

Predictive Learning on Hidden Tree-Structured Ising Models

Konstantinos E. Nikolakakis

K.NIKOLAKAKIS@RUTGERS.EDU

*Department of Electrical & Computer Engineering
Rutgers, The State University of New Jersey
94 Brett Road, Piscataway, NJ 08854, USA*

Dionysios S. Kalogieras

DIONYSIS@MSU.EDU

*Department of Electrical & Computer Engineering
Michigan State University
428 S. Shaw Lane, MI 48824, USA*

Anand D. Sarwate

ANAND.SARWATE@RUTGERS.EDU

*Department of Electrical & Computer Engineering
Rutgers, The State University of New Jersey
94 Brett Road, Piscataway, NJ 08854, USA*

Abstract

We provide high-probability sample complexity guarantees for exact *structure recovery* and accurate *predictive learning* using noise-corrupted samples from an acyclic (tree-shaped) graphical model. The hidden variables follow a tree-structured Ising model distribution, whereas the observable variables are generated by a binary symmetric channel taking the hidden variables as its input (flipping each bit independently with some constant probability $q \in [0, 1/2)$). In the absence of noise, predictive learning on Ising models was recently studied by [Bresler and Karzand \(2020\)](#); this paper quantifies how noise in the hidden model impacts the tasks of structure recovery and marginal distribution estimation by proving upper and lower bounds on the sample complexity. Our results generalize state-of-the-art bounds reported in prior work, and they exactly recover the noiseless case ($q = 0$). In fact, for any tree with p vertices and probability of incorrect recovery $\delta > 0$, the sufficient number of samples remains logarithmic as in the noiseless case, i.e., $\mathcal{O}(\log(p/\delta))$, while the dependence on q is $\mathcal{O}(1/(1 - 2q)^4)$, for both aforementioned tasks. We also present a new equivalent of Isserlis' Theorem for sign-valued tree-structured distributions, yielding a new low-complexity algorithm for higher-order moment estimation.

Keywords: Ising Model, Chow-Liu Algorithm, Structure Learning, Predictive Learning, Distribution Estimation, Noisy Data, Hidden Markov Random Fields

1. Introduction

Graphical models are a useful tool for modeling high-dimensional structured data. The graph captures structural dependencies: its edge set corresponds to (often physical) interactions between variables. There is a long and deep literature on graphical models (see [Koller and Friedman \(2009\)](#) for a comprehensive introduction), and they have found wide applications in areas such as image processing and vision ([Schwing and Urtasun, 2015](#); [Li and Wand, 2016](#); [Lin et al., 2016a](#); [Liu et al., 2017](#); [Morningstar and Melko, 2018](#); [Wu et al., 2017](#)), artificial intelligence more broadly ([Wainwright et al., 2003](#); [Wang et al., 2017](#)), signal processing ([Kim](#)

and Smaragdis, 2013; Wisdom et al., 2016), and gene regulatory networks (Zuo et al., 2017; Banf and Rhee, 2017), to name a few.

An *undirected graphical model*, or *Markov random field* (MRF) in particular, is defined in terms of a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, that models the Markov properties of a joint distribution on $p \triangleq |\mathcal{V}|$ node variables $(X_1, X_2, \dots, X_p) \triangleq \mathbf{X}$. A *tree-structured graphical model* is one in which \mathcal{G} is a tree. We denote the tree-structured model as $T = (\mathcal{V}, \mathcal{E})$. In this paper, we consider binary models on $2p$ variables (\mathbf{X}, \mathbf{Y}) , where the joint distribution $p(\cdot)$ of \mathbf{X} is a tree-structured Ising model distribution on $\{-1, +1\}^p$ and $\mathbf{Y} = (Y_1, Y_2, \dots, Y_p)$ is a noisy version of \mathbf{X} , such that $Y_i = N_i X_i$ and $\{N_i\}$ are independent and identically distributed (i.i.d.) Rademacher noise with $\mathbb{P}(N_i = -1) = 1 - \mathbb{P}(N_i = +1) = q$, for all $i \in \mathcal{V}$. We refer to \mathbf{X} as the *hidden layer* and \mathbf{Y} as the *observed layer*. Under this setting, our objective is to recover the underlying tree structure and accurately estimate the distribution of the hidden layer \mathbf{X} (with high probability) using only the noisy observations \mathbf{Y} . This is non-trivial because \mathbf{Y} does *not* itself follow any tree structure; this is similar to more traditional problems in nonlinear filtering, where a Markov process of known distribution (and thus, of known structure) is observed through noisy measurements (Arulampalam et al., 2002; Jazwinski, 2007; Van Handel, 2009; Douc et al., 2011; Kalogierias and Petropulu, 2016). The sample complexity of the noiseless version of our model was recently studied by Bresler and Karzand (2020), where the well-known Chow-Liu algorithm (Chow and Liu, 1968) is employed for tree reconstruction. Like them, we also analyze the Chow-Liu algorithm.

1.1 Applications and Motivating Examples

Models for joint distributions characterized by pairwise variable interactions have found many applications, with the Ising model being a popular model for binary variables. Our work is primarily motivated by examples of Ising models *corrupted by noise*. In many cases, the underlying graph-structured process cannot be observed directly; instead, only a noisy version of the process is available. Examples abound in physics, computer science, biology, medicine, psychology, social sciences, and finance. Some applications motivating this work include the following:

1) *Statistical mechanics of population, social and pedestrian dynamics* (see related work by Matsuda et al. (1992); Castellano et al. (2009)): The Ising model can be used to represent the statistical properties of the spreading of a feeling, behavior or the change of an emotional state among individuals in a crowd, where each individual interacts with his neighbors.

2) *Epidemic dynamics and epidemiological models* by Barnett et al. (2013); Erten et al. (2017): Disease spread can be modeled through the Ising model, where each individual is susceptible (spin down) or ineffective (spin up).

3) *Neoplastic transitions* and related applications in biology (Torquato (2011)): Each cell interacts with neighboring cells. Different cases are studied in the literature, for instance, healthy versus cancerous cells, malignant versus benign cells, where both can be modeled as spin up and spin down observations. The probability of diagnostic error is not zero which gives rise to the hidden model that we consider.

4) *Differential Privacy*, originally proposed by Dwork et al. (2006a,b): In computer science, differential privacy is used to guarantee privacy for individuals. A hidden model

describes data gathered using a *locally differentially private mechanism* (Warner, 1965; Kasiviswanathan et al., 2008) such as randomized response.

5) *Trading* and related applications in economics (see related work by Zhou and Sornette (2007); Takaishi (2015)): The Ising model has been considered in the literature to model increasing (spin up) or decreasing (spin down) price trends in a market.

1.2 Structure Learning for Undirected Graphical Models and Related Work

For a detailed review of methods for structure learning involving undirected and directed graphical models, see the relevant article by Drton and Maathuis (2017). In general, learning the structure of a graphical model from samples can be intractable (Karger and Srebro, 2001; Højsgaard et al., 2012). For general graphs, neighborhood selection methods (Jalali et al., 2011; Bresler, 2015; Ray et al., 2015) estimate the conditional distribution for each vertex in order to learn the neighborhood of each node and therefore the full structure. These approaches may use greedy search or ℓ_1 regularization. For Gaussian or Ising models, ℓ_1 -regularization (Ravikumar et al., 2010), the GLasso (Yuan and Lin, 2007; Banerjee et al., 2008), or coordinate descent approaches (Friedman et al., 2008) have been proposed, focusing on estimating the non-zero entries of the precision (or interaction) matrix. Model selection can also be performed using score matching methods (Hyvärinen, 2005, 2007; Nandy et al., 2015; Lin et al., 2016b), or Bayesian information criterion methods (Foygel and Drton, 2010; Gao et al., 2012; Barber et al., 2015). Other works address non-Gaussian models such as elliptical distributions, t -distribution models or latent Gaussian data (Finegold and Drton, 2011; Vogel and Fried, 2011; Vogel and Tyler, 2014; Bilodeau, 2014), or even mixed data (Fan et al., 2017).

For tree- or forest-structured models, exact inference and the structure learning problem are significantly simpler: the Chow-Liu algorithm provides an estimate of the tree or forest structure of the underlying graph (Chow and Liu, 1968; Wainwright et al., 2008; Edwards et al., 2010; Tan et al., 2011; Liu et al., 2011; Daskalakis et al., 2018; Bresler and Karzand, 2020). Furthermore, marginal distributions and maximum values are simpler to compute using a variety of algorithms (sum-product, max-product, message passing, variational inference) (Pearl, 1988; Lauritzen, 1996; Wainwright et al., 2003, 2008)).

The *noiseless counterpart* of the model considered in this paper was studied recently by Bresler and Karzand (2020); in this paper, we extend their results to the hidden case, where samples from a tree-structured Ising model are passed through a binary symmetric channel with crossover probability $q \in [0, 1/2)$. Of course, in the special case of a linear graph, our model reduces to a hidden Markov model. Latent variable models are often considered in the literature when some variables of the graph are deterministically unobserved (Chandrasekaran et al., 2010; Anandkumar and Valluvan, 2013; Ma et al., 2013; Anandkumar et al., 2014). Our model is most similar to that studied by Chaganty et al. (Chaganty and Liang, 2014), in which a hidden model is considered with a discrete exponential distribution and Gaussian noise. They solve the parameter estimation problem by using moment matching and pseudo-likelihood methods; the structure can be recovered indirectly using the estimated parameters.

