,

$$\begin{aligned} \left|\{u \in B_v : \widehat{\sigma}_u^{(-v)} \neq s\sigma_u^*\}\right| &= o(n/\log n), \\ |B_v \cap \{u : \sigma_u^* = +1\}| &= \frac{\zeta n}{2} + o(n), \\ |B_v \cap \{u : \sigma_u^* = -1\}| &= \frac{\zeta n}{2} + o(n). \end{aligned} \tag{3}$$

[4] See Appendix A for a more detailed version of Algorithm 1.

In other words, Definition 1 conveys that after the first stage, vertices that may still be mislabeled are confined to a sublinear set $\mathcal{A}$, while every screened vertex $v \in \mathcal{A}$ comes with a large reference set B_v whose labels are almost correct and can therefore be used to classify v in the second stage. The parameter ζ measures how large this reference set is. The screening set C_v is kept separate from B_v so that the information used to decide whether v needs further refinement is different from the information used to determine its final label.

The next result, in Theorem 3, shows that any first-stage transcript satisfying Definition 1 can be extended to achieve exact recovery by querying only the screened vertices in $\mathcal{A}$.

Theorem 3 (Adaptive labeling from a screened initialization)**.** *Assume $\Delta^2 > 2$ and $w \in (0,1]$. Suppose the first-stage oracle transcript admits a γ-screening initialization with reference parameter ζ for some $\gamma \in (0,1)$ and $\zeta \in (0,1]$. Choose a fixed integer $L \geq 1$ such that*

$$\frac{\zeta}{2}\big(1-(1-w)^L\big)\Delta^2 > \gamma. \tag{4}$$

Then there is an adaptive second-stage strategy using at most $L\,|\mathcal{A}|$ additional oracle queries that recovers all labels in $\mathcal{A}$ correctly with probability $1-o(1)$. Consequently, the total query budget is

$$k = k_0 + L\,|\mathcal{A}| = x_0\, n + o(n)$$

on the event (2)–(3)*, and exact recovery succeeds w.h.p.*

Proof sketch. On the high-probability screening event, all vertices outside $\mathcal{A}$ are already correctly labeled, so it remains only to repair the vertices in $\mathcal{A}$. In the second stage, fix $v \in \mathcal{A}$, query v an additional L times and use only these fresh oracle observations from v to the reserved reference set B_v. For each $u \in B_v$, let $Y_u = 1$ if u is revealed at least once among these L fresh queries to v, and $Y_u = 0$ otherwise. We then classify v by a likelihood-ratio test based on the variables $(Y_u)_{u\in B_v}$, using the leave-one-out reference labels $\widehat{\sigma}_{B_v}^{(-v)}$. Intuitively, the test compares whether the fresh observations from v are more consistent with v belonging to the $+1$ side of the reference population or the -1 side. Since L fresh queries reveal an existing edge with probability $a_L := 1-(1-w)^L$, this local experiment is a one-vertex SBM test against an almost-correct reference population of size $|B_v| = \zeta n + o(n)$, with approximately $\zeta n/2$ vertices from each community. Therefore, the Bhattacharyya error exponent for one screened vertex is $\frac{\zeta}{2}\, a_L\, \Delta^2 + o(1) = \frac{\zeta}{2}\big(1-(1-w)^L\big)\Delta^2 + o(1)$. Hence the error probability for a single screened vertex is at most $n^{-\frac{\zeta}{2}(1-(1-w)^L)\Delta^2+o(1)}$. Since $|\mathcal{A}| \leq n^\gamma$ and the exponent is strictly larger than γ, a union bound implies that all vertices in $\mathcal{A}$ are recovered correctly w.h.p. The additional second-stage budget is $L\,|\mathcal{A}| = O(n^\gamma) = o(n)$, because L is fixed and $\gamma < 1$. □

A remaining question is whether the screened initialization can be obtained using fewer queries than balanced uniform exact recovery requires. The answer is yes whenever the first-stage budget is below the uniform threshold but still large enough to produce a screened initialization. Indeed, let $x_\star$ be the uniform exact recovery threshold, i.e., $\theta(x_\star, w)\Delta^2 = 2$. If, for some $x_0 < x_\star$, the first-stage uniform oracle graph generated using $x_0 n + o(n)$ queries admits a γ-screening initialization with reference parameter ζ for some $\gamma < 1$, and if L satisfies (4), then Theorem 3 gives exact recovery with $x_0 n + L|A| = x_0 n + o(n)$ queries. Since $x_0 < x_\star$, this is below the balanced-uniform exact recovery threshold by a linear margin $(x_\star - x_0)n - o(n)$.

Achieving screening below the uniform threshold. Next, we show that the screened-initialization condition can actually be achieved and it is not merely a formal abstraction. For an explicit parameter choice, a first-stage uniform query budget that is still below the balanced-uniform exact recovery threshold already provides two things: a leave-one-out labeling that is accurate up to only $n^{1/4+o(1)}$ errors, and a small screened set containing the remaining vertices with uncertain labels. An important aspect is that, although this first-stage budget is not large enough for exact recovery by itself, it is still large enough for a standard maximum-likelihood argument to produce a sufficiently accurate coarse estimate, and that weaker guarantee is enough for the second-stage repair step.

Recall that for parameters $a, b > 0$, we write

$$D_+(\alpha,\beta) := \frac{1}{2}\Delta^2(\alpha,\beta) = \frac{1}{2}(\sqrt{\alpha}-\sqrt{\beta})^2 = \frac{\alpha+\beta}{2} - \sqrt{\alpha\beta}.$$

For each $v \in [n]$, let $G^{(-v)}$ be the graph obtained by deleting all edges incident to v, and let

$$\mathcal{B}_{-v} := \{\tau \in \{\pm 1\}^{[n]\setminus\{v\}} : |\sum_{u\neq v} \tau_u| = 1\}.$$

Here, $\mathcal{B}_{-v}$ is the set of balanced labelings for the leave-one-out graph. The balanced leave-one-out MLE $\widehat{\sigma}^{(-v)}$ is any maximum-likelihood estimator for $G^{(-v)}$ over $\mathcal{B}_{-v}$. Its two global signs are fixed by the following deterministic anchor convention. Let $D_n = \{1, \dots, \lceil n^{3/4} \rceil\}$. For each v, orient $\widehat{\sigma}^{(-v)}$ so that $\sum_{u\in D_n\setminus\{v\}} \widehat{\sigma}_u^{(-v)} \geq 0$, breaking ties arbitrarily.

Lemma 4 (Screening estimate for the concrete first stage)**.** *Consider the oracle-only model with $\alpha = 16$ and $\beta = 1$. Let $\theta_0 := 1 - 2^{-1/3}$, and $w_0 := 1 - \sqrt{1-\theta_0}$. If the first stage queries every vertex exactly once with oracle reliability w_0, then the first-stage transcript admits a 0.88-screening initialization with reference parameter $1/2$, in the sense of* (2)–(3)*.*

Combining Lemma 4 with Theorem 3 gives the following concrete parameter choice, where the first-stage budget is below the balanced-uniform threshold and the adaptive two-stage repair still succeeds. Example 1 provides a concrete parameter regime to demonstrate that adaptivity can improve the constant in the linear query budget.

Example 1. Let $\alpha = 16$, $\beta = 1$ and set the oracle reliability parameter to $w = 1 - 2^{-1/6}$. In the first stage, we query every vertex exactly once, so we use $k_0 = n$ queries. The edge-retention factor of the corresponding uniform subsampled graph is $\theta_0 = 1 - (1-w)^2 = 1 - 2^{-1/3}$ and the graph is a homogeneous SBM with effective parameters $a = 16\,(1 - 2^{-1/3})$ and $b = 1 - 2^{-1/3}$. Applying Lemma 4 on this graph, we obtain a 0.88-screening initialization with reference parameter $\zeta = 1/2$.

The relaxed balanced-uniform exact recovery threshold is determined by $\theta(x, w)\Delta^2 = 2$, where $\theta(x, w) = 1-(1-w)^{2x}$ and $\Delta^2 = (\sqrt{\alpha}-\sqrt{\beta})^2 = 9$. Thus $x_\star$ is the solution of $9\,(1-(1-w)^{2\,x_\star}) = 2$, precisely $x_\star = 3\,\log_2(9/7) > 1$. Since the first stage uses $x_0 = 1$ query per vertex, it operates strictly below the balanced-uniform exact recovery threshold.

Now take $L = 5$. Using $\zeta = 1/2$ and $\Delta^2 = 9$, we obtain $\frac{\zeta}{2}(1-(1-w)^L)\Delta^2 = \frac{1}{4}(1-2^{-5/6})\cdot 9 > 0.88$. The adaptive strategy succeeds with the query budget $k = n + 5n^{0.88+o(1)} = n + o(n)$. In contrast, balanced uniform querying requires $x_\star n = 3\log_2(9/7)n \approx 1.087n$ queries. Thus, for this parameter choice and for sufficiently large n, adaptive querying achieves exact recovery with strictly fewer linear-order queries than uniform querying. ◇

Next, we establish a stronger result that this is the strongest possible type of separation in the oracle-only model, i.e., any adaptive strategy still requires a linear number of queries for exact recovery.

Theorem 5. *In Model 1 every oracle-only exact recovery strategy requires $k = \Omega(n)$ queries.*

Proof sketch. We prove the contrapositive. If an oracle-only strategy uses $k = o(n)$ queries, then only $o(n)$ vertices are ever queried. Since the underlying SBM has maximum degree $O(\log n)$ with high probability, the total number of vertex appearances in all oracle responses is only $o(n \log n)$. Hence there are $n - o(n)$ vertices that are never queried and appear only $o(\log n)$ times in the transcript; by exact balance, this set contains linearly many vertices from each community. For any such $+$-vertex i and $-$-vertex j, compare the truth with the balanced labeling obtained by swapping i and j. Since neither vertex is queried and both appear rarely, this swap changes the transcript likelihood by only a small polynomial factor, while preserving the balance constraint. Thus there are linearly many distinct balanced swapped labelings with likelihood comparable to that of the truth, so the posterior mass of the true labeling is $o(1)$. Consequently even the MAP estimator fails with probability tending to one, and exact recovery is impossible when $k = o(n)$. Therefore every successful oracle-only strategy requires $k = \Omega(n)$. □

4 Community Recovery with Additional Correlated Information

We now turn to Model 2, where the learner first observes a subsampled graph G_{sub} and then uses the noisy neighborhood oracle for refinement. The main question is whether this additional correlated side information G_{sub} makes adaptive querying more effective than uniform querying, and how this contrasts with the oracle-only setting. We show that, under a sublinear budget, uniform querying

cannot improve the exact recovery threshold beyond that of G_{sub} alone. In contrast, under the same sublinear budget, an adaptive strategy can use G_{sub} to identify uncertain vertices and concentrate the oracle budget on them, leading to exact recovery. Thus, unlike in the oracle-only model where the adaptivity gap is in linear order, the presence of correlated side information results in a genuine sublinear adaptivity gap.

4.1 Why does Uniform Querying Under a Sublinear Budget Fail?

The non-adaptive baseline spreads the oracle budget as uniformly as possible, independently of the observed graph. In the sublinear regime this should be interpreted in the integer balanced sense: only $k = o(n)$ vertices receive an oracle query, while $n - o(n)$ vertices receive no oracle information at all. Thus the combined observation is not exactly a homogeneous SBM with a new edge-retention parameter. Nevertheless, the conclusion is the same as the homogeneous heuristic would suggest. Since the queried set is data-independent and contains only $o(n)$ vertices, uniform querying cannot target the vertices that are ambiguous under G_{sub}. Hence it cannot change the exact recovery threshold determined by the subsampled graph alone. The following lemma formalizes this statement.

Lemma 6 (Uniform querying fails)**.** *Consider Model 2 with subsampling rate $s \in (0,1)$, oracle reliability $w \in (0,1]$, and query budget $k = o(n)$. Suppose the learner uses a data-independent balanced uniform querying strategy: the queried set is fixed in advance, or chosen independently of G_{sub} and the oracle responses. Then, in the sublinear-budget regime, uniform querying does not improve the exact recovery threshold beyond that of G_{sub} alone. In particular, exact recovery under uniform querying is impossible w.h.p. whenever $s\,\Delta^2 < 2$.*

Proof sketch. With $k = o(n)$, a balanced uniform strategy queries only $o(n)$ vertices, independently of G_{sub} and of the labels. Hence $n - o(n)$ vertices receive no oracle information at all. For any two unqueried vertices of opposite communities, the oracle can affect their comparison to only the $o(n)$ queried vertices; their evidence from the remaining $n - o(n)$ vertices is exactly the evidence available in G_{sub}. Thus the local exact recovery exponent for a balanced swap of two unqueried opposite-community vertices is the same as for the subsampled graph alone, namely $s\Delta^2$. If $s\Delta^2 < 2$, the usual ABH second-moment converse gives, with high probability, at least one likelihood-competitive balanced swap among unqueried vertices. Therefore the posterior cannot concentrate on the true balanced labeling, and uniform sublinear querying cannot achieve exact recovery. □

4.2 Adaptive Targeting Succeeds

The failure in Lemma 6 is not necessarily due to the oracle being too weak; rather, it may stem from the fact that uniform querying ignores where the subsampled graph remains (un)certain. Once the querying strategy can exploit the correlated information from G_{sub}, this limitation disappears.