Connection with Phylogenetic Estimation. In phylogenetic estimation problems the goal is to learn the structure of tree given only observations from the leaves (Erdős et al.,

1999). The sample complexity of phylogenetic reconstruction algorithms grows exponentially with respect to the depth of the tree (Erdős et al., 1999), however if we are interested in reconstructing only parts of the tree which are “close” to the leaves then the depth of tree does not affect the sample complexity (Daskalakis et al., 2009). The hidden structure learning problem that we consider in this paper is a special case of phylogeny estimation problem with constant depth; there is exactly one noisy observable for each hidden node of the tree. In contrast with phylogenetic estimation approaches, Chow-Liu algorithm is simple and computationally more efficient, while the sample complexity is of the same order¹ with the well-known phylogenetic reconstruction methods, to name a few “Dyadic Closure” method by Erdős et al. (1999), the “Contractor-Extender” and “Cherry-picking” algorithms by Daskalakis et al. (2006, 2009, 2013). On the other hand, the approach of distribution estimation by matching the structure and the correlations (Bresler and Karzand, 2020) has not been considered in the phylogenetic estimation literature. Based on the above discussion, the following interesting question naturally rises: How well can we estimate the distribution of a hidden tree structured model while having access only to the leaves of the tree? The latter remains open problem for future work.

1.3 Statement of Contributions

We are interested in answering the following general question: *How does noise affect the sample complexity of the structure and predictive learning procedure?* That is, given *only* noisy observations, our goal is to learn the tree structure of the hidden layer in a well-defined and meaningful sense. The MLE-structure from tree-structured (noiseless) data is the output of the Chow-Liu algorithm (Chow and Liu, 1968). However, the MLE-structure from noisy data is not consistent with the hidden structure in general because the graphical model of the observables is a complete graph. Further, the (latent) MLE of the actual interaction parameters θ of the hidden layer is intractable. In Sections 2.4 and 2.6 we explain the importance of Chow-Liu algorithm in our setting, we show why the classical MLE approach fails, and we discuss the connection between the output of the Chow-Liu algorithm and an alternative, projection-based MLE approach.

The estimated structure is an essential statistic for estimating the underlying distribution of the hidden layer, allowing for predictive learning. Specifically, based on the structure estimate, we are also interested in appropriately approximating the tree-structured distribution under study, which can then be used for accurate predictions. We also consider the problem of hidden layer higher-order moment estimation of tree-structured Ising models and, in particular, how such estimation can be efficiently performed, on the basis of noisy observations.

A summary of the main contributions of this paper is as follows:

- A lower bound on the sufficient number of samples needed to recover the exact hidden structure with high probability, by using the Chow-Liu algorithm. We also show an upper bound on the necessary number of samples for any algorithm to estimate the hidden structure. The proof of the lower bound follows the general structure of Lemmata 8.1-8.4 by Bresler and Karzand (2020), however we need to extend the necessary events and prove new concentration bounds for the noisy setting. Although

1. while considering the depth fixed

the graphical model of the observables is a complete graph we show that the Chow-Liu algorithm (with input a finite number of noisy samples) returns the exact tree of the hidden layer with high probability and we characterize its sample complexity. The proof of the upper bound uses the same construction of the approach in Section 7.1 by [Bresler and Karzand \(2020\)](#) but requires the combination of Fano’s inequality and a strong data processing inequality (SDPI) by [Polyanskiy and Wu \(2017\)](#). Specifically, we show that SDPI’s can be a useful tool to derive minimax bounds when closed form expressions or upper bounds of the KL-divergence are hard to be found. The later is of independent interest and it can be applied to other machine learning problems that involve noisy observations.

- Determination of the sufficient and necessary number of samples for accurate predictive learning. We analyze the sample complexity of learning distribution estimates, which can accurately provide predictions on the hidden tree. The estimates are computed using the noisy data. Predictive learning under noisy samples is challenging because structural properties such as the independence of random variables $X_i X_j$ and correlation estimates $\hat{\mathbb{E}}[X_i X_j]$ for $(i, j) \in \mathcal{E}$ do not hold for the noisy observable \mathbf{Y} . To overcome this we evaluate the required conditional distributions of the dependent variables, construct a martingale difference sequence, and prove a high probability bound of the event that involves these variables by applying a concentration bound for supermartingales (generalized Bennet’s inequality ([Fan et al., 2012](#))). We refer the reader to Section 4.3 for a detailed discussion about the technical contributions and a sketch of proof of the main result.
- A closed-form expression *and* a computationally efficient estimator for higher-order moment estimation in tree-structured Ising models. This result corresponds to an equivalent statement of Isserlis’ theorem for sign-valued tree models. Given pair-wise correlations and the tree (or estimates of both, from noisy or noiseless data) we provide an algorithm that runs on the tree and returns the expression of high-order moments. The proof involves the existence and identification of (minimum length) disjoint paths among any set of pairs of nodes. The proposed algorithm (Algorithm 2) identifies these paths that yield the expression of the moments. The results may be of independent interest for a computational efficient exact or approximated higher-moment evaluation.

Our main results Theorem 5 and Theorem 7 provide the amount of finite samples needed for exact structure recovery and accurate predictive learning with high probability. Although we are interested in the finite sample complexity bounds, our results are also asymptotically optimal. That is, for any fixed (constant) $q \in [0, 1/2)$ the order of the upper bound (necessary number of samples) matches the corresponding (lower) minimax bound. The sample complexity bounds that we provide are the extended form of state of the art (noiseless setting) bounds by [Bresler and Karzand \(2020\)](#). By setting $q = 0$, our bounds reduce to the noiseless setting bounds. Further, the explicit version of our results (see Section 3) are continuous functions of the cross-over probability q .

1.4 Notation

Boldface indicates a vector or tuple and calligraphic face for sets and trees. The sets of even and odd natural numbers are $2\mathbb{N}$ and $2\mathbb{N} + 1$ respectively. For an integer n , define

Symbol	Meaning
p	number of variables nodes in the tree
$p(\mathbf{x})$	$\exp(\sum_{(i,j) \in \mathcal{E}} \theta_{ij} x_i x_j) / Z(\theta)$, $\mathbf{x} \in \{-1, +1\}^p$, $Z(\theta)$: partition function
α	minimum $ \theta_{ij} $ in the Ising model, $\min_{i,j \in \mathcal{V}} \theta_{ij} $
β	maximum $ \theta_{ij} $ in the Ising model, $\max_{i,j \in \mathcal{V}} \theta_{ij} $
T	Original tree of the model
$\mathcal{P}_T(\alpha, \beta)$	set of tree-structured Ising models with $\alpha \leq \theta_{ij} \leq \beta$
n	number of samples
q	crossover probability of the BSC, $q \in [0, 1/2]$
c_q	$1 - 2q$
$p(\cdot)$	distribution of the hidden node variables, $\mathbf{X} \sim p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$
$p_{\dagger}(\cdot)$	distribution of the observable node variables $\mathbf{Y} \sim p_{\dagger}(\cdot)$
$\mathbf{1}_A$	indicator function of the set A
D_{KL}	KL divergence
S_{KL}	symmetric KL divergence
$I(X, Y)$	mutual information of X, Y
d_{TV}	total variation distance
$\mathcal{L}^{(2)}(P, Q)$	$\sup_{i,j \in \mathcal{V}} d_{\text{TV}}(P_{ij}, Q_{ij})$, and P_{ij}, Q_{ij} the pairwise marginals of P, Q
$\mathbf{X}^{1:n}$	n independent observations of \mathbf{X}
$\mathbf{Y}^{1:n}$	n independent observations of \mathbf{Y}
T^{CL}	Chow-Liu-estimated structure from noiseless data $\mathbf{X}^{1:n}$
T_{\dagger}^{CL}	Chow-Liu-estimate of the hidden tree structure T from noisy data $\mathbf{Y}^{1:n}$
$\text{path}_T(w, \tilde{w})$	the set of edges which connects the nodes $w, \tilde{w} \in \mathcal{V}_T$
$\hat{\mu}_{i,j}$	$\frac{1}{n} \sum_{k=1}^n X_i^{(k)} X_j^{(k)}$
$\hat{\mu}_{i,j}^{\dagger}$	$\frac{1}{n} \sum_{k=1}^n Y_i^{(k)} Y_j^{(k)}$
$\Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})$	estimator of the distribution $p(\cdot)$ from noisy data $\mathbf{Y}^{1:n}$
η	maximum error on the distribution estimation: $\mathcal{L}^{(2)}(P, \hat{P}) \leq \eta$
δ	maximum probability of error, the notation depends on the task in structure estimation: $\mathbb{P}(T_{\dagger}^{\text{CL}} \neq T) \leq \delta$ in predictive learning: $\mathbb{P}\left(\mathcal{L}^{(2)}\left(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})\right) \leq \eta\right) \geq 1 - \delta$.

Table 1: Notation/Definitions.

$[n] \triangleq \{1, 2, \dots, n\}$. The indicator function of a set A is $\mathbf{1}_A$. For a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $\mathcal{V} = [p]$ indexes the set of variables $\{X_1, X_2, \dots, X_p\}$, for any pair of vertices $i, j \in \mathcal{V}$ the correlation $\mu_{ij} \triangleq \mathbb{E}[X_i X_j]$ and for any edge $e = (i, j) \in \mathcal{E}$ it is $\mu_e \triangleq \mathbb{E}[X_i X_j]$. For two nodes w, \tilde{w} of a tree, the term $\text{path}(w, \tilde{w})$ denotes the set of edges in the unique path with endpoints w and \tilde{w} . Further, $\text{BSC}(q)^p$ denotes a binary symmetric channel with crossover probability q and block-length p . The $\text{BSC}(q)^p$ is a conditional distribution from $\{-1, 1\}^p \rightarrow \{-1, 1\}^p$ that acts componentwise independently on \mathbf{X} to generate \mathbf{Y} , such that $X_i = N_i Y_i$ and \mathbf{N} is a vector of i.i.d. Rademacher variables equal to $+1$ with probability $1 - q$. We use the symbol \dagger to indicate the corresponding quantity for the observable (noisy) layer. For instance, $p_{\dagger}(\cdot)$

is the probability mass function of \mathbf{Y} and $\mu_{i,j}^\dagger \triangleq \mathbb{E}[Y_i Y_j]$ corresponds to the correlation of variables Y_i, Y_j . For our readers' convenience, we summarize the notation in Table 1.




Figure 1: The simulation corresponds to structure learning. Comparison of the experimental results (heat-map) and the theoretical bound of Theorem 5, the bound that yields (1). The colored regions denote different values of the estimated probability of error δ (at least one edge has been missed). The value of δ varies between 0 and 1 while the parameters $\alpha = 0.2, \beta = 1.1, p = 100$ are fixed. The red line shows the bound from Theorem 5 (the explicit form of Theorem 1). The code of the experiment is available at <https://github.com/KonstantinosNikolakis/Structure-Learning>.

1.5 Summary of the Results

In this section, we present a summary of the main results of our work up to constant factors $C, C' > 0$. We refer the reader to Table 1 for the definition of the model parameters. We provide the explicit statements of the results, and we specify the constants in Section 3. Recall that, the random vector $\mathbf{Y} \in \{-1, +1\}^p$ is the output of the binary symmetric channel $\text{BSC}(q)^p$ with input the random vector $\mathbf{X} \sim p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$.

1.5.1 STRUCTURE LEARNING

The first results provides the sufficient number of samples for exact structure recovery.

Theorem 1 (Sample Complexity for Structure Learning.) *The Chow-Liu algorithm with input n noisy samples $\mathbf{Y}^{1:n}$ exactly estimates the hidden tree structure $\mathbf{T}_\dagger^{CL} \equiv \mathbf{T}$ with probability at least $1 - \delta \in (0, 1)$, as long as*

$$n > C \frac{e^{2\beta(1+\mathbf{1}_{q \neq 0})}}{(1-2q)^4 \tanh^2(\alpha)} \log(p/\delta). \quad (1)$$

The order with respect to β is $\mathcal{O}(e^{4\beta})$ for all $q > 0$. The bound in (1) exactly reduces to the noiseless case (Bresler and Karzand, 2020, Theorem 3.2). Additionally, the explicit form of the result, Theorem 5, shows that the bound is also a continuous function of $q \in [0, 1/2)$. The next proposition gives the necessary number of samples for exact structure recovery.

Proposition 2 *No algorithm can recover the structure with probability great than 1/2 if*

$$n < C' \frac{e^{2\beta} [1 - (4q(1-q))^p]^{-1}}{\alpha \tanh(\alpha)} \log(p). \quad (2)$$

Note that the terms $(1-2q)^{-4}$ and $[1 - (4q(1-q))^p]^{-1}$ introduce a gap between the sample complexity of (1) and (2). However, the sample complexity of Theorem 1 is indeed accurate. To illustrate this experimentally, we show that the theoretical and experimental bounds exactly match, see Figure 1. The latter indicates that the Chow-Liu algorithm requires exactly the number of samples that our theoretical result suggest (see Figure 1). On the other hand, Proposition 2 provides the necessary number of samples, for any algorithm. Finally, we conjecture that the bound of Proposition 1.2 is tight only under the low temperature regime $|\theta_{i,j}| \rightarrow \infty$ for all $i, j \in \mathcal{E}$. The derivation of generalized tighter forms of the bound in (2) is challenging and left for future work.

1.5.2 PREDICTIVE LEARNING

To learn the tree-shaped distribution $p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$ of \mathbf{X} from n noisy samples $\mathbf{Y}^{1:n}$, we first estimate the correlations $\hat{\mu}_{i,j}^\dagger$ for all $i, j \in \mathcal{V}$. We then estimate the tree structure T_\dagger^{CL} by running the Chow-Liu algorithm with input the candidate edge weights $\hat{\mu}_{i,j}^\dagger$ and finally evaluate the estimator of $p(\cdot)$ (by matching correlations) as follows²

$$\Pi_{T_\dagger^{\text{CL}}}(\hat{p}_\dagger) \triangleq \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_{T_\dagger^{\text{CL}}}} \frac{1 + x_i x_j \frac{\hat{\mu}_{i,j}^\dagger}{(1-2q)^2}}{2}, \quad \mathbf{x} \in \{-1, +1\}^p. \quad (3)$$

Note that one restriction of our approach is that the distribution estimator requires the value q to be known. The same restriction appears in other structure learning from noisy data approaches (Goel et al., 2019). However, in our setting q is required only for the predictive learning, while the Chow-Liu algorithm and the structure estimation does not require q to be known. Under the assumption that q is unknown, one can first learn its value through an independent procedure (Goel et al., 2019, Section 5). The accuracy of the estimated distribution in (3) is measured by the small-set Total Variation (ssTV), that captures the estimation error on the k^{th} -order marginals (Georgii, 2011; Rebeschini et al., 2015; Bresler and Karzand, 2020). Let P_S, Q_S denote the marginals of P, Q on a set $\mathcal{S} \subset \mathcal{V}$, and $|\mathcal{S}| = k$. Then the k^{th} order ssTV of P and Q is defined as

$$\mathcal{L}^{(k)}(P, Q) \triangleq \sup_{\mathcal{S}: |\mathcal{S}|=k} d_{\text{TV}}(P_S, Q_S). \quad (4)$$

2. The distribution in (3) is a function of \mathbf{x} , however we suppress the notation for consistency with prior work and for sake of space.




Figure 2: The simulation corresponds to predictive learning. Comparison of the experimental results (heat-map) and the theoretical bound of Theorem 3. The colored regions denote different values of the estimated probability of error δ (ssTV to be greater than a fixed number η). The value of δ varies between 0 and 1 while the parameters $\eta = 0.03, \beta = 1.1, p = 31$ are fixed. The code of the experiment is available at <https://github.com/KonstantinosNikolakis/Predictive-Learning>.

The next results provides the necessary number of samples for accurate distribution estimation by guaranteeing that the $\mathcal{L}^{(2)}$ is less than a small positive number η with high probability. We provide guarantees on higher-order marginals ($k > 2$) in Section 3.3.

Theorem 3 (Sample Complexity for Predictive Learning) *Fix $\delta \in (0, 1)$. Choose $\eta > 0$ (independent of δ). If*

$$n \geq C \max \left\{ \frac{1}{\eta^2(1-2q)^4}, \frac{e^{2\beta(1+\mathbb{1}_{q \neq 0})}}{(1-2q)^4}, \frac{e^{4\beta}}{\eta^2} \mathbb{1}_{q \neq 0} \right\} \log \left(\frac{p}{\delta} \right) \quad (5)$$

then

$$\mathbb{P} \left(\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger}) \right) \leq \eta \right) \geq 1 - \delta. \quad (6)$$

Note that the dependence on β is $\mathcal{O}(e^{4\beta})$ for accurate distribution learning from noisy data (similarly to the structure learning task, Theorem 1). The bound in (5) *exactly* reduces to the noiseless setting bound by (Bresler and Karzand, 2020, Theorem 3.3). Theorem 3 is a short version of the main result of the paper. The explicit statement, Theorem 7, shows that the bound is also continuous at $q \rightarrow 0$.

Conversely, the following proposition gives an upper bound on the necessary number of samples for accurate marginal distributions' estimation under the assumption $\beta > \alpha$.

Sufficient Number of Samples		
Task/Setting	Noiseless (prior work)	Noisy
Structure Learning	$C \frac{e^{2\beta}}{\tanh^2(\alpha)} \log(p/\delta)$	$C \frac{e^{2\beta(1+\mathbb{1}_{q \neq 0})}}{(1-2q)^4 \tanh^2(\alpha)} \log(p/\delta)$
Predictive Learning	$C \max\{\eta^{-2}, e^{2\beta}\} \log(p/\delta)$	$C \max\left\{\frac{\eta^{-2}}{(1-2q)^4}, \frac{e^{2\beta(1+\mathbb{1}_{q \neq 0})}}{(1-2q)^4}, \frac{e^{4\beta} \mathbb{1}_{q \neq 0}}{\eta^2}\right\} \log(p/\delta)$

Table 2: Sufficient number of samples for accurate structure and predictive learning.

Necessary Number of Samples		
Task/Setting	Noiseless (prior work)	Noisy
Structure learning	$C' \frac{e^{2\beta}}{\alpha \tanh(\alpha)} \log(p)$	$C' \frac{e^{2\beta[1-(4q(1-q))^p]-1}}{\alpha \tanh(\alpha)} \log(p)$
Predictive learning	$C' \eta^{-2} \log(p)$	$C' \eta^{-2} [1 - (4q(1-q))^p]^{-1} \log(p)$

Table 3: Necessary number of samples for structure and predictive learning.