Lemma 7. *Consider Model 2, a subsampling rate $s \in (0,1)$, and oracle reliability $w \in (0,1]$. There exist constants (α, β, s, w) and a budget $k = o(n)$ such that a two-stage strategy that uses G_{sub} to* target *queries succeeds at exact recovery w.h.p.*

Proof sketch. Take, for example, $\alpha = 16$, $\beta = 1$, $w = 1/2$, and choose $s = 2(1-\epsilon)/9$ for a sufficiently small constant $\epsilon > 0$. Then G_{sub} is below the exact recovery threshold, but it still gives a leave-one-out preliminary labeling with only $n^{\epsilon+o(1)}$ errors. Using this preliminary labeling, compute for each vertex a local likelihood score from a reserved part of its subsampled incident edges, and screen only vertices whose score has small margin. Large-deviation bounds show that the screened set has size at most $n^{0.77}$ w.h.p., while every unscreened vertex is already correctly labeled. Query each screened vertex once and classify it using fresh oracle observations against a disjoint reference set. The oracle likelihood test has error probability $n^{-0.7875+o(1)}$ per screened vertex, so a union bound over $n^{0.77}$ screened vertices succeeds. Thus $k = \lceil n^{0.78} \rceil = o(n)$ queries suffice for exact recovery w.h.p. □

Combining the uniform-querying lower bound with the adaptive construction yields the main separation of this section: with the same sublinear oracle budget, data-independent querying cannot improve on the subsampled graph, whereas adaptive querying achieves exact recovery.

Theorem 8. *In Model 2, there exist SBM parameters $\alpha, \beta > 0$, a subsampling rate $s \in (0,1)$, an oracle reliability $w \in (0,1]$ and a query budget $k = o(n)$ such that simultaneously:*

1. *Every data-independent uniform querying strategy fails to achieve exact recovery w.h.p.*
2. *There exists a two-stage adaptive querying strategy that uses G_{sub} to target queries and achieves exact recovery with high probability.*

5 Simulations

To illustrate the efficacy of our adaptive scheme, we compare simulations using adaptive uncertainty-based querying with balanced, data-independent uniform querying in balanced two-community SBMs with $p = \frac{16 \log n}{n}$ and $q = \frac{\log n}{n}$. Each oracle call independently returns each true neighbor with probability w, without resampling the underlying graph. Success means exact recovery of the entire partition, up to a global label swap. We focus on Model 2, which exhibits a qualitative separation at sublinear additional query budgets. Model 1 instead improves the leading constant of a linear budget; directly illustrating this asymptotic improvement using the sufficient theoretical bounds is impractical at reasonable graph sizes.

We start with a subsampled graph that retains each true edge independently with probability $s = 1.9/9$, and the oracle reliability is $w = 1/2$. Both active policies make $k = \lceil n^{0.4} \rceil$ additional calls, one per selected vertex. In 30 trials per graph size instance, the adaptive policy ranges between 93% and 97% success rates in just 13 to 21 additional calls, compared to uniform querying that never exceeds 24% (Table 2). Adaptive selection covers approximately 98% of the initial labeling errors.

Table 2: Model 2 success rates with Wilson 95% intervals from 30 independent paired trials per size.

n	k	Passive	Uniform	Adaptive
500	13	$0.267\,[0.142, 0.444]$	$0.233\,[0.118, 0.409]$	$0.967\,[0.833, 0.994]$
1000	16	$0.167\,[0.073, 0.336]$	$0.233\,[0.118, 0.409]$	$0.967\,[0.833, 0.994]$
2000	21	$0.133\,[0.053, 0.297]$	$0.167\,[0.073, 0.336]$	$0.933\,[0.787, 0.982]$

6 Conclusion

We studied exact community recovery in the two-community SBM under budgeted noisy-neighborhood-oracle access. Our models make data acquisition part of the inference problem: the learner must decide not only how to recover communities from partial observations, but also where to acquire those observations. In Model 1, balanced uniform querying gives a sharp non-adaptive benchmark, but a two-stage adaptive strategy can succeed with $n + o(n)$ queries in regimes where uniform querying requires mn queries for some $m > 1$. In Model 2, a subsampled graph can identify locally ambiguous vertices, allowing adaptive querying to focus a sublinear budget and yielding a genuine adaptivity gap. Thus, under the same observation budget, different acquisition rules can lead to fundamentally different exact recovery behavior.

Scope and Limitations. Our results are asymptotic and apply to the canonical balanced two-community logarithmic-degree SBM with independent one-sided oracle noise. In this context, we isolate the role of adaptive data acquisition and prove concrete information-theoretic gains over uniform querying. This work leaves open the problem of identifying the (α, β) pairs for which adaptivity leads to strictly fewer queries, and characterizing the optimal number of adaptive queries.

Future Work. By treating data acquisition as part of the inference problem, we open several directions for future work, including design of general adaptive querying frameworks that match or outperform uniform querying across broader SBM regimes, and closing the remaining exact threshold gaps.

Societal Impact. This work is theoretical and does not involve real data, human subjects, or deployed systems. Its potential positive impact is to clarify when limited local measurements can reduce the cost of network recovery. At the same time, improved community recovery from partial observations could be misused in privacy-sensitive network settings. Any practical deployment should therefore account for consent, privacy constraints, and the sensitivity of inferred communities.

Funding and Acknowledgments

Suhas Thejaswi acknowledges support from the Technology Industries of Finland Centennial Foundation through a grant awarded to Aalto University. The authors declare that the funding source does not create any conflict of interest in relation to this work. We thank the anonymous reviewers of NeurIPS 2026 for their thoughtful and detailed feedback, which helped improve this manuscript.

References

[1] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.

[2] Emmanuel Abbe and Colin Sandon. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In *Symposium on Foundations of Computer Science*, pages 670–688. IEEE, 2015.

[3] Emmanuel Abbe and Colin Sandon. Recovering communities in the general stochastic block model without knowing the parameters. *Advances in neural information processing systems*, 28, 2015.

[4] Emmanuel Abbe, Afonso S. Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, 2016.

[5] Emmanuel Abbe, Jianqing Fan, Kaizheng Wang, and Yiqiao Zhong. Entrywise eigenvector analysis of random matrices with low expected rank. *Annals of Statistics*, 48(3):1452, 2020.

[6] Jess Banks, Cristopher Moore, Joe Neeman, and Praneeth Netrapalli. Information-theoretic thresholds for community detection in sparse networks. In *Conference on Learning Theory*, pages 383–416. PMLR, 2016.

[7] Zalán Bodó, Zsolt Minier, and Lehel Csató. Active learning with clustering. In *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*, pages 127–139. JMLR Workshop and Conference Proceedings, 2011.

[8] Cécile Bothorel, Juan David Cruz, Matteo Magnani, and Barbora Micenkova. Clustering attributed graphs: models, measures and methods. *Network Science*, 3(3):408–444, 2015.

[9] Narasinga Rao Chaganty and Jayaram Sethuraman. Strong large deviation and local limit theorems. *The Annals of Probability*, pages 1671–1690, 1993.

[10] Shuwen Chai and Miklós Z Rácz. Efficient graph matching for correlated stochastic block models. *Advances in Neural Information Processing Systems*, 37:116388–116461, 2024.

[11] Guanyi Chen, Jian Ding, Shuyang Gong, and Zhangsong Li. A computational transition for detecting correlated stochastic block models by low-degree polynomials. *The Annals of Statistics*, 54(1):226–251, 2026.

[12] 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*, pages 391–423. PMLR, 2015.

[13] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 84(6):066106, 2011.

[14] Shaofeng Deng, Shuyang Ling, and Thomas Strohmer. Strong consistency, graph Laplacians, and the stochastic block model. *Journal of Machine Learning Research*, 22(117):1–44, 2021.

[15] Jingqiu Ding, Yiding Hua, Lucas Slot, and David Steurer. Low-degree evidence for computational transition of recovery rate in stochastic block model. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[16] Akshay Gadde, Eyal En Gad, Salman Avestimehr, and Antonio Ortega. Active learning for community detection in stochastic block models. In *International Symposium on Information Theory*, pages 1889–1893. IEEE, 2016.

[17] Akshay Gadde, Eyal En Gad, Salman Avestimehr, and Antonio Ortega. Active learning for community detection in stochastic block models. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 1889–1893. IEEE, 2016.

[18] Chao Gao, Zongming Ma, Anderson Y. Zhang, and Harrison H. Zhou. Achieving optimal misclassification proportion in stochastic block models. *J. Mach. Learn. Res.*, 18:60:1–60:45, 2017.

[19] Julia Gaudio and Nirmit Joshi. Community detection in the hypergraph SBM: Exact recovery given the similarity matrix. In *The Thirty Sixth Annual Conference on Learning Theory*, pages 469–510. PMLR, 2023.

[20] Julia Gaudio, Miklos Z. Racz, and Anirudh Sridhar. Exact community recovery in correlated stochastic block models. In *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 2183–2241, 2022.

[21] Mark S Handcock and Krista J Gile. Modeling social networks from sampled data. *The Annals of Applied Statistics*, 4(1):5, 2010.

[22] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.

[23] Di Jin, Zhizhi Yu, Pengfei Jiao, Shirui Pan, Dongxiao He, Jia Wu, Philip S Yu, and Weixiong Zhang. A survey of community detection approaches: From statistical modeling to deep learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(2):1149–1170, 2021.

[24] Varun Kanade, Elchanan Mossel, and Tselil Schramm. Global and local information in clustering labeled block models. *IEEE Trans. Inf. Theory*, 62(10):5906–5917, 2016.

[25] Brian Karrer and Mark EJ Newman. Stochastic blockmodels and community structure in networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 83(1):016107, 2011.

[26] Chuyang Ke and Jean Honorio. Information-theoretic limits for community detection in network models. *Advances in Neural Information Processing Systems*, 31, 2018.

[27] Dan Kushnir and Benjamin Mirabelli. Active community detection with maximal expected model change. In *International Conference on Artificial Intelligence and Statistics*, pages 724–734. PMLR, 2020.

[28] Timothy LaRock, Timothy Sakharov, Sahely Bhadra, and Tina Eliassi-Rad. Understanding the limitations of network online learning. *Applied Network Science*, 5(1):60, 2020.

[29] Elchanan Mossel, Joe Neeman, and Allan Sly. Consistency thresholds for the planted bisection model. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 69–75, 2015.

[30] Chung Kyong Nguen, Oscar Hernan Madrid Padilla, and Arash A Amini. Network two-sample test for block models. *arXiv preprint arXiv:2406.06014*, 2024.

[31] Tiago P Peixoto. Hierarchical block structures and high-resolution model selection in large networks. *Physical Review X*, 4(1):011047, 2014.

[32] Miklos Racz and Anirudh Sridhar. Correlated stochastic block models: Exact graph matching with applications to recovering communities. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

[33] Miklós Z. Rácz and Jifan Zhang. Harnessing multiple correlated networks for exact community recovery. In *Advances in Neural Information Processing Systems*, volume 37. Curran Associates, Inc., 2024.

[34] Hussein Saad and Aria Nosratinia. Community detection with side information: Exact recovery under the stochastic block model. *J. Sel. Top. Signal Process.*, 12(5):944–958, 2018.

[35] Jeffrey A Smith, James Moody, and Jonathan H Morgan. Network sampling coverage ii: The effect of non-random missing data on network measurement. *Social networks*, 48:78–99, 2017.

[36] Sucheta Soundarajan, Tina Eliassi-Rad, Brian Gallagher, and Ali Pinar. ϵ - WGX: adaptive edge probing for enhancing incomplete networks. In Peter Fox, Deborah L. McGuinness, Lindsay Poirier, Paolo Boldi, and Katharina Kinder-Kurlanda, editors, *Proceedings of the ACM on Web Science Conference*, pages 161–170.

[37] Michael PH Stumpf and Mason A Porter. Critical truths about power laws. *Science*, 335(6069): 665–666, 2012.

[38] Michael PH Stumpf, Carsten Wiuf, and Robert M May. Subnets of scale-free networks are not scale-free: sampling properties of networks. *Proceedings of the National Academy of Sciences*, 102(12):4221–4224, 2005.

[39] Yuanchen Wu and Yubai Yuan. Robust offline active learning on graphs. *Advances in Neural Information Processing Systems*, 37:58955–58983, 2024.

[40] Tzu-Chi Yen and Daniel B Larremore. Community detection in bipartite networks with stochastic block models. *Physical Review E*, 102(3):032309, 2020.

[41] Se-Young Yun and Alexandre Proutière. Community detection via random and adaptive sampling. In *Proceedings of the Conference on Learning Theory*, volume 35, pages 138–175. JMLR.org, 2014.

[42] Anderson Ye Zhang. Fundamental limits of spectral clustering in stochastic block models. *IEEE Transactions on Information Theory*, 70(10):7320–7348, 2024.