Proposition 4 *Fix $\eta > 0$ such that $\eta \leq (\tanh(\beta) - \tanh(\alpha))/2$. Then no algorithm can accurately estimate the distribution of the hidden variables (ssTV less than $\eta > 0$) with probability greater than $1/2$ if*

$$n < C' \eta^{-2} [1 - (4q(1-q))^p]^{-1} \log(p). \quad (7)$$

A quick comparison of (5) and (7) shows that there is a gap between the sufficient and necessary number of samples. Our experiments (Figure 2) confirm the accuracy of our theoretical results. For instance the bound of Theorem 3 exactly matches the experimental curve. For further discussion related to the gap between the upper and lower bounds see Section 2.6. Further, we conjecture that bound in (7) is tight only under the low temperature regime, similarly to the Proposition 2. The derivation of tighter characterization of the necessary number of samples Propositions 2 and 4 remains an problem for future work. Additional plots of the experiments are provided in Section 3.4. Finally, Table 2 and 3 summarize the state-of-the-art bounds of the noiseless setting by [Bresler and Karzand \(2020\)](#) and the extended version under the noisy setting that we study in this paper.

To summarize, the following holds for both structure and predictive learning: the dependence on the parameter β is of the order $\mathcal{O}(e^{2\beta})$ for $q = 0$ and becomes $\mathcal{O}(e^{4\beta})$ for positive values of q . Further, the bounds are continuous functions of q , as our results suggest (for the continuity see the explicit form of the results Theorem 5 and Theorem 7.) Similarly to the noiseless case, the following statement holds when noise exists as well: Under the high temperature regime (α close to zero), structure learning requires much more data than the predictive learning task, because of the $\tanh^2(\alpha)$ in the denominator of the bound in (1). On the contrary, the required number of samples for predictive learning (5) does not depend on α . Specifically, exact structure recovery is not necessary for learning the distribution

efficiently, that is, *weak* edges' identification failure does not affect the predictive learning task. We refer the reader to Section 4.3 for the definition of *weak/strong* edges and additional explanation. Finally, for $q > 0$ an extra term that involves both β and η appears in the bound of Theorem 3, while for values of q close to zero and $q = 0$ vanishes.

The pairwise correlations of end-point vertices ($E[X_i X_j]$: $(i, j) \in \mathcal{E}$) are sufficient statistics, and as expected, the accuracy of pairwise marginals corresponds to accuracy of higher order marginals and accurate estimation of higher order moments. In Sections 3.3 and 4.4 we provide a method for evaluating higher order moments (and marginals) from noisy observations. Our approach is based on an equivalent of Isserlis' theorem for tree-structured Ising models that is also of independent interest.

2. Preliminaries and Problem Statement

In this section, we introduce our model of *hidden sign-valued Markov random fields on trees*.

2.1 Undirected Graphical Models

We consider *sign-valued graphical models* where the joint distribution $p(\cdot)$ has support $\{-1, +1\}^p$. Let $\mathbf{X} = (X_1, X_2, \dots, X_p) \in \{-1, +1\}^p$ be a collection of sign-valued (binary) random variables. Then, $\mathbb{1}_{X_i=x_i} \equiv (1 + x_i X_i)/2$, and the distribution of \mathbf{X} is

$$p(\mathbf{x}) = \mathbb{E} \left[\prod_{i=1}^p \mathbb{1}_{X_i=x_i} \right] = \frac{1}{2^p} \left[1 + \sum_{k \in [p]} \sum_{\mathcal{S} \subset \mathcal{V}: |\mathcal{S}|=k} \mathbb{E} \left[\prod_{s \in \mathcal{S}} X_s \right] \prod_{s \in \mathcal{S}} x_s \right], \quad \mathbf{x} \in \{-1, +1\}^p. \quad (8)$$

In this paper we assume that the marginal distributions of the X_i are uniform, that is,

$$\mathbb{P}(X_i = \pm 1) = \frac{1}{2}, \quad \forall i \in \mathcal{V}. \quad (9)$$

Thus, $\mathbb{E}[X_i] = 0$, for all $i \in \mathcal{V}$. A distribution is Markov with respect to a hypergraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ if for every node i in the set \mathcal{V} it is true that $\mathbb{P}(X_i | \mathbf{x}_{\mathcal{V} \setminus \{i\}}) = \mathbb{P}(X_i | \mathbf{x}_{\mathcal{N}(i)})$, where $\mathcal{N}(i)$ is the set of neighbors of i in \mathcal{G} . One subclass of distributions for which the Markov property holds is the *Ising model*, in which the random variables X_i are sign-valued and the hypergraph is a simple undirected graph, indicating that variables have only pairwise and unary interactions. The joint distribution for the *Ising model* with zero external field is given by

$$p(\mathbf{x}) = \frac{1}{Z(\theta)} \exp \left\{ \sum_{(s,t) \in \mathcal{E}} \theta_{st} x_s x_t \right\}, \quad \mathbf{x} \in \{-1, 1\}^p. \quad (10)$$

$\{\theta_{st} : (s, t) \in \mathcal{E}\}$ are parameters of the model representing the interaction strength of the variables and $Z(\cdot) \in (0, \infty)$ is the *partition function*. These interactions are expressed through potential functions $\exp(\theta_{st} x_s x_t)$ that ensure that the Markov property holds with respect to the graph $G = (\mathcal{V}, \mathcal{E})$. Next, we discuss the properties of distributions of the form of (8), *which are Markov with respect to a tree*.

2.2 Sign-Valued Markov Fields on Trees

From prior work by [Lauritzen \(1996\)](#), it is known that any distribution $p(\cdot)$ that is Markov with respect to a tree (or forest) $T = (\mathcal{V}, \mathcal{E})$ factorizes as

$$p(\mathbf{x}) = \prod_{i \in \mathcal{V}} p(x_i) \prod_{(i,j) \in \mathcal{E}} \frac{p(x_i, x_j)}{p(x_i)p(x_j)}, \quad \mathbf{x} \in \{-1, +1\}^p, \quad (11)$$

and we call $p(\cdot)$ as tree (forest) structured distribution, to indicate the factorization property. If the distribution $p(\cdot)$ has the form of (8) with $\mathbb{P}(X_i = \pm 1) = 1/2$, for all $i \in \mathcal{V}$, and is Markov with respect to a tree T , then

$$p(\mathbf{x}) = \frac{1}{2} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{2} \quad (12)$$

and

$$\mathbb{E}[X_i X_j] = \prod_{e \in \text{path}(i,j)} \mu_e, \quad \text{for all } i, j \in \mathcal{V}. \quad (13)$$

(see Appendix A, Lemma 12). Additionally, let us state the definition of the so-called *Correlation (coefficient) Decay Property (CDP)*, that will be of central importance in our analysis.

Definition 1 *The CDP holds if and only if $|\mathbb{E}[X_i X_k]| \geq |\mathbb{E}[X_\ell X_m]|$ for all tuples $\{i, k, \ell, m\} \subset \mathcal{V}$ such that $\text{path}(i, k) \subset \text{path}(\ell, m)$.*

The CDP is a well known attribute of acyclic Markov fields (see, e.g., [Tan et al. \(2010\)](#), [Bresler and Karzand \(2020\)](#)). Further, it is true that the products $X_i X_j$ for all $(i, j) \in \mathcal{E}$ are independent and the CDP holds for every $p(\cdot)$ of the form of (8), that factorizes with respect to a tree (see Lemma 13, Appendix A). This is a consequence of property (13) and the inequality $|\mu_e| \leq 1$, for all $e \in \mathcal{E}$. We can interpret the CDP as a type of data processing inequality (see [Cover and Thomas \(2012\)](#)). The connection is clear through the relationship between the mutual information $I(X_i, X_j)$ and the correlations $\mathbb{E}[X_i X_j]$, namely,

$$I(X_i, X_j) = \frac{1}{2} \log_2 \left((1 - \mathbb{E}[X_i X_j])^{1 - \mathbb{E}[X_i X_j]} (1 + \mathbb{E}[X_i X_j])^{1 + \mathbb{E}[X_i X_j]} \right), \quad (14)$$

for any pair of nodes $i, j \in \mathcal{V}$. This expression shows that the mutual information is a symmetric function of $\mathbb{E}[X_i X_j]$ and increasing with respect to $|\mathbb{E}[X_i X_j]|$ (see also Lemma 17, Appendix A).

Tree-structured Ising models: Despite its simple form, the Ising model has numerous useful properties. In particular, (12), (13) hold for any tree-structured Ising model with uniform marginal distributions and $\theta_r = 0$ for all $r \in \mathcal{V}$. Furthermore,

$$\mathbb{E}[X_i X_j] = \tanh \theta_{ij}, \quad \forall (i, j) \in \mathcal{E}_T, \quad (15)$$

the latter implies that

$$p(\mathbf{x}) = \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_T} \frac{1 + x_i x_j \tanh \theta_{ij}}{2}, \quad \mathbf{x} \in \{-1, 1\}^p, \quad \alpha \leq |\theta_{ij}| \leq \beta, \quad (16)$$

$$\mathbb{E}[X_i X_j] = \prod_{e \in \text{path}(i,j)} \mu_e = \prod_{e \in \text{path}(i,j)} \tanh(\theta_e), \quad \forall i, j \in \mathcal{V}. \quad (17)$$

A short argument showing (15) and (16) is included in Appendix A, Lemma 14. For the rest of the paper, we assume a tree-structured Ising model for the hidden variable \mathbf{X} , that is, the distribution of \mathbf{X} has the form of (12). We also impose a reasonable compactness assumption on the respective interaction parameters, as follows.

Assumption 1 *There exist α and β such that for the distribution $p(\cdot)$, $0 < \alpha \leq |\theta_{st}| \leq \beta < \infty$ for all $(s, t) \in \mathcal{E}$.*

For a fixed tree structure T , and for future reference, we hereafter let $\mathcal{P}_T(\alpha, \beta)$ be the class of Ising models satisfying Assumption 1.