A Algorithm 1 with more details

Recall that, we write $V = [n]$, where n is even, and $p = \alpha\frac{\log n}{n}, q = \beta\frac{\log n}{n}, \alpha > \beta > 0$. The model parameters are supplied to the likelihood tests. Define $\mathrm{sgn}_+(x) = +1$ for $x \geq 0$ and $\mathrm{sgn}_+(x) = -1$ otherwise. For a deterministic first-stage allocation $(r_u)_{u\in V}$, let $A^{(0)}_{ij}$ indicate whether the edge $\{i, j\}$ is revealed at least once by a first-stage query to either endpoint. The corresponding edge-retention probability is

$$\theta_{ij} = 1 - (1 - w)^{r_i + r_j}.$$

For a candidate labeling $z \in \{\pm 1\}^{V\setminus\{v\}}$, define

$$P_{ij}(z) = \theta_{ij}\left(p\,\mathbf{1}\{z_i = z_j\} + q\,\mathbf{1}\{z_i \neq z_j\}\right).$$

The balanced leave-one-out likelihood is

$$\mathcal{L}_{-v}(z) = \prod_{\substack{i<j\\ i,j\neq v}} P_{ij}(z)^{A^{(0)}_{ij}}\left(1 - P_{ij}(z)\right)^{1-A^{(0)}_{ij}}, \qquad z \in \mathcal{B}_{-v}, \tag{5}$$

where

$$\mathcal{B}_{-v} = \left\{z \in \{\pm 1\}^{V\setminus\{v\}} : \left|\sum_{u\neq v} z_u\right| = 1\right\}.$$

Each factor is interpreted as a Bernoulli probability mass, including at parameter values zero or one. Thus, $\mathcal{L}_{-v}$ uses only pairs not incident to v. Using this notation, we give a more detailed version of Algorithm 1 in Algorithm 2.

B Omitted proofs from Section 3

For the sharp homogeneous-SBM benchmark, we impose the divisibility convention $k/n \in \mathbb{N}_0$. Writing $m := k/n$, the uniform policy queries every vertex exactly m times. Hence an existing edge $\{i, j\}$ is missed by all oracle queries at both endpoints with probability $(1 - w)^{2m}$, and is therefore revealed with probability: $\theta_{\mathrm{unif}}(k, w) := 1-(1-w)^{2m} = 1-(1-w)^{2k/n}$. Consequently, the observed graph is exactly an SBM with effective parameters

$$\widetilde{\alpha} = \alpha\theta_{\mathrm{unif}}(k, w), \qquad \widetilde{\beta} = \beta\theta_{\mathrm{unif}}(k, w).$$

Applying the Abbe–Bandeira–Hall threshold gives the following benchmark.

Proposition 1 (Uniform oracle access benchmark)**.** *Let $G \sim SBM(n, \alpha\log n/n, \beta\log n/n)$. Suppose $k/n \in \mathbb{N}_0$ and the learner uses the uniform non-adaptive oracle policy with reliability $w \in [0, 1]$. Then exact recovery from the observed graph is possible w.h.p. if and only if*

$$\theta_{\mathrm{unif}}(k, w)(\sqrt{\alpha} - \sqrt{\beta})^2 > 2, \qquad \theta_{\mathrm{unif}}(k, w) := 1 - (1 - w)^{2k/n}.$$

Proof. Under the uniform policy, each endpoint of an existing edge is queried $m = k/n$ times. Thus the edge is revealed with probability $\theta_{\mathrm{unif}}(k, w) = 1 - (1 - w)^{2m}$, independently across edges conditional on the labels. Hence the observed graph is an SBM with parameters

$$\widetilde{p} = \alpha\theta_{\mathrm{unif}}(k, w)\frac{\log n}{n}, \qquad \widetilde{q} = \beta\theta_{\mathrm{unif}}(k, w)\frac{\log n}{n}.$$

The Abbe–Bandeira–Hall threshold applied to this attenuated SBM gives

$$(\sqrt{\widetilde{\alpha}} - \sqrt{\widetilde{\beta}})^2 > 2,$$

which is equivalent to

$$\theta_{\mathrm{unif}}(k, w)(\sqrt{\alpha} - \sqrt{\beta})^2 > 2.$$

□

Algorithm 2: Two-stage adaptive screening strategy (detailed version)

Input: Vertex set $V = [n]$;
oracle with reliability $w \in (0,1]$;
SBM parameters $\alpha > \beta > 0$;
first-stage budget $k_0 \in \mathbb{N}_0$;
screening threshold $\tau > 0$;
refinement count $L \in \mathbb{N}$.

Output: Community labels $\widehat{\sigma} \in \{\pm 1\}^V$.

`// Stage 1: deterministic balanced allocation`

$m \leftarrow \lfloor k_0/n \rfloor$; $r \leftarrow k_0 - mn$

foreach $u \in V$ **do**

  $r_u \leftarrow m + \mathbf{1}\{u \le r\}$

  Query u exactly r_u times and record the responses

Form $A^{(0)}$, with $A^{(0)}_{ij} = 1$ if either endpoint revealed the other during Stage 1, and $A^{(0)}_{ij} = 0$ otherwise

$D_n \leftarrow \{1, \dots, \lceil n^{3/4} \rceil\}$

`// Leave-one-out references and screening`

foreach $v \in V$ **do**

  Compute $\widetilde{\sigma}^{(-v)} \in \arg\max_{z \in \mathcal{B}_{-v}} \mathcal{L}_{-v}(z)$ using (5)

  $h_v \leftarrow \mathrm{sgn}_+\Big(\sum_{u \in D_n \setminus \{v\}} \widetilde{\sigma}^{(-v)}_u\Big)$

  $\widehat{\sigma}^{(-v)} \leftarrow h_v \widetilde{\sigma}^{(-v)}$

  Let B_v be the first $\lfloor (n-1)/2 \rfloor$ vertices of $V \setminus \{v\}$ in increasing order

  $C_v \leftarrow V \setminus (B_v \cup \{v\})$

  $\Lambda^{\mathrm{loo}}_v \leftarrow \sum_{u \in C_v} \widehat{\sigma}^{(-v)}_u A^{(0)}_{uv}$

  $\widehat{\sigma}^{(0)}_v \leftarrow \mathrm{sgn}_+(\Lambda^{\mathrm{loo}}_v)$

$\mathcal{A} \leftarrow \Big\{ v \in V : |\Lambda^{\mathrm{loo}}_v| \le \tau \log n \Big\}$

Freeze $\mathcal{A}$, $(B_v, C_v)_{v \in V}$, and all leave-one-out reference labelings before Stage 2

$\widehat{\sigma} \leftarrow \widehat{\sigma}^{(0)}$

`// Stage 2: fresh observations on reserved pairs`

$a_L \leftarrow 1 - (1-w)^L$

foreach $v \in \mathcal{A}$ **do**

  Query v exactly L additional times; denote the returned neighbor sets by $N^{(2)}_{v,1}, \dots, N^{(2)}_{v,L}$

  foreach $u \in B_v$ **do**

    $Y_{vu} \leftarrow \mathbf{1}\Big\{ u \in \bigcup_{\ell=1}^{L} N^{(2)}_{v,\ell} \Big\}$

  Compute the local log-likelihood ratio

$$\Lambda^{(2)}_v \leftarrow \sum_{u \in B_v} \widehat{\sigma}^{(-v)}_u \left[Y_{vu} \log \frac{p}{q} + (1 - Y_{vu}) \log \frac{1 - a_L p}{1 - a_L q} \right].$$

  $\widehat{\sigma}_v \leftarrow \mathrm{sgn}_+(\Lambda^{(2)}_v)$

return $\widehat{\sigma}$

Corollary 2. *Assume $w \in (0,1)$ and $\Delta^2 > 2$. Under the divisibility convention $k = mn$, exact recovery under uniform oracle access is possible if and only if $1 - (1-w)^{2m} > \frac{2}{\Delta^2}$. Equivalently, the minimum number of queries per vertex is*

$$m_{\min} := \min\left\{ m \in \mathbb{N}_0 : 1 - (1-w)^{2m} > \frac{2}{\Delta^2} \right\},$$

and the corresponding minimum budget is $k_{\min} = n m_{\min}$. In the continuous relaxation, the threshold corresponds to $m > \frac{\log\left(1 - \frac{2}{\Delta^2}\right)}{2\log(1-w)}$. In particular, $k_{\min} = \Omega(n)$.

Proof. Under the uniform policy, each endpoint of an existing edge is queried $m = k/n$ times. Thus the edge is revealed with probability $\theta_{\mathrm{unif}}(k, w) = 1 - (1-w)^{2m}$, independently across edges

conditional on the labels. Hence the observed graph is an SBM with parameters

$$\widetilde{p} = \alpha\theta_{\mathrm{unif}}(k,w)\frac{\log n}{n}, \qquad \widetilde{q} = \beta\theta_{\mathrm{unif}}(k,w)\frac{\log n}{n}.$$

The Abbe–Bandeira–Hall threshold applied to this attenuated SBM gives

$$(\sqrt{\widetilde{\alpha}} - \sqrt{\widetilde{\beta}})^2 > 2,$$

which is equivalent to

$$\theta_{\mathrm{unif}}(k,w)(\sqrt{\alpha} - \sqrt{\beta})^2 > 2.$$

□

If $w = 1$, then $\theta_{\mathrm{unif}}(k,1) = 1$ for every $m = k/n \geq 1$. Thus, under the divisibility convention, one query per vertex suffices exactly when

$$(\sqrt{\alpha} - \sqrt{\beta})^2 > 2,$$

and in this case $k_{\min} = n$. If $w = 0$, then $\theta_{\mathrm{unif}}(k,0) = 0$, so no edge is ever revealed and exact recovery is impossible.

Remark (Non-divisible budgets). If $k/n \notin \mathbb{N}$, let $r = \lfloor k/n \rfloor$. A balanced allocation queries each vertex either r or $r+1$ times. Conditional on this allocation, the observed graph is generally not a homogeneous SBM, because $\rho_{ij} = 1 - (1-w)^{R_i+R_j}$ depends on the endpoint query counts.

The homogeneous theorem above nevertheless gives a useful bracket. The balanced allocation is at least as informative as the homogeneous allocation with r queries per vertex, since the learner may ignore any extra queries. Conversely, it is no more informative than the homogeneous allocation with $r+1$ queries per vertex. Therefore, with $\theta_\ell := 1 - (1-w)^{2\ell}$, we have

$$\theta_r(\sqrt{\alpha} - \sqrt{\beta})^2 > 2 \quad \Longrightarrow \quad \text{balanced uniform querying succeeds,}$$

whereas

$$\theta_{r+1}(\sqrt{\alpha} - \sqrt{\beta})^2 < 2 \quad \Longrightarrow \quad \text{balanced uniform querying fails.}$$

Thus the non-divisible case is bracketed by the two neighboring homogeneous allocations. We discuss the genuinely sublinear-budget regime separately in Lemma 6.

Theorem 3 (Adaptive labeling from a screened initialization)**.** *Assume $\Delta^2 > 2$ and $w \in (0,1]$. Suppose the first-stage oracle transcript admits a γ-screening initialization with reference parameter ζ for some $\gamma \in (0,1)$ and $\zeta \in (0,1]$. Choose a fixed integer $L \geq 1$ such that*

$$\frac{\zeta}{2}\big(1 - (1-w)^L\big)\Delta^2 > \gamma. \tag{4}$$

Then there is an adaptive second-stage strategy using at most $L\,|\mathcal{A}|$ additional oracle queries that recovers all labels in $\mathcal{A}$ correctly with probability $1 - o(1)$. Consequently, the total query budget is

$$k = k_0 + L\,|\mathcal{A}| = x_0\, n + o(n)$$

on the event (2)–(3)*, and exact recovery succeeds w.h.p.*

Proof. On the event (2)–(3), fix the common global sign s. For each $v \in \mathcal{A}$, the adaptive stage queries v exactly L additional times. It ignores all responses from vertices outside the reference set B_v and classifies v using only the responses from v to B_v and the labels $\widehat{\sigma}_{B_v}^{(-v)}$. By the split measurability condition, the first-stage data that determine $\{v \in \mathcal{A}\}$ do not use observations on pairs $\{u,v\}$ with $u \in B_v$. Hence, after conditioning on the first-stage sigma-field that defines $\{v \in \mathcal{A}\}$, B_v, C_v, and $\widehat{\sigma}_{B_v \cup C_v}^{(-v)}$, the edge variables $\{G^*_{uv} : u \in B_v\}$ are still independent conditional on the labels, and the new oracle coins are fresh.

Let

$$a_L := 1 - (1-w)^L.$$

For $u \in B_v$, define $Y_u = 1$ if at least one of the L new queries to v returns u, and $Y_u = 0$ otherwise. Conditional on the labels, the variables $\{Y_u : u \in B_v\}$ are independent Bernoulli random variables

with parameter $a_L p$ when $\sigma_u^* = \sigma_v^*$ and $a_L q$ when $\sigma_u^* \neq \sigma_v^*$. By (3), B_v contains $\zeta n/2 + o(n)$ vertices of each sign and therefore $|B_v| = \zeta n + o(n)$.