2.3 Hidden Sign-Valued Tree-Structured Models

The problem considered in this paper is that of learning a tree-structured model from corrupted observations. Because we have no access to the original samples $\mathbf{X}^{1:n}$, we obtain the noisy observations $\mathbf{Y}^{1:n}$. To formalize this, consider a hidden Markov random field whose hidden layer \mathbf{X} is an Ising model with respect to a tree, i.e., $\mathbf{X} \sim p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$, as defined in (16). The observed variables \mathbf{Y} are formed by setting $Y_r = N_r X_r$ for all $r \in \mathcal{V}$, where $\{N_r\}$ are i.i.d. Rademacher(q) random variables. Let $p_{\dagger}(\cdot)$ be the distribution of the observed variables \mathbf{Y} . We can think of \mathbf{Y} as the result of passing \mathbf{X} through a binary symmetric channel $\text{BSC}(q)^p$. We have the following expressions

$$\mathbb{E}[N_r] = 1 - 2q \triangleq c_q, \quad \forall r \in \mathcal{V}, \text{ and } q \in [0, 1/2), \quad (18)$$

$$\mu_{r,s}^{\dagger} \triangleq \mathbb{E}[Y_r Y_s] = \mathbb{E}[N_r X_r N_s X_s] = (1 - 2q)^2 \mathbb{E}[X_r X_s], \quad \forall r, s \in \mathcal{V}. \quad (19)$$

The distribution $p_{\dagger}(\cdot)$ of \mathbf{Y} also has support $\{-1, +1\}^p$, and so the joint distribution satisfies the general form (8). Since the marginal distribution of each Y_r is also uniform, $\mathbb{E}[Y_r] = 0$ for all $r \in \mathcal{V}$, (8) and (18) yield

$$p_{\dagger}(\mathbf{y}) = \mathbb{E} \left[\prod_{i=1}^p \mathbf{1}_{Y_i=y_i} \right] = \frac{1}{2^p} \left[1 + \sum_{k \in [p] \cap 2\mathbb{N}} c_q^k \sum_{S \subset \mathcal{V}: |S|=k} \mathbb{E} \left[\prod_{s \in S} X_s \right] \prod_{s \in S} y_s \right], \quad \mathbf{y} \in \{-1, 1\}^p. \quad (20)$$

The moments of the hidden variables $\mathbb{E} \left[\prod_{s \in S} X_s \right]$ in (20) can be expressed as products of the pairwise correlations $\mathbb{E}[X_s X_t]$, for any $(s, t) \in \mathcal{E}_T$ (Section 3.3, Theorem 10). From (20) it is clear that the distribution $p_{\dagger}(\cdot)$ of \mathbf{Y} does not factorize with respect to any tree, that is, $p_{\dagger}(\cdot) \notin \mathcal{P}_T(\alpha, \beta)$ in general.³

2.4 Hidden Structure Estimation

We are interested in characterizing the *sample complexity* of structure recovery: given data generated from $p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$ for an unknown tree T , what is the minimum number n_{\dagger}

3. Lemma 30 shows the structure preserving property for the observable layer holds for the special case of single-edge forests.

Algorithm 1 Chow – Liu

Require: $\mathcal{D} = \{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(n)}\} \in \{-1, 1\}^{p \times n}$, where $\mathbf{y}^{(k)}$ is the k^{th} observation of \mathbf{Y}

- 1: Compute $\hat{\mu}_{i,j}^\dagger \leftarrow \frac{1}{n} \sum_{k=1}^n y_i^{(k)} y_j^{(k)}$, for all $i, j \in \mathcal{V}$
- 2: **return** $T_\dagger^{\text{CL}} \leftarrow \text{MaximumSpanningTree}\left(\cup_{i \neq j} \left\{ \left| \hat{\mu}_{i,j}^\dagger \right| \right\}\right)$

of samples $\{\mathbf{y}^{(i)}, i \in [n_\dagger]\}$ from $p_\dagger(\cdot)$ needed to recover the (unweighted) edge set of T with high probability? In particular, we would like to quantify how n_\dagger depends on the crossover probability q . Intuitively, noise makes “weak” edges to appear “weaker”, and the sample complexity is expected to be an increasing function of q . Because the distribution $p_\dagger(\cdot)$ of the observable variables does not factorize according to any tree, this problem does not follow directly from the noiseless case. Although the classical MLE is the standard approach for the noiseless case, for the noisy setting the MLE estimation of parameters θ of the hidden model is intractable, due to the summation over the support of \mathbf{X} . Additionally, the MLE structure estimate from noisy data is not in general consistent with the hidden structure as we explain in Section 2.6. However, for the model that we consider in this paper, the projected-MLE estimate of the observables onto the space of tree-structured models gives a consistent structure estimate. Additionally, that structure estimate is identical to the output of Chow-Liu algorithm (Algorithm 1) from noisy data. We refer the reader to Section 2.6 for the discussion about the MLE and the connection with the noisy Chow-Liu algorithm.

In this work, we use and analyze the sample complexity of the classical Chow-Liu algorithm (Algorithm 1) for the following reasons: We show that given finite number of noisy data as input, the Chow-Liu algorithm recovers the original tree T with high probability. Further the sample complexity is asymptotically optimal for fixed $q < 1/2$ (see Tables 2 and 3), and its order remains $\mathcal{O}(\log p)$ in the high dimensional regime. The algorithm is computationally efficient in comparison to other optimization techniques and it does not require the value q to be known. Additionally, Algorithm 1 solves the projected-MLE problem that we discuss in Section 2.6. The above reasons and our finite sample complexity bound Theorems 1 and 5 suggest that Algorithm 1 is an excellent approach for tree-structure learning from noisy data.

2.5 Evaluating the Accuracy of the Estimated Distribution

In addition to recovering the graph structure, we are interested in the “goodness of fit” of the estimated distribution. Let P_S, Q_S be the marginal distributions of P, Q on the set $S \subset \mathcal{V}$, let d_{TV} denote the total variation distance, and fix $k = 2$. We measure the error of distribution estimator through the “small set Total Variation” (or ssTV) distance as defined by [Bresler and Karzand \(2020\)](#)

$$\mathcal{L}^{(k)}(P, Q) \triangleq \sup_{S: |S|=k} d_{\text{TV}}(P_S, Q_S). \quad (21)$$

If Q is an estimate of P , the norm $\mathcal{L}^{(k)}$ guarantees predictive accuracy because ([Bresler and Karzand, 2020](#), Section 3, page 720)

$$\mathbb{E}_{X_S} \left[|P(X_i = +1 | X_S) - Q(X_i = +1 | X_S)| \right] \leq 2\mathcal{L}^{(|S|+1)}(P, Q). \quad (22)$$

The estimated (from noisy data) distribution of the hidden variables in (3) is a simple extension of the noiseless estimate. In fact the estimated distribution factorizes according to the estimated from noisy data tree structure, that is the output of Algorithm 1. Further, the pairwise correlations are normalized by the constant $(1 - 2q)$. As a result, the estimator is consistent because if $n \rightarrow \infty$ then $T_{\dagger}^{\text{CL}} \rightarrow T$, $\hat{\mu}_{i,j}^{\dagger}/(1 - 2q) \rightarrow \mu_{i,j}$, and as a consequence the estimate $\Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})$ convergence to the original distribution $p(\cdot)$ of \mathbf{X} . Our main result gives a lower bound on the number of samples needed to guarantee accurate estimation (in the sense of small ssTV), with high probability.

2.6 Maximum Likelihood Estimate

A natural first place to start in estimation is the maximum-likelihood estimate (MLE). We explain why this is problematic and show a method (the projected-MLE) which turns out to be equivalent to the Chow-Liu algorithm. This motivates why we study the Chow-Liu algorithm in the first place. To begin, the distribution of the observables parametrized over the interaction parameters θ of the hidden layer is

$$p_{\dagger}(\mathbf{y}) = \sum_{\mathbf{x} \in \{-1, +1\}^p} \frac{1}{Z(\theta)} \exp \left\{ \sum_{(s,t) \in \mathcal{E}_{\mathbf{G}}} \theta_{st} x_s x_t \right\} p(\mathbf{y}|\mathbf{x}), \quad \mathbf{y} \in \{-1, 1\}^p. \quad (23)$$

It is known that above expression is intractable in closed form and it can be evaluated only through approximations. Secondly, the log-likelihood of \mathbf{Y} can be written as

$$\log p_{\dagger}(\mathbf{y}) = \log \sum_{\mathbf{x} \in \{-1, +1\}^p} p(\mathbf{y}|\mathbf{x}) \prod_{i \in V} p(x_i) \prod_{(i,j) \in \mathcal{E}} \frac{p(x_i, x_j)}{p(x_i)p(x_j)}, \quad \mathbf{y} \in \{-1, 1\}^p, \quad (24)$$

and the logarithm of the summation cannot be expressed as summation of logarithms. Therefore we see the classical MLE structure estimation approach is not applicable for hidden models. Specifically, the structure of the observable layer is a complete graph and not a tree (there is no conditional independence between Y 's). The maximum likelihood structure estimate with respect to the parameters θ' of the observables in general will return a complete graph. Specifically, let $\mathbf{G} = (\mathbf{V}, \mathbf{E}_{\mathbf{G}})$ be the graph (which is complete) of the observable layer, then the distribution $p_{\dagger}(\cdot)$ is an Ising-Model distribution and it can be written as

$$p_{\dagger}(\mathbf{y}) = \frac{1}{Z'(\theta')} \exp \left\{ \sum_{(s,t) \in \mathcal{E}_{\mathbf{G}}} \theta'_{st} y_s y_t \right\}, \quad \mathbf{y} \in \{-1, 1\}^p. \quad (25)$$

Since all the edges exist in the edge set, none of the values θ'_{st} is zero. As a consequence, even asymptotically ($n \rightarrow \infty$) the maximum likelihood that estimates the parameters θ'_{st} gives a complete graph. Recall that we want to recover the structure of the hidden layer which is a tree. Thus, the maximum likelihood structure estimate directly applied on (25) is not consistent, because of the different hidden and observables' structure.