We now spell out how the second stage assigns a label to a fixed vertex $v \in \mathcal{A}$. The algorithm treats the labels in B_v as a reference set and uses only the fresh observations Y_u from the additional queries to v. The only quantity being decided in this calculation is whether v has the same label as the $+1$ side of the reference set or the -1 side. First suppose that the reference labels are exact, after applying the common global sign. For a candidate value $\tau \in \{\pm 1\}$ of $s\sigma_v^*$, the conditional log-likelihood of the observations $Y = (Y_u)_{u \in B_v}$ is, up to constants independent of τ,

$$\ell_v(\tau; Y) = \sum_{u \in B_v} \mathbf{1}\{\widehat{\sigma}_u^{(-v)} = \tau\} \left[Y_u \log(a_L p) + (1 - Y_u) \log(1 - a_L p) \right] + \mathbf{1}\{\widehat{\sigma}_u^{(-v)} \neq \tau\} \left[Y_u \log(a_L q) + (1 - Y_u) \log(1 - a_L q) \right].$$

Thus the log-likelihood ratio in favor of $s\sigma_v^* = +1$ over $s\sigma_v^* = -1$ is

$$\Lambda_v^{(2)} = \sum_{u \in B_v} \widehat{\sigma}_u^{(-v)} \left[Y_u \log \frac{p}{q} + (1 - Y_u) \log \frac{1 - a_L p}{1 - a_L q} \right]. \tag{6}$$

The second-stage decision rule is: set $\widehat{\sigma}_v^{(1)} = +1$ if $\Lambda_v^{(2)} \geq 0$ and $\widehat{\sigma}_v^{(1)} = -1$ otherwise, with the final labels interpreted up to the same common global sign. In words, the algorithm compares whether the fresh neighbors revealed by querying v look more like connections to vertices labeled $+1$ or to vertices labeled -1 in the reference set.

The corresponding Bhattacharyya coefficient is also explicit. For one reference vertex, the two hypotheses swap the Bernoulli parameters $a_L p$ and $a_L q$, so the one-coordinate coefficient is

$$b_L := \sqrt{a_L p \, a_L q} + \sqrt{(1 - a_L p)(1 - a_L q)}.$$

Since the coordinates are conditionally independent and $|B_v| = \zeta n + o(n)$, the product Bhattacharyya coefficient is

$$\mathrm{BC}_v = b_L^{|B_v|} = b_L^{\zeta n + o(n)}.$$

The likelihood-ratio test has error probability at most this coefficient.

We next return to the actual reference labels. Let $E_v = \{u \in B_v : \widehat{\sigma}_u^{(-v)} \neq s\sigma_u^*\}$. By (3), $|E_v| = o(n/\log n)$. Relative to the exact reference likelihood ratio, only the terms indexed by E_v are changed. The non-edge contribution of these terms is at most

$$|E_v| \left| \log \frac{1 - a_L p}{1 - a_L q} \right| = o(1),$$

while the number of newly revealed edges from v to E_v has mean $O(|E_v| \log n/n) = o(1)$. Hence the mislabeled reference vertices perturb (6) by $o(\log n)$ with probability $1 - o(1)$, uniformly over the screened vertices. Therefore the same Bhattacharyya exponent remains valid up to an $o(\log n)$ error. Using $p = \alpha \log n/n$ and $q = \beta \log n/n$,

$$\begin{aligned} b_L &= \sqrt{a_L p \, a_L q} + \sqrt{(1 - a_L p)(1 - a_L q)} \\ &= a_L \sqrt{pq} + 1 - \frac{a_L (p + q)}{2} + O(p^2 + q^2) \\ &= 1 - \frac{a_L}{2} (\sqrt{p} - \sqrt{q})^2 + O\left(\frac{\log^2 n}{n^2} \right) \\ &= 1 - \frac{a_L \Delta^2}{2} \frac{\log n}{n} + O\left(\frac{\log^2 n}{n^2} \right). \end{aligned}$$

Consequently,

$$\mathrm{BC}_v = \left(1 - \frac{a_L \Delta^2}{2} \frac{\log n}{n} + O\left(\frac{\log^2 n}{n^2} \right) \right)^{\zeta n + o(n)} = n^{-\zeta a_L \Delta^2/2 + o(1)}.$$

Therefore the conditional probability that the likelihood-ratio test misclassifies v is at most

$$n^{-\zeta a_L \Delta^2/2 + o(1)}.$$

By (4), choose $\eta > 0$ such that $\zeta a_L \Delta^2/2 \geq \gamma + 2\eta$. A union bound over the at most n^γ vertices in $\mathcal{A}$ gives

$$\Pr(\text{some } v \in \mathcal{A} \text{ is misclassified and the screening event holds}) \leq n^\gamma n^{-\zeta a_L \Delta^2/2 + o(1)} \leq n^{-\eta + o(1)} = o(1).$$

All vertices outside $\mathcal{A}$ are correct on (2) by assumption, so the combined labeling is correct up to the common global sign s w.h.p. The query count is $k_0 + L|\mathcal{A}| = x_0 n + O(n^\gamma) = x_0 n + o(n)$ on the same event. □

Before we can prove Lemma 4, we need the following result.

Lemma 9 (Uniform leave-one-out rough initialization)**.** *Consider the exactly balanced two-block SBM on an even number n of vertices, with edge probabilities $a \log n/n$ and $b \log n/n$. If $D_+(a,b) > 3/4$, then the balanced leave-one-out MLEs satisfy the following with probability $1 - o(1)$: there is a common global sign $s \in \{\pm 1\}$ such that, simultaneously for every $v \in [n]$,*

$$\left|\{u \neq v : \widehat{\sigma}_u^{(-v)} \neq s\sigma_u^*\}\right| \leq n^{1/4 + o(1)}.$$

In particular, the number of leave-one-out reference-label errors is $o(n/\log n)$ uniformly over v.

Proof. We use a union-bound proof following in the style of ABH, keeping track of the balanced constraint. Fix v and condition on the true labels. Write $N = n - 1$. The restriction of the truth to $[n] \setminus \{v\}$ belongs to $\mathcal{B}_{-v}$, since the full truth is exactly balanced. For $\tau \in \mathcal{B}_{-v}$, let $L_{-v}(\tau)$ denote the SBM likelihood of the graph on $[n] \setminus \{v\}$ under the labeling τ; that is, the product likelihood over pairs $\{i,j\} \subseteq [n] \setminus \{v\}$ with edge probability $a \log n/n$ when $\tau_i = \tau_j$ and $b \log n/n$ otherwise.

For $\tau \in \mathcal{B}_{-v}$, choose its global sign so that

$$m = m(\tau) := \left|\{u \neq v : \tau_u \neq \sigma_u^*\}\right| \leq N/2.$$

Let S be this disagreement set. A pair $\{i,j\} \subseteq [n] \setminus \{v\}$ has different within/between-community status under τ and under σ^* exactly when one endpoint lies in S and the other lies outside S. Thus the likelihood ratio between τ and the truth only involves

$$M(S) = m(N - m)$$

pairs; all other pair contributions cancel. For each such pair, the two hypotheses swap the Bernoulli parameters $a \log n/n$ and $b \log n/n$. Hence, by the Bhattacharyya bound applied to the likelihood-ratio test between the truth and this fixed alternative,

$$\Pr\{L_{-v}(\tau) \geq L_{-v}(\sigma_{-v}^*)\} \leq \left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)^{M(S)}, \qquad p = \frac{a \log n}{n}, \quad q = \frac{b \log n}{n}.$$

Since

$$\sqrt{pq} + \sqrt{(1-p)(1-q)} = 1 - D_+(a,b)\frac{\log n}{n} + O\left(\frac{\log^2 n}{n^2}\right),$$

we obtain

$$\Pr\{\tau \text{ has likelihood at least that of the truth}\} \leq \exp\left(-(D_+(a,b) - o(1))\frac{m(N-m)}{n}\log n\right). \quad (7)$$

It remains to union bound over balanced alternatives. The number of alternatives with Hamming distance m is at most $\binom{N}{m} \leq \exp(m \log(eN/m))$. For $n^{1/4} \leq m \leq n/\log n$, we have $(N-m)/n = 1 - o(1)$ and $\log(eN/m) \leq (3/4 + o(1)) \log n$, so the total contribution from this range is at most

$$\sum_{m=n^{1/4}}^{n/\log n} \exp\left(-\left(D_+(a,b) - \tfrac{3}{4} - o(1)\right) m \log n\right) = o(n^{-2}).$$

For $n/\log n < m \leq N/2$, we have $(N-m)/n \geq 1/2 - o(1)$ and $\log(eN/m) \leq \log(e \log n) = O(\log\log n)$. Therefore the likelihood exponent, which is of order $m \log n$, dominates the entropy term $m \log(eN/m)$, and this range contributes

$$\sum_{m > n/\log n} \exp\left(-\Omega(m \log n)\right) = o(n^{-2}).$$

Thus, for this fixed v, with probability $1 - o(n^{-2})$, no balanced labeling at distance at least $n^{1/4}$ has likelihood at least that of the truth. A union bound over the n choices of v gives, with probability $1 - o(1)$, that every leave-one-out maximizer is within $n^{1/4}$ of either $\sigma^{*(-v)}$ or $-\sigma^{*(-v)}$.

It remains only to check that the deterministic anchor convention chooses the same sign for all v. Since $|D_n| = n^{3/4} + O(1)$ and the true labels are sampled from the exactly balanced model, hypergeometric anti-concentration gives

$$\left| \sum_{u \in D_n} \sigma^*_u \right| > n^{1/3}$$

with probability $1 - o(1)$. Removing one vertex changes this anchor sum by at most one, and the leave-one-out error bound above contributes at most $n^{1/4+o(1)} = o(n^{1/3})$ errors on D_n. Hence the anchor orientation selects the same global sign for every leave-one-out maximizer. The claimed simultaneous $n^{1/4+o(1)}$ bound follows. □

Lemma 4 (Screening estimate for the concrete first stage)**.** *Consider the oracle-only model with $\alpha = 16$ and $\beta = 1$. Let $\theta_0 := 1 - 2^{-1/3}$, and $w_0 := 1 - \sqrt{1-\theta_0}$. If the first stage queries every vertex exactly once with oracle reliability w_0, then the first-stage transcript admits a 0.88-screening initialization with reference parameter $1/2$, in the sense of* (2)–(3)*.*

Proof. Let $G^{(0)}$ denote the first-stage observed graph, and let $A^{(0)}$ be its adjacency matrix. For notational convenience in the proof, write

$$a := 16\theta_0, \qquad b := \theta_0.$$

We first identify the distribution of the first-stage graph. The underlying graph is

$$G \sim \mathrm{SBM}\left(n, 16\frac{\log n}{n}, \frac{\log n}{n}\right).$$

Every vertex is queried exactly once. Thus an existing edge $\{u, v\}$ is missed by the first stage if it is missed when querying u and also missed when querying v. These two oracle coins are independent, so the miss probability is

$$(1 - w_0)^2 = 1 - \theta_0.$$

Therefore an existing edge is revealed with probability θ_0. Since the oracle never creates false edges, the first-stage observed graph is exactly an SBM with edge probabilities

$$a\frac{\log n}{n} = 16\theta_0\frac{\log n}{n}, \qquad b\frac{\log n}{n} = \theta_0\frac{\log n}{n}.$$

Equivalently,

$$A^{(0)} \sim \mathrm{SBM}\left(n, a\frac{\log n}{n}, b\frac{\log n}{n}\right).$$

The useful point is that this graph is not strong enough for exact recovery, but it is strong enough for rough recovery. Indeed,

$$D_+(a,b) = \frac{1}{2}(\sqrt{a} - \sqrt{b})^2 = \frac{1}{2}\theta_0(\sqrt{16} - \sqrt{1})^2 = \frac{9}{2}(1 - 2^{-1/3}) > 0.92 > \frac{3}{4}.$$

Hence Lemma 9 applies to the first-stage graph. It gives balanced leave-one-out rough labelings $\widehat{\sigma}^{(-v)}$, one for each $v \in [n]$, such that with probability $1 - o(1)$, there is a common global sign $s \in \{\pm 1\}$ satisfying

$$\max_{v \in [n]} \left| \left\{ u \neq v : \widehat{\sigma}^{(-v)}_u \neq s\sigma^*_u \right\} \right| \leq n^{1/4+o(1)} = o(n/\log n).$$

Here $\widehat{\sigma}^{(-v)}$ is computed after deleting all first-stage observations incident to v. This leave-one-out construction is important: it makes the reference labels $\widehat{\sigma}^{(-v)}$ independent of the first-stage observations from v, conditional on the true labels.

For each v, split the other vertices into two deterministic pieces. List $[n] \setminus \{v\}$ in increasing order. Let B_v be the first $\lfloor (n-1)/2 \rfloor$ vertices in this list, and let

$$C_v := ([n] \setminus \{v\}) \setminus B_v.$$

The set C_v will be used only to decide whether v is ambiguous. The set B_v will be reserved for the second-stage likelihood test. This split prevents us from using the same incident observations both to screen v and later to classify v.

Since the true labeling is sampled uniformly from the exactly balanced model, a deterministic set of size $n/2 + O(1)$ contains $n/4 + o(n)$ vertices from each community, with high probability. By hypergeometric concentration and a union bound over all v, we have simultaneously for every v,

$$|B_v \cap \{u : \sigma_u^* = +1\}| = \frac{n}{4} + o(n), \qquad |B_v \cap \{u : \sigma_u^* = -1\}| = \frac{n}{4} + o(n),$$

and also

$$|C_v \cap \{u : \sigma_u^* = +1\}| = \frac{n}{4} + o(n), \qquad |C_v \cap \{u : \sigma_u^* = -1\}| = \frac{n}{4} + o(n).$$

Now define the leave-one-out screening score

$$\Lambda_v^{\mathrm{loo}} := \sum_{u \in C_v} \widehat{\sigma}_u^{(-v)} A_{uv}^{(0)}.$$

This score has a simple interpretation. It compares how many revealed first-stage edges from v go to vertices currently labeled $+1$ and how many go to vertices currently labeled -1. If v is truly in the $+1$ community, then it should have more revealed edges to true $+1$ vertices than to true -1 vertices, so the score should usually be positive. If v is truly in the -1 community, the score should usually be negative. Vertices with scores close to zero are the ambiguous vertices that we will send to the second stage.

Fix the common global sign s from the rough-initialization guarantee. For readability, assume first that $s = +1$ and $\sigma_v^* = +1$. If the reference labels $\widehat{\sigma}^{(-v)}$ were exact on C_v, then Λ_v^{loo} would have the same large-deviation behavior as

$$X_v - Y_v,$$

where

$$X_v \sim \mathrm{Bin}\left(\frac{n}{4} + o(n), a\frac{\log n}{n}\right), \qquad Y_v \sim \mathrm{Bin}\left(\frac{n}{4} + o(n), b\frac{\log n}{n}\right).$$

Here X_v counts revealed edges from v to same-community vertices in C_v, while Y_v counts revealed edges from v to opposite-community vertices in C_v.

The reference labels are not perfectly exact, but the number of reference-label errors in C_v is $o(n/\log n)$, uniformly over v. Conditional on the leave-one-out graph, the number of first-stage revealed edges from v to these mislabeled reference vertices has mean

$$O\left(\frac{\log n}{n}\right) o(n/\log n) = o(1).$$

In the next claim, we show that the contribution of mislabeled reference vertices in C_v to the leave-one-out screening score Λ_v^{loo} is extremely small. More precisely, for every $\varepsilon > 0$, the probability that these mislabeled vertices perturb Λ_v^{loo} by more than $\varepsilon \log n$ is $o(n^{-2})$, uniformly over v.

Claim (Bounding the tail event)**.** *By Chernoff upper-tail bound, for every fixed $\varepsilon > 0$,*

$$\Pr\left(\textit{the mislabeled reference vertices change } \Lambda_v^{\mathrm{loo}} \textit{by more than } \varepsilon \log n\right) = o(n^{-2}),$$

uniformly in v.

Proof. Let $D_v \subseteq C_v$ be the set of mislabeled reference vertices in C_v, and let M_v be the number of first-stage revealed edges from v to vertices in D_v. By the earlier part of the proof, uniformly over v,

$$|D_v| = o(n/\log n).$$

Since each edge $\{u, v\}$ with $u \in D_v$ is revealed in the first stage graph with probability at most $c \log n/n$ for some constant $c > 0$, it follows that M_v is stochastically dominated by

$$\mathrm{Bin}\left(|D_v|, c\frac{\log n}{n}\right),$$

Hence, again uniformly over v,

$$\lambda_v := \mathbb{E}[M_v] = O\left(|D_v|\frac{\log n}{n}\right) = o(1).$$

Now let

$$|\Delta_v| := |\Lambda_v^{\mathrm{loo}} - (X_v - Y_v)|$$

so that Δ_v is the perturbation of the screening score caused by the mislabeled reference vertices in C_v. Each revealed edge to a mislabeled reference vertex changes the score by at most a constant; in fact, under the present definition of Λ_v^{loo}, each such edge changes the score by at most 2. Therefore, for some absolute constant $C > 0$,

$$|\Delta_v| \le CM_v.$$

Consequently, for every fixed $\varepsilon > 0$,

$$\Pr(|\Delta_v| > \varepsilon \log n) \le \Pr\left(M_v > \frac{\varepsilon}{C}\log n\right).$$

Set $t_v = \frac{\epsilon}{C}\log n$ and applying Chernoff upper-tail bound for the binomial random variable gives

$$\Pr(M_v \ge t_v) \le \left(\frac{e\lambda_v}{t_v}\right)^{t_v}$$

Since $\lambda_v = o(1)$ uniformly over v and $t_v = \Theta(\log n)$, we obtain

$$\left(\frac{e\lambda_v}{t_v}\right)^{t_v} = \exp\left(-\Theta(\log n \log\log n)\right) = n^{-\omega(1)}.$$

In particular, this is $o(n^{-2})$ uniformly over v. □

A union bound over v shows that, simultaneously for all vertices, the effect of mislabeled reference vertices is only $o(\log n)$ with probability $1 - o(1)$. Thus the large-deviation exponents of the actual score Λ_v^{loo} are the same as those of $X_v - Y_v$, up to $o(1)$ in the exponent.

Let $I_{a,b}(d)$ be the Cramér rate function for $X_v - Y_v$ at scale $\log n$:

$$I_{a,b}(d) := \sup_{t\in\mathbb{R}}\left\{td - \frac{a}{4}(e^t - 1) - \frac{b}{4}(e^{-t} - 1)\right\}.$$

This rate function describes the probability that $X_v - Y_v$ has an atypically small value:

$$\Pr(X_v - Y_v \le d\log n) = n^{-I_{a,b}(d)+o(1)}$$

in the lower-tail regime relevant here. In particular,

$$I_{a,b}(0) = \frac{1}{4}(\sqrt{a} - \sqrt{b})^2 = \frac{1}{4}\theta_0(\sqrt{16} - \sqrt{1})^2 = \frac{9}{4}(1 - 2^{-1/3}) > 0.46.$$

Choose

$$\delta := 0.32.$$

A direct evaluation of the displayed rate function gives

$$I_{a,b}(\delta) > 0.13, \qquad I_{a,b}(-\delta) > 1.02.$$

Define the observable screened set

$$\mathcal{A} := \left\{v : |\Lambda_v^{\mathrm{loo}}| \le \delta\log n\right\}.$$

For vertices outside the screen, define the preliminary label by

$$\widehat{\sigma}_v^{(0)} := \mathrm{sign}(\Lambda_v^{\mathrm{loo}}), \qquad v \notin \mathcal{A}.$$

Thus $\mathcal{A}$ contains exactly those vertices whose first-stage score has too small a margin to trust.

We now bound the size of $\mathcal{A}$. Suppose $\sigma_v^* = +1$. Since the typical value of Λ_v^{loo} is positive, the event $v \in \mathcal{A}$ is a lower-tail event:

$$\Pr(v \in \mathcal{A}) = \Pr(|\Lambda_v^{\mathrm{loo}}| \le \delta\log n) \le n^{-I_{a,b}(\delta)+o(1)} \le n^{-0.13+o(1)}.$$

The same estimate holds when $\sigma_v^* = -1$, by symmetry. Therefore

$$\mathbb{E}|\mathcal{A}| \le n^{0.87+o(1)}.$$

By Markov's inequality,

$$|\mathcal{A}| \le n^{0.88}$$

with high probability.

Next we show that every vertex outside the screen is labeled correctly. If $\sigma_v^* = +1$, a wrong sign outside the screen requires

$$\Lambda_v^{\mathrm{loo}} < -\delta \log n.$$

The corresponding large-deviation bound gives

$$\Pr(\Lambda_v^{\mathrm{loo}} < -\delta \log n) \le n^{-I_{a,b}(-\delta)+o(1)} \le n^{-1.02+o(1)}.$$

Equivalently, for both possible true labels,

$$\Pr(\sigma_v^* \Lambda_v^{\mathrm{loo}} < -\delta \log n) \le n^{-1.02+o(1)}.$$

A union bound over all vertices gives

$$\Pr\left(\exists v \notin \mathcal{A} : \widehat{\sigma}_v^{(0)} \ne s\sigma_v^*\right) = o(1).$$

Thus, with high probability, every vertex outside $\mathcal{A}$ is correctly labeled, up to the same global sign s.

It remains to check the reference condition for vertices inside the screen. For each $v \in \mathcal{A}$, we use B_v as the second-stage reference set and use $\widehat{\sigma}^{(-v)}$ restricted to B_v as the reference labeling. The rough leave-one-out guarantee gives

$$\left|\left\{u \in B_v : \widehat{\sigma}_u^{(-v)} \ne s\sigma_u^*\right\}\right| = o(n/\log n)$$

simultaneously for every v. The hypergeometric concentration estimates above give

$$|B_v \cap \{u : \sigma_u^* = +1\}| = \frac{n}{4} + o(n), \qquad |B_v \cap \{u : \sigma_u^* = -1\}| = \frac{n}{4} + o(n).$$

Therefore (3) holds with reference parameter

$$\zeta = \frac{1}{2}.$$

Finally, the measurability and independence requirements in the definition of screened initialization hold by construction. The event $\{v \in \mathcal{A}\}$ uses first-stage observations incident to v only through vertices in C_v. The second-stage likelihood test will use fresh oracle queries from v to vertices in B_v. Since $B_v \cap C_v = \varnothing$, screening does not condition on the first-stage observations from v to the reference set used later for refinement. Also, the complement of $B_v \cup \{v\}$ has linear size, so the exact balance constraint does not determine σ_v^* from the labels on B_v alone.

Combining the preceding estimates, with probability $1 - o(1)$,

$$|\mathcal{A}| \le n^{0.88}, \qquad \widehat{\sigma}_u^{(0)} = s\sigma_u^* \quad \text{for all } u \notin \mathcal{A},$$

and for every $v \in \mathcal{A}$, the leave-one-out reference set B_v satisfies (3) with $\zeta = 1/2$. Hence the first-stage graph admits a 0.88-screening initialization with reference parameter $1/2$. □

Theorem 5. *In Model 1 every oracle-only exact recovery strategy requires $k = \Omega(n)$ queries.*

Proof. We prove the contrapositive. Fix an arbitrary adaptive oracle-only strategy with query budget $k = o(n)$. We show that exact recovery fails with probability tending to one.

The main intuition is the following. A query at a vertex can only reveal neighbors of that queried vertex. If $k = o(n)$, then only $o(n)$ vertices can be queried at all. Moreover, since the graph has logarithmic maximum degree with high probability, the total number of vertex appearances in all oracle responses is only $o(n \log n)$. Hence most vertices are both unqueried and appear only rarely in the entire transcript. Swapping the labels of two such vertices, one from each community, changes

the likelihood of the transcript by only a small polynomial factor. Since there are linearly many such balanced swaps, the true labeling cannot dominate the posterior.

Let T denote the full oracle transcript, including the algorithm's internal randomness. Define

$$Q(T) := \{v : \text{vertex } v \text{ is queried at least once}\}.$$

Since the strategy makes at most k queries,

$$|Q(T)| \le k = o(n).$$

We next control how often vertices appear in oracle responses. For each $v \in [n]$, let $A_v(T)$ be the number of times v is returned as a neighbor in the oracle transcript, counted with multiplicity over repeated queries. Since the oracle never returns non-neighbors, every response is a subset of the true neighborhood of the queried vertex.

With high probability, the underlying SBM has maximum degree at most $C_{\deg} \log n$, for a constant $C_{\deg} = C_{\deg}(\alpha, \beta)$. On this event, every oracle response has size at most $C_{\deg} \log n$, and hence

$$\sum_{v=1}^{n} A_v(T) \le C_{\deg} k \log n = o(n \log n).$$

Fix a small constant $\varepsilon > 0$, to be chosen later, and define the set of vertices that appear too often in the transcript:

$$H(T) := \{v : A_v(T) > \varepsilon \log n\}.$$

Then

$$|H(T)| \le \frac{1}{\varepsilon \log n} \sum_{v=1}^{n} A_v(T) \le \frac{C_{\deg}}{\varepsilon} k = o(n).$$

Now define the good set

$$G(T) := [n] \setminus \big(Q(T) \cup H(T)\big).$$

Thus a vertex in $G(T)$ is never queried and is returned as a neighbor at most $\varepsilon \log n$ times in the entire transcript. Since both $Q(T)$ and $H(T)$ have size $o(n)$, we have

$$|G(T)| = n - o(n)$$

with high probability. Because the true labeling is exactly balanced, removing only $o(n)$ vertices leaves linearly many vertices in each community:

$$|G(T) \cap V^+| = \frac{n}{2} - o(n), \qquad |G(T) \cap V^-| = \frac{n}{2} - o(n).$$

We now fix two good vertices of opposite labels,

$$i \in G(T) \cap V^+, \qquad j \in G(T) \cap V^-.$$

Let $\sigma^{(i,j)}$ be the balanced labeling obtained by swapping the labels of i and j:

$$\sigma_i^{(i,j)} = -1, \qquad \sigma_j^{(i,j)} = +1, \qquad \sigma_u^{(i,j)} = \sigma_u^\star \quad \text{for } u \notin \{i, j\}.$$

This is the natural local obstruction in the balanced SBM: a one-vertex flip is not an admissible balanced labeling, but swapping one + vertex with one − vertex is.