To overcome the inconsistency that is introduced by the noise, we can project the distribution $p_{\dagger}(\mathbf{y})$ to a set of tree-structured distributions and then find the maximum likelihood structure estimate. We denote the projection of $p_{\dagger}(\mathbf{y})$ onto the space of trees as $p_{\dagger}^T(\mathbf{y})$

and we call the MLE with respect to $p_{\dagger}^{\mathcal{T}}(\mathbf{y})$ as projected-MLE (PMLE). Then the following questions are natural: Is the PMLE always *consistent* with respect to structure of the hidden layer? Is the PMLE *asymptotically optimal* ($n \rightarrow \infty$)? Is the PMLE *optimal for finite values of n* ? (by optimal we mean that the sample complexity bound matches the minimax bound). We continue by answering the questions above. First we present the structural consistency and then we continue by discussing the asymptotic optimality and optimality for finite n .

Although, the PMLE is not in general consistent with structure of the hidden layer (see also related work by [Nikolakakis et al. \(2020\)](#)), for the setting of the BSC channel with i.i.d noise we do have $\hat{T}_{\text{PMLE}} \rightarrow T$ when $n \rightarrow \infty$. In fact, the projected distribution as $p_{\dagger}^{\mathcal{T}}(\mathbf{y})$ is given by

$$p_{\dagger}^{\mathcal{T}}(\mathbf{y}) \triangleq \underset{Q(\cdot) \in \mathcal{P}_T(\alpha, \beta)}{\operatorname{argmin}} D_{\text{KL}}(p_{\dagger}(\mathbf{y}) || Q(\mathbf{y})). \quad (26)$$

The proof of the claim follows by a standard argument (see also Lemma 1 and Lemma 2 by [Bresler and Karzand, 2020](#), Supplemetary material, Appendix A)) and it gives

$$D_{\text{KL}}(p_{\dagger}(\mathbf{y}) || p_{\dagger}^{\mathcal{T}}) = 1 - H(p_{\dagger}(\mathbf{y})) + \sum_{(i,j) \in \mathcal{E}} H_B \left(\frac{1 + (1 - 2q)^2 \mu_{i,j}}{2} \right). \quad (27)$$

As a consequence the projected-MLE \hat{T}_{PMLE} is

$$\hat{T}_{\text{PMLE}} = \underset{T \in \mathcal{T}}{\operatorname{argmin}} \sum_{(i,j) \in \mathcal{E}_T} H_B \left(\frac{1 + \hat{\mu}_{i,j}^{\dagger}}{2} \right) \equiv T_{\dagger}^{\text{CL}}, \quad (28)$$

and the following

$$\underset{T \in \mathcal{T}}{\operatorname{argmin}} \sum_{(i,j) \in \mathcal{E}_T} H_B \left(\frac{1 + (1 - 2q)^2 \mu_{i,j}}{2} \right) \equiv \underset{T \in \mathcal{T}}{\operatorname{argmin}} \sum_{(i,j) \in \mathcal{E}_T} H_B \left(\frac{1 + \mu_{i,j}}{2} \right) \quad (29)$$

gives that $\hat{T}_{\text{PMLE}} \equiv T_{\dagger}^{\text{CL}} \rightarrow T$ (almost surely) when $n \rightarrow \infty$. Although, the above discussion of the consistency for $n \rightarrow \infty$ shows the connection with MLE, our results for instance Theorem 3.1 shows that the Chow-Liu algorithm returns the original tree for finite n with probability $1 - \delta$.

Additionally, the PMLE is asymptotically optimal, however for finite n it may be not optimal. For our structure/predictive learning problem our bounds are asymptotically optimal (up to constants). That is, for fixed q the upper and lower bounds match as $n \rightarrow \infty$. Nevertheless for finite n the PMLE is not optimal in general. It is known that under the presence of noise the MLE approach may be non-robust and sub-optimal and extra steps should be considered including pre-processing, statistical learning of the noise by using pilot samples, and detecting and rejecting bad samples (for further information see also [Zoubir et al. \(2012, page 62\)](#) and [Nikolakakis et al. \(2020\)](#)). The reason that we consider Chow-Liu algorithm in our work is that it is computationally efficient, while its sample complexity remains logarithmic with respect to p even when noise exists. The latter makes the Chow-Liu algorithm useful in practice when only noisy observations are available. Finally, to give

further insight about the gap between the upper and lower bounds we present an example in Section H.1 (Appendix), for which perfect denoising is possible for $p \rightarrow \infty$ before running the Chow-Liu algorithm. As consequence, for $p \rightarrow \infty$ the bounds in Propositions 2 and 4 reduce to the noiseless case as they should. This example is a marginal case (since perfect denoising is not possible in general) and it affects our converse results which are universal and owe to include corner cases.

3. Main Results

The main question asked by this paper is as follows: *what is the impact of noise on the sample complexity of learning a tree-structured graphical model in order to make predictions?* This corresponds to sampling variables \mathbf{Y} generated by sampling \mathbf{X} from the model (10) and randomly flipping each sign independently with probability q . We use the Chow-Liu algorithm to estimate the hidden structure using the noise-corrupted samples. We first find upper (Theorem 5) and lower bounds (Theorem 6) on the sample complexity for exact hidden structure recovery using the Chow-Liu algorithm on noisy observations.

Secondly, we use the structure statistic to derive an accurate estimate of the hidden layer's probability distribution. The distribution estimate is computed to be accurate under the ssTV utility measure, that was introduced by [Bresler and Karzand \(2020\)](#). Furthermore, the estimator of the distribution factorizes with respect to the structure estimate, while the ssTV metric ensures that the estimated distribution is a trustworthy predictor. Theorem 7 and Theorem 8 give the sufficient and necessary sample complexity for accurate distribution estimation from noisy samples. These theorems generalize the results for the noiseless case ($q = 0$) by [Bresler and Karzand \(2020\)](#) and lead to interesting connections between structure learning on hidden models and data processing inequalities ([Raginsky, 2016](#); [Polyanskiy and Wu, 2017](#)).

The third part of the results includes Theorem 10, which gives an equivalent of Isserlis' theorem by providing closed form expressions for higher order moments of sign-valued Markov fields on trees. Based on Theorem 10 we propose a low complexity algorithm to estimate any higher order moment of the hidden variables given the estimated tree structure and estimates of the pairwise correlations (both evaluated from observations corrupted by noise).

Finally, Theorem 11 gives the sufficient number of samples for distribution estimation, when the symmetric KL divergence is considered as utility measure. These give rise to extensions of testing algorithms [Daskalakis et al. \(2018\)](#) under a hidden model setting.

3.1 Tree Structure Learning from Noisy Observations

Our goal is to learn the tree structure T of an Ising model with parameters $|\theta_{st}| \in [\alpha, \beta]$, when the nodes X_i are hidden variables and we observe $Y_i \triangleq N_i X_i$, $i \in \mathcal{V}$, where $N_i \sim \text{Rademacher}(q)$ are i.i.d, for all $i \in \mathcal{V}$ and for all $q \in [0, 1/2)$. We derive the estimated structure T_{\dagger}^{CL} by applying the Chow-Liu algorithm (Algorithm 1) ([Chow and Liu, 1968](#)).

Instead of mutual information estimates, our Chow-Liu algorithm (Algorithm 1) requires correlation estimates; these are sufficient statistics because of (14). Further, it can consistently recover the hidden structure through noisy observations. The latter is true because of the *order preserving* property of the mutual information. That is, the stochastic map-

ping $\mathbf{X} \xrightarrow{\text{BSC}(q)^p} \mathbf{Y}$ allows structure recovery of \mathbf{X} by observing \mathbf{Y} , because for any tuple $X_i, X_j, X_{i'}, X_{j'}$ such that $I(X_i; X_j) \leq I(X_{i'}, X_{j'})$, it is true that $I(Y_i; Y_j) \leq I(Y_{i'}, Y_{j'})$. The proof directly comes from (14) and (19). In addition, the monotonicity of mutual information with respect to the absolute values of correlations allows us to apply the Chow-Liu algorithm directly on the estimated correlations $\hat{\mu}_{i,j}^\dagger \triangleq 1/n_\dagger \sum_{k=1}^{n_\dagger} (Y_i)^{(k)} (Y_j)^{(k)}$. Notice that because of (19), $\hat{\mu}_{i,j}^\dagger$ can be used as an alternative of $\hat{\mu}_{i,j}$. The algorithm returns the maximum spanning tree T_\dagger^{CL} . Further discussion about the Chow-Liu algorithm is given in Section 4.1. The following theorem provides the sufficient number of samples for exact structure recovery through noisy observations.

Theorem 5 (Sufficient number of samples for structure learning) *Let \mathbf{Y} be the output of a $\text{BSC}(q)^p$, with input variable $\mathbf{X} \sim p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$. Fix a number $\delta \in (0, 1)$. If the number of samples n_\dagger of \mathbf{Y} satisfies the inequality*

$$n_\dagger \geq \frac{32 \left[1 - (1 - 2q)^4 \tanh \beta \right]}{(1 - 2q)^4 (1 - \tanh \beta)^2 \tanh^2 \alpha} \log \frac{2p^2}{\delta}, \quad (30)$$

then Algorithm 1 returns $T_\dagger^{\text{CL}} = T$ with probability at least $1 - \delta$.