Define the likelihood ratio

$$L_{ij}(T) := \frac{\mathbb{P}_{\sigma^{(i,j)}}(T)}{\mathbb{P}_{\sigma^\star}(T)}.$$

We now lower-bound $L_{ij}(T)$. The important point is that we do this directly on the original adaptive transcript. We do not compare the original run to a counterfactual run. The algorithm's future query choices may depend on past oracle responses, but once a transcript T is fixed, the query choices are deterministic functions of the past transcript and the internal randomness. Thus the query choices themselves contribute no likelihood factor. Only the probabilities of the observed oracle responses matter.

Under the swap $\sigma^\star \mapsto \sigma^{(i,j)}$, all pairs not incident to i or j have the same within/between-community status as before. The pair $\{i,j\}$ is also between-community under both labelings, so it contributes no likelihood difference. Therefore the only possible likelihood differences come from pairs $\{u,i\}$ and $\{u,j\}$ with $u \notin \{i,j\}$.

Since $i,j \in G(T)$, neither i nor j is ever queried. Thus the only way the transcript can see i or j is if some queried vertex u returns i or j as a neighbor.

Consider one such pair $\{u,v\}$, where $v \in \{i,j\}$ and u is a queried vertex. Suppose u is queried r times in the realized transcript. If v is returned at least once in those r queries, then the likelihood ratio contribution of the pair $\{u,v\}$ is bounded below by a positive constant depending only on α,β. Indeed, the two possible edge probabilities are p and q, so the return event changes likelihood by at worst a factor comparable to p/q or q/p.

If v is never returned in those r queries, then the probability of this non-return event is

$$1-\rho p \qquad \text{or} \qquad 1-\rho q, \qquad \rho := 1-(1-w)^r \in [0,1],$$

depending on whether the pair is within-community or between-community. Since $p,q = O(\log n/n)$, uniformly over $\rho \in [0,1]$,

$$\left|\log \frac{1-\rho p}{1-\rho q}\right| = O\Big(\frac{\log n}{n}\Big).$$

Thus a non-return contributes only $O(\log n/n)$ log-likelihood loss.

Combining these two cases, there exists a constant $C_0 = C_0(\alpha,\beta,w)$ such that, for the realized transcript,

$$\log L_{ij}(T) \geq -C_0\big(A_i(T)+A_j(T)\big) - C_0|Q(T)|\frac{\log n}{n}.$$

The first term accounts for actual returns of i or j, each of which can cost only a constant factor in likelihood. The second term accounts for all queried vertices that did not return i or j, each of which costs only $O(\log n/n)$.

Because $i,j \in G(T)$,

$$A_i(T) \leq \varepsilon \log n, \qquad A_j(T) \leq \varepsilon \log n.$$

Also $|Q(T)| \leq k = o(n)$, so

$$|Q(T)|\frac{\log n}{n} = o(\log n).$$

Therefore

$$\log L_{ij}(T) \geq -2C_0\varepsilon \log n - o(\log n),$$

or equivalently

$$L_{ij}(T) \geq n^{-2C_0\varepsilon - o(1)}.$$

Choose $\varepsilon > 0$ small enough that

$$2C_0\varepsilon < \frac{1}{2}.$$

Then every balanced swap between two good vertices of opposite communities has likelihood at least

$$L_{ij}(T) \geq n^{-1/2-o(1)}.$$

On the high-probability event above, choose an arbitrary matching

$$\mathcal{M}(T) \subset \big(G(T)\cap V^+\big) \times \big(G(T)\cap V^-\big)$$

of size

$$|\mathcal{M}(T)| = \frac{n}{2} - o(n).$$

For every $(i,j) \in \mathcal{M}(T)$, the labeling $\sigma^{(i,j)}$ is balanced and distinct from $\sigma^\star$, and satisfies

$$\frac{\mathbb{P}_{\sigma^{(i,j)}}(T)}{\mathbb{P}_{\sigma^\star}(T)} = L_{ij}(T) \geq n^{-1/2-o(1)}.$$

Since the prior over balanced labelings is uniform, the posterior mass assigned to the true labeling is

$$\begin{aligned}\Pi(\sigma^\star \mid T) &= \frac{\mathbb{P}_{\sigma^\star}(T)}{\sum_{\tau \text{ balanced}} \mathbb{P}_\tau(T)} \\ &\le \frac{\mathbb{P}_{\sigma^\star}(T)}{\mathbb{P}_{\sigma^\star}(T) + \sum_{(i,j)\in\mathcal{M}(T)} \mathbb{P}_{\sigma^{(i,j)}}(T)} \\ &= \frac{1}{1 + \sum_{(i,j)\in\mathcal{M}(T)} L_{ij}(T)} \\ &\le \frac{1}{1 + \left(\frac{n}{2} - o(n)\right) n^{-1/2-o(1)}} \\ &= o(1).\end{aligned}$$

Thus, with probability $1 - o(1)$, the true balanced labeling has vanishing posterior mass. To translate this into an impossibility statement for every estimator, recall the standard Bayes-risk identity. Conditional on the transcript T, the best possible estimator is the MAP estimator, whose conditional success probability is

$$\max_{\tau \text{ balanced}} \Pi(\tau \mid T).$$

Moreover,

$$\mathbb{E}\big[\Pi(\sigma^\star \mid T) \mid T\big] = \sum_{\tau \text{ balanced}} \Pi(\tau \mid T)^2.$$

Since $\Pi(\sigma^\star \mid T) = o(1)$ with high probability and is always at most one, we have

$$\mathbb{E} \sum_{\tau \text{ balanced}} \Pi(\tau \mid T)^2 = o(1).$$

Therefore

$$\mathbb{E} \max_{\tau \text{ balanced}} \Pi(\tau \mid T) \le \mathbb{E}\left[\left(\sum_{\tau \text{ balanced}} \Pi(\tau \mid T)^2\right)^{1/2}\right] = o(1).$$

Hence even the MAP estimator succeeds with probability $o(1)$, and no estimator can achieve exact recovery with probability tending to one.

We conclude that every oracle-only exact recovery strategy must use $k = \Omega(n)$ queries. □

C Omitted proofs from Section 4

Lemma 6 (Uniform querying fails)**.** *Consider Model 2 with subsampling rate $s \in (0,1)$, oracle reliability $w \in (0,1]$, and query budget $k = o(n)$. Suppose the learner uses a data-independent balanced uniform querying strategy: the queried set is fixed in advance, or chosen independently of G_{sub} and the oracle responses. Then, in the sublinear-budget regime, uniform querying does not improve the exact recovery threshold beyond that of G_{sub} alone. In particular, exact recovery under uniform querying is impossible w.h.p. whenever $s\,\Delta^2 < 2$.*

Proof. The proof has one simple idea. When $k = o(n)$, a balanced uniform querying rule touches only $o(n)$ vertices. Hence almost every vertex is unqueried. For two unqueried vertices of opposite communities, the extra oracle observations affect only $o(n)$ possible comparison vertices, while the $n - o(n)$ remaining comparisons are exactly those from the subsampled graph. Thus the local exact recovery exponent for such vertices is the same as for G_{sub} alone.

Write

$$p = \alpha\frac{\log n}{n}, \qquad q = \beta\frac{\log n}{n}, \qquad I = (\sqrt{\alpha} - \sqrt{\beta})^2.$$

For all sufficiently large n, since $k = o(n)$, we have $k < n$. A balanced uniform non-adaptive strategy queries k vertices once and leaves the remaining $n - k$ vertices unqueried. Let

$$Q := \{i \in [n] : R_i = 1\}, \qquad U := [n] \setminus Q.$$

Thus $|Q| = k = o(n)$ and $|U| = n - o(n)$. We condition throughout on the queried set Q. Since the querying rule is data-independent, Q is fixed in advance, or randomized independently of both the graph and the labels.

For an existing edge $\{i,j\}$, the edge is observed in the final data if it is retained in G_{sub}, or if it is revealed by an oracle query to one of its endpoints. Conditional on the query counts,

$$\rho_{ij} := \Pr(\{i,j\} \text{ is observed} \mid \{i,j\} \in E) = 1 - (1-s)(1-w)^{R_i+R_j}.$$

In particular, if $i,j \in U$, then $\rho_{ij} = s$, whereas if $i \in U$ and $j \in Q$, then

$$\rho_{ij} = 1 - (1-s)(1-w) = s + (1-s)w.$$

Therefore, for every unqueried vertex $i \in U$,

$$\sum_{j \neq i} \rho_{ij} = \sum_{j \in U \setminus \{i\}} \rho_{ij} + \sum_{j \in Q} \rho_{ij} \tag{8}$$

$$= s(|U| - 1) + \big(s + (1-s)w\big)|Q| \tag{9}$$

$$= s(n - k - 1) + \big(s + (1-s)w\big)k \tag{10}$$

$$= sn + o(n). \tag{11}$$

Thus an unqueried vertex has total exposure $sn + o(n)$, and the extra oracle exposure coming from queried vertices contributes only $o(n)$ possible comparison vertices.

We now prove impossibility using balanced two-vertex swaps. Let

$$U_+ := U \cap \{i : \sigma_i^\star = +1\}, \qquad U_- := U \cap \{i : \sigma_i^\star = -1\}.$$

Because Q is independent of the labels and $|Q| = o(n)$, the exact balance of the true labeling implies, with probability $1 - o(1)$,

$$|U_+| = \frac{n}{2} - o(n), \qquad |U_-| = \frac{n}{2} - o(n). \tag{12}$$

We work on this event.

Let Y denote the final observation, consisting of G_{sub} and the oracle-revealed edges. Conditional on Q and on a candidate labeling σ, the final observed edge indicators are independent over unordered pairs, and

$$Y_{ab} \sim \begin{cases} \mathrm{Bernoulli}(\rho_{ab}p), & \sigma_a = \sigma_b, \\ \mathrm{Bernoulli}(\rho_{ab}q), & \sigma_a \neq \sigma_b. \end{cases}$$

For $i \in U_+$, define the signed one-vertex log-likelihood ratio in favor of assigning i the wrong label:

$$X_i := \sum_{\ell \neq i} \sigma^\star_\ell \left[Y_{i\ell} \log \frac{q}{p} + (1 - Y_{i\ell}) \log \frac{1 - \rho_{i\ell} q}{1 - \rho_{i\ell} p} \right].$$

Indeed, if $\sigma^\star_\ell = +1$, then under the true label $\sigma_i = +1$ the pair $\{i, \ell\}$ is within-community, while under the wrong label $\sigma_i = -1$ it is between-community. The corresponding single-pair log-likelihood ratio is

$$Y_{i\ell} \log \frac{\rho_{i\ell} q}{\rho_{i\ell} p} + (1 - Y_{i\ell}) \log \frac{1 - \rho_{i\ell} q}{1 - \rho_{i\ell} p} = Y_{i\ell} \log \frac{q}{p} + (1 - Y_{i\ell}) \log \frac{1 - \rho_{i\ell} q}{1 - \rho_{i\ell} p}.$$

If instead $\sigma^\star_\ell = -1$, the roles of p and q are reversed, and the log-likelihood ratio changes sign. This is exactly captured by the factor $\sigma^\star_\ell$. Hence X_i is the log-likelihood ratio between assigning i the wrong label -1 and assigning it the true label $+1$, with all other labels fixed at their true values.

For $j \in U_-$, define analogously

$$X_j := -\sum_{\ell \neq j} \sigma^\star_\ell \left[Y_{j\ell} \log \frac{q}{p} + (1 - Y_{j\ell}) \log \frac{1 - \rho_{j\ell} q}{1 - \rho_{j\ell} p} \right],$$

so that X_j is the log-likelihood ratio in favor of assigning j the wrong label $+1$.

We first estimate the probability that one such unqueried vertex has atypically strong evidence for the wrong label. Fix $i \in U_+$. Conditional on the labels, X_i is a sum of independent terms. The terms corresponding to queried vertices contribute negligibly to the logarithmic large-deviation scale. More precisely, for any fixed θ in a neighborhood of 0, each $\ell \in Q$ contributes

$$\log \mathbb{E} \exp \left\{ \theta \, \sigma^\star_\ell \left[Y_{i\ell} \log \frac{q}{p} + (1 - Y_{i\ell}) \log \frac{1 - \rho_{i\ell} q}{1 - \rho_{i\ell} p} \right] \right\} = O\left(\frac{\log n}{n} \right),$$

uniformly in $\rho_{i\ell} \in [0, 1]$. This is because $\Pr(Y_{i\ell} = 1) = O(\log n / n)$, the edge-present log-ratio is $O(1)$, and the edge-absent log-ratio is $O(\log n / n)$. Since $|Q| = k = o(n)$, the total contribution of queried vertices to $\log \mathbb{E} e^{\theta X_i}$ is $o(\log n)$. Hence the queried vertices contribute $o(1)$ to the normalized cumulant generating function at scale $\log n$. Equation (11) also shows that their first-moment contribution is $o(n)$, so they do not shift the typical value of X_i at the relevant scale.