Theorem 5 characterizes the finite-sample performance of the Chow-Liu estimator and by taking $n \rightarrow \infty$ we can see that Algorithm 1 is consistent in the noisy setting. As a consequence of (30) and the inequality $1 - \tanh(\beta) \geq e^{-2\beta}$, if the number of samples satisfies the following bound

$$n > C \frac{e^{2\beta}}{\tanh^2(\alpha)} \left[\mathbb{1}_{q=0} + e^{2\beta} ((1 - 2q)^{-4} - \tanh(\beta)) \mathbb{1}_{q \neq 0} \right] \log(p/\delta), \quad (31)$$

then the structure is exactly recovered with probability at least $1 - \delta$. The latter gives the statement of Theorem 1.

Complementary to Theorem 5, our next result characterizes the necessary number of samples required for exact structure recovery. Specifically, we prove a lower bound on the sample complexity that characterizes the necessary number of samples for any estimator ψ .

Theorem 6 (Necessary number of samples for structure learning) *Let \mathbf{Y} be the output of a $\text{BSC}(q)^p$, with input variable $\mathbf{X} \sim p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$. If the given number of samples of \mathbf{Y} satisfies the inequality*

$$n_\dagger < \frac{[1 - (4q(1 - q))^p]^{-1}}{16\alpha \tanh(\alpha)} e^{2\beta} \log(p), \quad (32)$$

then for any estimator ψ , it is true that

$$\inf_{\psi} \sup_{\substack{T \in \mathcal{T} \\ p(\cdot) \in \mathcal{P}_T(\alpha, \beta)}} \mathbb{P}(\psi(\mathbf{Y}_{1:n_\dagger}) \neq T) > \frac{1}{2}. \quad (33)$$

It can be shown that the right hand-side of (30) is greater than the right-hand side of (32) for any q in $[0, 1/2)$ (and for all possible values of p, β, α), by simply comparing the two terms. Theorems 5 and 6 reduce to the noiseless setting by setting $q = 0$ (Bresler and Karzand (2020)). The sample complexity is increasing with respect to q , and structure learning is always feasible as long as $q \neq 1/2$. Let n denote the required samples under a noiseless setting assumption, then for a fixed probability of exact recovery, we always need $n_{\dagger} \geq n$ because

$$\frac{[1 - (1 - 2q)^4 \tanh(\beta)]}{[(1 - 2q)^4 (1 - \tanh(\beta))]} \geq 1, \quad \forall q \in [0, \frac{1}{2}) \text{ and } \beta \in \mathbb{R}. \quad (34)$$

Furthermore,

$$\frac{1}{1 - (4q(1 - q))^p} \geq 1, \quad \forall q \in [0, 1/2) \text{ and } p \in \mathbb{N}, \quad (35)$$

the latter shows that the sample complexity in a hidden model is greater than the noiseless case ($q = 0$), for any measurable estimator (Theorem 6). When q approaches $1/2$, the sample complexity approaches infinity, $n_{\dagger} \rightarrow \infty$, and the structure learning is impossible. Theorem 6 extends Theorem 3.1 by Bresler and Karzand (2020) to our hidden model. Our results combines Bresler’s and Karzand’s method and a strong data processing inequality (SDPI) by Polyanskiy and Wu (2017, Evaluation of the BSC). Upper bounds on the symmetric KL divergence for the output distribution $p_{\dagger}(\cdot)$ can not be found in a closed form. However, by using the SDPI, we manage to capture the dependence of the bound on the parameters α, β, q and derive a non-trivial result. When $p \rightarrow \infty$, the bound becomes trivial since $\lim_{p \rightarrow \infty} 1/[1 - (4q(1 - q))^p] \rightarrow 1$, giving the classical data processing inequality (contraction of KL divergence for finite alphabets, (Raginsky, 2016; Polyanskiy and Wu, 2017)). While direct application of the SDPI is simple and provides an upper bound which is almost insensitive to q (for sufficiently large p), it introduces a gap between the lower and upper bounds. Nevertheless, it is important because it indicates a possible non-optimal performance of the classical Chow-Liu algorithm (under a hidden model). We conjecture that the sample complexity bounds in (Theorem 6 and Theorem 8) are tight only under the low temperature regime $|\theta_{i,j}| = \beta \rightarrow \infty$ for all $i, j \in \mathcal{E}$, while in general ($\theta_{i,j} \in [\alpha, \beta]$) the inequalities hold but they are not tight. For further explanation related to the gap between the upper and lower bounds see Section 2.6. The latter is a consequence of the SDPI, which is tight for the repetition code (Polyanskiy and Wu, 2017, Evaluation for the BSC, page 12). We performed extensive simulations (c.f. Figures 1, 2) that suggests that our bound does indeed accurately characterize the performance of Chow-Liu. These simulations choose $p = 100$, but our evidence shows that the dependence on q is not affected for larger ($p = 200$) or smaller ($p = 50$) values of q . We believe that the term $1/[1 - (4q(1 - q))^p]$ does not characterize the Chow-Liu algorithm, but possibly a more complicated algorithm.

3.2 Predictive Learning from Noisy Observations

In addition to recovering the structure of the hidden Ising model, we are interested in estimating the distribution $p(\cdot) \in \mathcal{P}_{\mathbf{T}}(\alpha, \beta)$ itself. If the $\mathcal{L}^{(2)}$ distance between the estimator and the true distribution is sufficiently small, then the estimated distribution is appropriate

for predictive learning because of (22). For consistency, this distribution should factorize according to the structure estimate T_{\dagger}^{CL} and for the predictive learning part, the estimate T_{\dagger}^{CL} is considered the output of the Chow-Liu algorithm (see Algorithm 1). We continue by defining the distribution estimator of $p(\cdot)$ as

$$\Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger}) \triangleq \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_{T_{\dagger}^{\text{CL}}}} \frac{1 + x_i x_j \frac{\hat{\mu}_{i,j}^{\dagger}}{(1-2q)^2}}{2}. \quad (36)$$

The estimator (36) can be defined for any $q \in [0, 1/2)$. For $q = 0$ it reduces to that in the noiseless case, since $T_{\dagger}^{\text{CL}} \equiv T^{\text{CL}}$, $\hat{\mu}_{i,j}^{\dagger} \equiv \hat{\mu}_{i,j}$, and thus $\Pi_{T_{\dagger}^{\text{CL}}}(\hat{P}^{\dagger}) \equiv \Pi_{T^{\text{CL}}}(\hat{P})$. It is also closely related to the reverse information projection onto the tree-structured Ising models (Bresler and Karzand, 2020, supplementary material, Appendix A), in the sense that

$$\Pi_T(P) = \underset{Q \in \mathcal{P}_T(\alpha, \beta)}{\text{argmin}} \mathbf{D}_{\text{KL}}(P || Q), \quad P \in \mathcal{P}_T(\alpha, \beta). \quad (37)$$

To compute $\Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})$, two sufficient statistics are required: the structure T_{\dagger}^{CL} and the set of second order moments (Chow and Liu, 1968; Bresler and Karzand, 2020), under the assumption that q is known. The next result provides a sufficient condition on the number of samples to guarantee that the $\mathcal{L}^{(2)}$ distance between the true distribution and the estimated distribution is small with probability at least $1 - \delta$.

Note that the dependence on β changes from $e^{2\beta}$ to $e^{4\beta}$ when the data are noisy $q > 0$, while for $q = 0$ our bound exactly recovers the noiseless case (Bresler and Karzand, 2020). A key component of the bound is the following function

$$\Gamma(\beta, q) \triangleq \left(\frac{1 - (1 - 2q)^2}{1 - (1 - 2q)^4 \tanh^2(\beta)} \right)^2, \quad \beta > 0 \text{ and } q \in [0, 1/2). \quad (38)$$

Further, notice that $\Gamma(\beta, q) \in [0, 1]$ for all $\beta > 0$ and $q \in [0, 1/2)$, and $\Gamma(\beta, 0) = 0$ for all $\beta > 0$. Additionally, we define the functions

$$K(\beta, q) \triangleq \frac{10(1 - \tanh^2(\beta))}{9 + (1 - 2q)^2 - \tanh^2(\beta)(1 - 2q)^2(9(1 - 2q)^2 + 1)}, \quad (39)$$

and

$$B(\beta, q) \triangleq \max \left\{ \frac{1}{K(\beta, q)}, \left(1 + 2e^{\beta} \sqrt{2(1 - q)q \tanh \beta} \right)^2 \right\}. \quad (40)$$

The latter constitute additional components of the main result that follows.

Theorem 7 Fix $\delta \in (0, 1)$ and choose $\eta > 0$. If

$$n \geq \max \left\{ \frac{512}{\eta^2(1 - 2q)^4}, \frac{1152e^{2\beta}B(\beta, q)}{(1 - 2q)^4}, \frac{48e^{4\beta}\Gamma(\beta, q)}{\eta^2} \right\} \log \left(\frac{6p^3}{\delta} \right) \quad (41)$$

then

$$\mathbb{P} \left(\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger}) \right) \leq \eta \right) \geq 1 - \delta. \quad (42)$$

(41) and the inequalities $\Gamma(\beta, q) \leq \mathbb{1}_{q \neq 0}$, $B(\beta, q) \leq (1 + 3\sqrt{q})^2 e^{2\beta \mathbb{1}_{q \neq 0}}$ give Theorem 3. We provide the proof of Theorem 3 and Theorem 7 in Section E (Appendix). As we mentioned in Section 1.5.2, the sample complexity for accurate predictive learning does not depend on α , that is, even in the high temperature regime $\alpha \rightarrow 0$ (and in contrast with the structure learning), the number of required samples does not increase.