The remaining $n - o(n)$ comparison vertices are unqueried and therefore have retention probability exactly s. Thus the usual one-vertex Cramér calculation for the logarithmic-degree SBM applies with effective parameters $s\alpha$ and $s\beta$. Consequently,

$$\Pr_{\sigma^\star}(X_i \geq 0) = n^{-sI/2 + o(1)}.$$

More generally, for fixed t in a neighborhood of 0,

$$\Pr_{\sigma^\star}(X_i \geq t \log n) = n^{-J(t) + o(1)},$$

where J is the corresponding Cramér rate function and

$$J(0) = \frac{sI}{2}.$$

The matching lower bound is a standard strong large-deviation estimate for triangular arrays. In particular, it follows from the Bahadur–Rao-type theorem of Chaganty and Sethuraman [9], applied to the log-moment generating functions of the independent Bernoulli summands. In the logarithmic-degree SBM, the same logarithmic asymptotic is computed explicitly in the converse proof of Abbe, Bandeira, and Hall [4]; see also Mossel, Neeman, and Sly [29] for the corresponding exact recovery converse.

Since $sI < 2$, we have $J(0) < 1$. By continuity of J, choose a sufficiently small fixed $\delta > 0$ such that

$$\mu := J(\delta) < 1.$$

Then, uniformly for $i \in U_+$,

$$\Pr_{\sigma^\star}(X_i \geq \delta \log n) = n^{-\mu + o(1)}. \tag{13}$$

The same estimate holds symmetrically for $j \in U_-$.

Define

$$Z_+ := \big|\{i \in U_+ : X_i \geq \delta \log n\}\big|, \qquad Z_- := \big|\{j \in U_- : X_j \geq \delta \log n\}\big|.$$

By (12) and (13),

$$\mathbb{E}_{\sigma^\star} Z_+ = n^{1-\mu+o(1)} \to \infty, \qquad \mathbb{E}_{\sigma^\star} Z_- = n^{1-\mu+o(1)} \to \infty.$$

We next show concentration. For distinct $i, i' \in U_+$, the score X_i depends on the edge variables $\{Y_{i\ell} : \ell \neq i\}$, while $X_{i'}$ depends on $\{Y_{i'\ell} : \ell \neq i'\}$. These two collections share exactly one unordered edge variable, namely $Y_{ii'}$. All other edge observations entering the two sums are disjoint and hence independent conditional on the labels.

Write

$$X_i = S_i + T_{ii'}, \qquad X_{i'} = S_{i'} + T_{i'i},$$

where $T_{ii'}$ and $T_{i'i}$ are the contributions of the shared edge $\{i, i'\}$. Conditional on $Y_{ii'}$, the variables S_i and $S_{i'}$ are independent. Moreover,

$$\Pr_{\sigma^\star}(Y_{ii'} = 1) = O\left(\frac{\log n}{n}\right).$$

If $Y_{ii'} = 0$, then $T_{ii'}, T_{i'i} = O(\log n/n)$, so the threshold $\delta \log n$ is shifted only by $o(1)$. The sharp large-deviation estimate above is stable under such an $o(1)$ shift, giving the same tail probability up to a factor $1 + o(1)$. If $Y_{ii'} = 1$, then $T_{ii'}, T_{i'i} = O(1)$; this changes the logarithmic exponent only by $o(1)$, while the event $Y_{ii'} = 1$ itself has probability $O(\log n/n)$, making this contribution negligible compared to the product of the two one-vertex tail probabilities. Therefore, uniformly over distinct $i, i' \in U_+$,

$$\Pr_{\sigma^\star}(X_i \geq \delta \log n,\ X_{i'} \geq \delta \log n) = (1 + o(1)) \Pr_{\sigma^\star}(X_i \geq \delta \log n) \Pr_{\sigma^\star}(X_{i'} \geq \delta \log n).$$

It follows that

$$\mathrm{Var}(Z_+) = o((\mathbb{E} Z_+)^2).$$

Hence, by Chebyshev's inequality, $Z_+ \geq 1$ with probability $1 - o(1)$. The same argument gives $Z_- \geq 1$ with probability $1 - o(1)$. By a union bound, with probability $1 - o(1)$ both events occur.

On the event $\{Z_+ \geq 1,\ Z_- \geq 1\}$, choose $i \in U_+$ and $j \in U_-$ such that

$$X_i \geq \delta \log n, \qquad X_j \geq \delta \log n.$$

Let $\sigma^{(i,j)}$ be the balanced labeling obtained from $\sigma^\star$ by swapping the labels of i and j:

$$\sigma_i^{(i,j)} = -1, \qquad \sigma_j^{(i,j)} = +1, \qquad \sigma_\ell^{(i,j)} = \sigma_\ell^\star \quad \text{for all } \ell \notin \{i, j\}.$$

Let

$$L_{ij}(Y) := \frac{\mathbb{P}_{\sigma^{(i,j)}}(Y)}{\mathbb{P}_{\sigma^\star}(Y)}$$

be the likelihood ratio between the swapped labeling and the true labeling.

Pairs $\{u, v\}$ with $u, v \notin \{i, j\}$ have identical distributions under $\sigma^\star$ and $\sigma^{(i,j)}$, so they contribute 0 to $\log L_{ij}(Y)$. Pairs $\{i, \ell\}$ with $\ell \notin \{i, j\}$ contribute the corresponding terms in X_i, and pairs $\{j, \ell\}$ with $\ell \notin \{i, j\}$ contribute the corresponding terms in X_j. The pair $\{i, j\}$ is between-community under both $\sigma^\star$ and $\sigma^{(i,j)}$, so it contributes 0 to the likelihood ratio of the balanced swap.

The only bookkeeping issue is that the one-vertex scores X_i and X_j include the terms corresponding to the pair $\{i, j\}$, even though this pair does not contribute to the balanced-swap likelihood ratio. Therefore

$$\log L_{ij}(Y) = X_i + X_j - \big[\text{the } \ell = j \text{ term in } X_i\big] - \big[\text{the } \ell = i \text{ term in } X_j\big].$$

Each of the two subtracted terms is $O(1)$ if $Y_{ij} = 1$, and $O(\log n/n)$ if $Y_{ij} = 0$. Hence, uniformly in Y,

$$\log L_{ij}(Y) = X_i + X_j + O(1).$$

Consequently,

$$\log L_{ij}(Y) \geq 2\delta \log n - O(1) > 0$$

for all sufficiently large n. Thus

$$\mathbb{P}_{\sigma^{(i,j)}}(Y) > \mathbb{P}_{\sigma^\star}(Y).$$

Therefore, with probability $1 - o(1)$, there exists a balanced labeling $\sigma^{(i,j)} \neq \sigma^\star$ whose likelihood is strictly larger than the likelihood of the true labeling. Since the prior over balanced labelings is uniform, the maximum-a-posteriori estimator cannot return $\sigma^\star$ on this event. Hence exact recovery is impossible with probability tending to one.

The argument was carried out after conditioning on the queried set Q and on the high-probability event (12). Removing this conditioning changes the failure probability by only $o(1)$. The use of data-independence of Q is essential in two places: first, it allows us to condition on Q and retain a product measure over the edge observations with deterministic exposures ρ_{ab}; second, it gives (12), since Q is independent of the true labels. Thus the conclusion applies to any non-adaptive allocation that is fixed in advance or randomized independently of $(G_{\mathrm{sub}}, \sigma^\star)$. It does not apply to allocations that depend on G_{sub}.

We conclude that exact recovery under balanced uniform non-adaptive querying is impossible with high probability whenever

$$s(\sqrt{\alpha} - \sqrt{\beta})^2 = sI < 2.$$

This proves the lemma. □

Lemma 7. *Consider Model 2, a subsampling rate $s \in (0,1)$, and oracle reliability $w \in (0,1]$. There exist constants (α, β, s, w) and a budget $k = o(n)$ such that a two-stage strategy that uses G_{sub} to* target *queries succeeds at exact recovery w.h.p.*

Proof. All label recovery statements are understood up to the unavoidable global sign flip. We prove the lemma by giving one concrete choice of parameters and one concrete two-stage strategy.

Step 1: Choose parameters. Set

$$\alpha = 16, \qquad \beta = 1, \qquad w = \frac{1}{2}, \qquad (\sqrt{\alpha} - \sqrt{\beta})^2 = 9.$$

Choose a sufficiently small fixed $\epsilon > 0$ and set

$$s = \frac{2(1-\epsilon)}{9}.$$

Then the subsampled graph G_{sub} has effective exact recovery signal

$$s(\sqrt{\alpha} - \sqrt{\beta})^2 = 2(1-\epsilon) < 2.$$

Thus G_{sub} alone is below the sharp Abbe–Bandeira–Hall threshold for exact recovery in the logarithmic-degree SBM. This is useful context: the subsampled graph alone cannot exactly recover, so the oracle queries must be targeted to repair the remaining ambiguous vertices.

Let

$$c := \frac{(\sqrt{\alpha} - \sqrt{\beta})^2}{2} = \frac{9}{2}.$$

Then

$$cs = 1 - \epsilon.$$

Step 2: A leave-one-out near-exact initial estimator. We use the following standard near-exact recovery fact for the logarithmic-degree SBM. There exist estimators $\widehat{\sigma}^{(0,-v)}$, one for each $v \in [n]$, such that $\widehat{\sigma}^{(0,-v)}$ is computed from G_{sub} after deleting all edges incident to v, and such that, after orienting all estimators by a common deterministic anchor convention,

$$\max_{v \in [n]} \mathrm{Ham}\big(\widehat{\sigma}^{(0,-v)}, \sigma^\star\big) \le n^{1-cs+o(1)} = n^{\epsilon+o(1)} \tag{14}$$

with high probability.

Intuitively, G_{sub} is below the exact recovery threshold but still contains enough information to recover all but a subpolynomially small fraction of the labels. The leave-one-out construction is used only to ensure independence: $\widehat{\sigma}^{(0,-v)}$ is computed without looking at any subsampled edge incident to v, so it is independent of those incident edges conditional on the true labels.

Step 3: Split the neighbors of each vertex into a screening set and a refinement set. We now separate the information used to decide whether v is ambiguous from the information used later to refine v by oracle queries. This separation is the key technical point.

Set

$$b_0 := \frac{7}{20}, \qquad c_0 := \frac{13}{20}.$$

For every v, choose deterministic disjoint sets $B_v, C_v \subset [n] \setminus \{v\}$ satisfying

$$|B_v| = b_0 n + o(n), \qquad |C_v| = c_0 n + o(n).$$

For example, one can choose these sets by a fixed cyclic rule independent of the graph and of the labels. The set C_v will be used only to screen v, and the set B_v will be used only for the second-stage oracle refinement.

Because the true labels are balanced and the sets B_v, C_v are fixed independently of the labels, hypergeometric concentration and a union bound over v imply that, with high probability, simultaneously for all v,

$$|B_v \cap \{u : \sigma_u^\star = +1\}| = \frac{b_0 n}{2} + o(n), \qquad |B_v \cap \{u : \sigma_u^\star = -1\}| = \frac{b_0 n}{2} + o(n),$$

and similarly

$$|C_v \cap \{u : \sigma_u^\star = +1\}| = \frac{c_0 n}{2} + o(n), \qquad |C_v \cap \{u : \sigma_u^\star = -1\}| = \frac{c_0 n}{2} + o(n).$$

Step 4: Define the screening score. Write

$$p = \alpha \frac{\log n}{n}, \qquad q = \beta \frac{\log n}{n}.$$

For each vertex v, define the leave-one-out screening score

$$\Lambda_v^{\mathrm{loo}} := \sum_{u \in C_v} \widehat{\sigma}_u^{(0,-v)} \left[A_{uv}^{\mathrm{sub}} \log \frac{p}{q} + (1 - A_{uv}^{\mathrm{sub}}) \log \frac{1 - sp}{1 - sq} \right].$$

This is the log-likelihood evidence, computed from the subsampled graph, for whether v has label $+1$ or -1, using the leave-one-out estimates of the labels of vertices in C_v as references.

We screen vertices whose evidence is not strong. Let

$$\tau := \frac{21}{20}$$

and define

$$\mathcal{A} := \left\{ v : |\Lambda_v^{\mathrm{loo}}| \le \tau \log n \right\}.$$

The intended interpretation is simple: vertices outside $\mathcal{A}$ have a large positive or negative likelihood score and will be trusted; vertices in $\mathcal{A}$ are ambiguous and will be queried.

Step 5: Analyze the ideal genie-aided screening score. First replace the estimated reference labels $\widehat{\sigma}_u^{(0,-v)}$ by the true labels $\sigma_u^\star$. Define

$$\Lambda_v^\star := \sum_{u \in C_v} \sigma_u^\star \left[A_{uv}^{\mathrm{sub}} \log \frac{p}{q} + (1 - A_{uv}^{\mathrm{sub}}) \log \frac{1 - sp}{1 - sq} \right].$$

We compute the large-deviation rate of $\Lambda_v^\star$ conditional on $\sigma_v^\star = +1$.

For a reference vertex u with $\sigma_u^\star = +1$, the subsampled edge indicator satisfies

$$A_{uv}^{\mathrm{sub}} \sim \mathrm{Bernoulli}(sp).$$

For such a vertex, the single-edge log-likelihood contribution is

$$X_p = A_{uv}^{\mathrm{sub}} \log \frac{p}{q} + (1 - A_{uv}^{\mathrm{sub}}) \log \frac{1 - sp}{1 - sq}.$$

For fixed θ, using $p = \alpha \log n/n$ and $q = \beta \log n/n$, we have

$$\log \mathbb{E} e^{\theta X_p} = \frac{s \log n}{n} \left[\alpha \left(\left(\frac{\alpha}{\beta} \right)^{\theta} - 1 - \theta \right) + \theta\beta \right] + o\left(\frac{\log n}{n} \right).$$

For a reference vertex u with $\sigma_u^\star = -1$, we have $A_{uv}^{\mathrm{sub}} \sim \mathrm{Bernoulli}(sq)$ and the score contributes the negative of the same log-likelihood expression. Hence

$$\log \mathbb{E} e^{\theta X_q} = \frac{s \log n}{n} \left[\beta \left(\left(\frac{\beta}{\alpha} \right)^{\theta} - 1 - \theta \right) + \theta\alpha \right] + o\left(\frac{\log n}{n} \right).$$

The linear terms in θ cancel when we sum over asymptotically equal numbers of $+1$ and -1 reference labels in C_v. Since $|C_v| = c_0 n + o(n)$ and C_v contains asymptotically half of each community, the normalized cumulant generating function of $\Lambda_v^\star$ is

$$\psi_{c_0}(\theta) := \lim_{n\to\infty} \frac{1}{\log n} \log \mathbb{E}\left[e^{\theta \Lambda_v^\star} \,\middle|\, \sigma_v^\star = +1 \right]$$

with

$$\psi_{c_0}(\theta) = c_0 \frac{s}{2} \left[\alpha \left(\left(\frac{\alpha}{\beta} \right)^{\theta} - 1 \right) + \beta \left(\left(\frac{\beta}{\alpha} \right)^{\theta} - 1 \right) \right].$$

Let

$$I_{c_0}(t) := \sup_{\theta \in \mathbb{R}} \{\theta t - \psi_{c_0}(\theta)\}$$

be the associated rate function. The Chernoff bound gives, for fixed t below the typical value of $\Lambda_v^\star / \log n$,

$$\Pr\left(\Lambda_v^\star \le t \log n \,|\, \sigma_v^\star = +1 \right) \le n^{-I_{c_0}(t) + o(1)}.$$

We now evaluate this rate function for our constants. At $\epsilon = 0$, $s = 2/9$, and therefore

$$\psi_{c_0}(\theta) = \frac{13}{180} \left[16(16^\theta - 1) + 16^{-\theta} - 1 \right].$$

The optimizer in

$$I_{c_0}(t) = \sup_{\theta \in \mathbb{R}} \{\theta t - \psi_{c_0}(\theta)\}$$

is the unique solution of $\psi'_{c_0}(\theta) = t$. Since

$$\psi'_{c_0}(\theta) = \frac{13}{180} \log 16 \left[16 \cdot 16^\theta - 16^{-\theta} \right],$$

this is a one-dimensional equation. Direct evaluation at $t = \tau = 21/20$ and $t = -\tau = -21/20$ gives

$$I_{c_0}\left(\frac{21}{20} \right) \approx 0.24515, \qquad I_{c_0}\left(-\frac{21}{20} \right) \approx 1.29515.$$

Hence, for $\epsilon > 0$ sufficiently small, by continuity of the rate function in $s = 2(1-\epsilon)/9$,

$$I_{c_0}(\tau) > 0.24, \qquad I_{c_0}(-\tau) > 1.25. \tag{15}$$

Step 6: Replace the genie labels by the leave-one-out labels. We now compare Λ_v^{loo} with $\Lambda_v^\star$. By (14), uniformly over v, the number of reference labels in C_v on which $\widehat{\sigma}^{(0,-v)}$ disagrees with $\sigma^\star$ is at most $n^{\epsilon + o(1)}$.

The effect of these wrong reference labels is negligible on the $\log n$ scale. Indeed, for any fixed θ, changing the sign of a single reference label changes the corresponding single-vertex log-moment generating function by at most $O(\log n/n)$, because the edge itself appears with probability only $O(\log n/n)$. Therefore the total change in the log-moment generating function caused by all wrong reference labels is at most

$$n^{\epsilon + o(1)} O\left(\frac{\log n}{n} \right) = o(1).$$

In particular, the large-deviation exponents for Λ_v^{loo} are the same as those for the genie-aided score $\Lambda_v^\star$, up to $o(1)$.

Combining this perturbation bound with (15), we get

$$\Pr\left(|\Lambda_v^{\mathrm{loo}}| \le \tau \log n\right) \le n^{-0.24+o(1)}.$$

Indeed, conditional on $\sigma_v^\star = +1$, the event $|\Lambda_v^{\mathrm{loo}}| \le \tau \log n$ is contained in the lower-tail event $\Lambda_v^{\mathrm{loo}} \le \tau \log n$, and the case $\sigma_v^\star = -1$ is symmetric.

Similarly,

$$\Pr\left(\Lambda_v^{\mathrm{loo}} \le -\tau \log n \,\middle|\, \sigma_v^\star = +1\right) \le n^{-1.25+o(1)},$$

and, by symmetry,

$$\Pr\left(\Lambda_v^{\mathrm{loo}} \ge \tau \log n \,\middle|\, \sigma_v^\star = -1\right) \le n^{-1.25+o(1)}.$$

Step 7: The screen is small and all unscreened vertices are correct. Since

$$\mathcal{A} = \{v : |\Lambda_v^{\mathrm{loo}}| \le \tau \log n\},$$

the previous bound gives

$$\mathbb{E}|\mathcal{A}| \le n^{0.76+o(1)}.$$

By Markov's inequality,

$$|\mathcal{A}| \le n^{0.77} \tag{16}$$

with high probability.

Next consider a vertex outside the screen. If $\sigma_v^\star = +1$ and $v \notin \mathcal{A}$ but the score has the wrong sign, then necessarily

$$\Lambda_v^{\mathrm{loo}} < -\tau \log n.$$

This event has probability at most $n^{-1.25+o(1)}$. The same argument applies symmetrically when $\sigma_v^\star = -1$. A union bound over all vertices gives

$$\Pr\left(\exists v \notin \mathcal{A} : \mathrm{sign}(\Lambda_v^{\mathrm{loo}}) \ne \sigma_v^\star\right) \le n \cdot n^{-1.25+o(1)} = n^{-0.25+o(1)} = o(1).$$

Thus, with high probability, every vertex outside $\mathcal{A}$ is already correctly labeled by $\mathrm{sign}(\Lambda_v^{\mathrm{loo}})$.

Step 8: Refine the screened vertices using fresh oracle queries. We now query every vertex in $\mathcal{A}$ once. The second-stage classifier for $v \in \mathcal{A}$ uses only oracle responses from v to vertices in B_v, together with the reference labels $\widehat{\sigma}_u^{(0,-v)}$ for $u \in B_v$.

The separation between B_v and C_v is crucial. The event $\{v \in \mathcal{A}\}$ depends on subsampled incident edges from v only through pairs $\{u, v\}$ with $u \in C_v$. The second-stage oracle likelihood test uses only pairs $\{u, v\}$ with $u \in B_v$, and $B_v \cap C_v = \varnothing$. Moreover, $\widehat{\sigma}^{(0,-v)}$ was computed after deleting all subsampled edges incident to v. Therefore, conditional on the first-stage data and on the event $v \in \mathcal{A}$, the underlying edges from v to B_v are still independent SBM edges, and the oracle coins used in the second stage are fresh.

For one oracle query, an edge from v to u is revealed with probability w if it exists. Thus the effective second-stage edge probabilities are

$$wp = w\alpha\frac{\log n}{n}, \qquad wq = w\beta\frac{\log n}{n}.$$

Using the same single-vertex Bhattacharyya calculation as in the standard SBM exact recovery analysis, but now restricted to the reference set B_v, the error probability for classifying a fixed screened vertex is

$$\Pr\left(\widehat{\sigma}_v^{(1)} \ne \sigma_v^\star \,\middle|\, v \in \mathcal{A}\right) \le n^{-b_0 w(\sqrt{\alpha}-\sqrt{\beta})^2/2+o(1)}.$$

Here B_v has asymptotic size $b_0 n$ and contains asymptotically half of each community.

The leave-one-out reference labels on B_v are not perfectly correct, but by (14) they contain at most $n^{\epsilon+o(1)} = o(n/\log n)$ errors, uniformly in v. As before, these wrong reference labels perturb the likelihood exponent by only $o(1)$ and do not change the displayed exponent.

For our constants,

$$\frac{b_0 w}{2}(\sqrt{\alpha}-\sqrt{\beta})^2 = \frac{7/20}{2} \cdot \frac{1}{2} \cdot 9 = \frac{63}{80} = 0.7875.$$

Therefore

$$\Pr\left(\widehat{\sigma}_v^{(1)} \neq \sigma_v^\star \,\middle|\, v \in \mathcal{A}\right) \leq n^{-0.7875+o(1)}.$$

On the high-probability event (16), $|\mathcal{A}| \leq n^{0.77}$. Hence a union bound over all screened vertices gives

$$\Pr\left(\exists v \in \mathcal{A} : \widehat{\sigma}_v^{(1)} \neq \sigma_v^\star\right) \leq n^{0.77} \cdot n^{-0.7875+o(1)} = n^{-0.0175+o(1)} = o(1).$$

Thus all screened vertices are correctly recovered after one oracle query each.

Step 9: Query budget and conclusion. Choose

$$k = \left\lceil n^{0.78} \right\rceil.$$

By (16), with high probability this budget is enough to query every vertex in $\mathcal{A}$ once. Since $0.78 < 1$, we have $k = o(n)$.

With high probability, all vertices outside $\mathcal{A}$ are correctly labeled by the screening score, and all vertices inside $\mathcal{A}$ are correctly labeled by the second-stage oracle refinement. Therefore the two-stage targeted strategy achieves exact recovery with probability $1 - o(1)$ using $k = o(n)$ oracle queries. □

Theorem 8. *In Model 2, there exist SBM parameters $\alpha, \beta > 0$, a subsampling rate $s \in (0,1)$, an oracle reliability $w \in (0,1]$ and a query budget $k = o(n)$ such that simultaneously:*

1. *Every data-independent uniform querying strategy fails to achieve exact recovery w.h.p.*
2. *There exists a two-stage adaptive querying strategy that uses G_{sub} to target queries and achieves exact recovery with high probability.*

Proof. By Lemma 6, when $k = o(n)$, uniform non-adaptive querying does not improve the exact recovery threshold beyond that of the subsampled graph G_{sub} alone. Hence, for parameter choices satisfying $s(\sqrt{\alpha} - \sqrt{\beta})^2 < 2$, any uniform querying strategy fails with high probability.

On the other hand, by Lemma 7, there exist constants (α, β, s, w) and a sublinear budget $k = o(n)$ for which a two-stage strategy that uses G_{sub} to target its queries succeeds at exact recovery with high probability. Therefore, for the same observation model and the same asymptotic budget regime, uniform querying fails while a G_{sub}-dependent non-uniform strategy succeeds. This establishes a strict adaptivity gap for Model 2. □

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The claims made match theoretical results.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The paper includes a dedicated paragraph in the final section (Section 6) discussing the scope of the threshold characterizations, the uniform-querying baseline, and the modeling assumptions.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All omitted and full proofs are given in Appendices B and C. We provide proof sketches with the key ideas in the main part of the paper.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [N/A]

   Justification: Our paper is theoretical and does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [N/A]

Justification: Our paper is theoretical and does not include experiments requiring code.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [N/A]

Justification: Our paper is theoretical and does not include experiments.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [N/A]

Justification: Our paper is theoretical and does not include experiments.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [N/A]

   Justification: Our paper is theoretical and does not include experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The paper is theoretical and does not involve human subjects, datasets, deployed systems, released models, or experimental interventions. We are not aware of any aspect of the work that conflicts with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Section 6 discusses both the potential benefit of reducing the cost of network recovery and the possible privacy risks of improved inference from partial network observations.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper does not release data, models, or software artifacts with high risk for misuse; it is a theoretical work.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: The paper does not use existing assets.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were not used as an important, original, or non-standard component of the core methodology. Primary use was limited to writing, editing, or formatting assistance. After the rebuttal, simulations were added, and coding agents were used to generate the code.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.