Conversely, the following result provides the necessary number of samples for small $\mathcal{L}^{(2)}$ distance by a minimax bound, that characterizes any possible estimator ψ . In other words, it provides the necessary number of samples required for accurate distribution estimation, appropriate for predictive learning (small $\mathcal{L}^{(2)}(\cdot)$).

Theorem 8 (Necessary number of samples for inference) *Fix a number $\delta \in (0, 1)$. Choose $\eta > 0$ such that $\tanh(\alpha) + 2\eta < \tanh(\beta)$. If the given number of samples satisfies the inequality*

$$n_{\dagger} < \frac{1 - [\tanh(\alpha) + 2\eta]^2}{16\eta^2[1 - (4q(1 - q))^p]} \log p, \quad (43)$$

then for any algorithm ψ , it is true that

$$\inf_{\psi} \sup_{\substack{\mathbf{T} \in \mathcal{T} \\ \mathbf{p}(\cdot) \in \mathcal{P}_{\mathbf{T}}(\alpha, \beta)}} \mathbb{P}(\mathcal{L}^2(\mathbf{p}(\cdot), \psi(\mathbf{Y}_{1:n})) > \eta) > \frac{1}{2}.$$

Theorems 7 and 8 reduce to the noiseless setting for $q = 0$, that has been studied earlier by [Bresler and Karzand \(2020\)](#). Similarly to our structure learning results, presented previously (Theorem 5, Theorem 6), when $q \rightarrow 1/2$ we have $n_{\dagger} \rightarrow \infty$, the latter indicates that the learning task becomes impossible for $q = 1/2$.

Remark 9 *Theorem 8 requires the assumption $\alpha < \beta$. The special case $\alpha = \beta$ can be derived by applying the same proof technique of Theorem 8 combined with Theorem 3 by ([Bresler and Karzand, 2020](#), supplementary material) and the SDPI by [Polyanskiy and Wu \(2017\)](#).*

Further details and proof sketches of Theorems 7 and 8 are provided in Section 4.3.

3.3 Estimating Higher Order Moments of Signed-Valued Trees

A collection of moments is sufficient to represent completely any probability mass function. For many distributions, the first and second order moments are sufficient statistics; this is true, for instance, for the *Gaussian distribution* or the *Ising model* with unitary and pairwise interactions. Even further, in the Gaussian case, the well-known Isserlis' Theorem ([Isserlis \(1918\)](#)) gives a closed form expression for all moments of every order. As part of this work, we derive the corresponding moment expressions, *for any tree-structured Ising model*. To derive the expression of higher order moments, we first prove a key property of tree structures: for any tree structure $\mathbf{T} = (\mathcal{V}, \mathcal{E})$ and a even-sized set of nodes $\mathcal{V}' \subset \mathcal{V}$, we can partition \mathcal{V}' into $|\mathcal{V}'|/2$ pairs of nodes, such that the path along any pair is disjoint with the path of any other pair (see Appendix A, Lemma 15). We denote as $\mathcal{C}_{\mathbf{T}}(\mathcal{V}')$ the set of distinct $|\mathcal{V}'|/2$ pairs of nodes in \mathcal{V}' , such that $\text{path}(u, u') \cap \text{path}(w, w') = \emptyset$, for all

Algorithm 2 Matching Pairs

Require: Tree structure $T = (\mathcal{V}, \mathcal{E})$, any set $\mathcal{V}' \subset \mathcal{V} : |\mathcal{V}'| \in 2\mathbb{N}$

```

1:  $\mathcal{CP}_T \leftarrow \emptyset$ 
2: for  $i \in \mathcal{V}$  do
3:   if  $i \in \mathcal{V}'$  then
4:      $p(i) \leftarrow 1$ 
5:   else
6:      $p(i) \leftarrow 0$ 
7: for  $k \in [d]$  do  $\triangleright d$  is the depth of the tree
8:   Store all nodes at level  $k$  to  $L(k)$ 
9: for  $k \in [d]$  do
10:  for  $i \in L(d+1-k)$  do  $\triangleright$  Visit each of the nodes at level  $d+1-k$ 
11:    if  $p(i) = 1$  then
12:       $\mathcal{V}' \leftarrow \mathcal{V}' \setminus \{i\}$ 
13:       $\mathcal{CP}_T \leftarrow \mathcal{CP}_T \cup (i, \text{ancestor}(i))$ 
14:      if  $p(\text{ancestor}(i)) = 1$  then
15:         $\mathcal{V}' \leftarrow \mathcal{V}' \setminus \{\text{ancestor}(i)\}$ 
16:         $p(\text{ancestor}(i)) \leftarrow 0$ 
17:      else
18:         $p(\text{ancestor}(i)) \leftarrow 1$ 
19:    if  $\mathcal{V}' \equiv \emptyset$  then
20:      return  $\mathcal{CP}_T$ 
    
```

$\{u, u'\}, \{w, w'\} \in \mathcal{C}_T(\mathcal{V}')$. Let $\mathcal{CP}_T(\mathcal{V}')$ be the set of all edges in all mutually edge-disjoint paths with endpoints the pairs of nodes in \mathcal{V}' , that is,

$$\mathcal{CP}_T(\mathcal{V}') \triangleq \bigcup_{\{w, w'\} \in \mathcal{C}_T(\mathcal{V}')} \text{path}_T(w, w'). \quad (44)$$

For any tree T , the set $\mathcal{CP}_T(\mathcal{V}')$ can be computed via the Matching Pairs algorithm, Algorithm 2. By using the notation above, we can now present the equivalent of Isserlis' Theorem. The closed form expression of moments is given by the next theorem.

Theorem 10 *For any distribution of the form of (11), which factorizes according to a tree T and has support $\{-1, +1\}^p$, it is true that*

$$\mathbb{E}[X_{i_1} X_{i_2} \dots X_{i_k}] = \begin{cases} 0 & k \text{ odd} \\ \prod_{e \in \mathcal{CP}_T(i_1, i_2, \dots, i_k)} \mu_e & k \text{ even.} \end{cases} \quad (45)$$

Theorem 10 is an equivalent of Isserlis' theorem for tree-structured sign-valued distributions. Equation (45) is used later to define an estimator of higher order moments that requires two sufficient statistics: the estimated structure T_{\dagger}^{CL} and the correlation estimates $\hat{\mu}_e^{\dagger}$, for any $e \in T_{\dagger}^{\text{CL}}$. Together with the parameter q , the higher order moments completely characterize the distribution of the noisy variables of the hidden model (20). We provide the proof of Theorem 10 in Appendix A.

A similar expression to (45) has been introduced in prior work. Specifically, Algorithm 2 solves the problem of finding the optimal matching, see Definition 1 by (Bresler and Karzand, 2020, supplementary material). The evaluation of higher order moments requires an explicit expression or a way to compute the set \mathcal{CP}_T . For a given tree $T = (\mathcal{E}, \mathcal{V})$ and a set $((i_1, i_2, \dots, i_k) \subset \mathcal{V}$, there is a unique set $\mathcal{CP}_T(i_1, i_2, \dots, i_k)$ (see Appendix A, proof of Theorem 10). Given a set of edges \mathcal{E} , we show that the set \mathcal{CP}_T can be evaluated by running a matching pair algorithm. For that purpose, we provide Algorithm 2 (with complexity $(\mathcal{O}(\mathcal{E}))$) and we prove its consistency (See Appendix, Lemma 15). The latter yields to an explicit expression of higher order moments; the Theorem 10. Furthermore, it provides a concrete higher order moments estimator, that is based on the estimated structure T^{CL} (or T_{\dagger}^{CL}) and the set of estimated correlations $\{\hat{\mu}_e : e \in \mathcal{CP}_{T^{\text{CL}}}\}$.

High Order Moments Estimator: A higher order moment is the expected value of the product of the hidden tree-structured Ising model variables $\{X_i : i \in \mathcal{V}'\}$ where $\mathcal{V}' \subset \mathcal{V}$. Theorem 10 gives the closed form solution for such moments. We have the following estimator for higher order moments using only noisy observations and known q . In particular, we have

$$\hat{\mathbb{E}}[X_{i_1} X_{i_2} \dots X_{i_k}] \equiv 0, \quad k \in 2\mathbb{N} + 1, \quad (46)$$

$$\hat{\mathbb{E}}[X_{i_1} X_{i_2} \dots X_{i_k}] \triangleq \prod_{e \in \mathcal{CP}_{T_{\dagger}^{\text{CL}}}(i_1, i_2, \dots, i_k)} \frac{\hat{\mu}_e^{\dagger}}{(1 - 2q)^2}, \quad k \in 2\mathbb{N}. \quad (47)$$

The estimated structure and pairwise correlations are sufficient statistics: given those, (47) suggests a computationally efficient estimator for higher order moments. First we run the classical Chow-Liu algorithm to estimate the tree structure T_{\dagger}^{CL} , and then we run Algorithm 2 with input the estimate T_{\dagger}^{CL} to evaluate the set $\mathcal{CP}_{T_{\dagger}^{\text{CL}}}$. Thus, by estimating T_{\dagger}^{CL} , $\mathcal{CP}_{T_{\dagger}^{\text{CL}}}$ and $\hat{\mu}_e^{\dagger}$ for any $e \in \mathcal{CP}_{T_{\dagger}^{\text{CL}}}$, we can in turn estimate any higher order moment through (47). Considering the absolute estimation error, we have

$$\left| \hat{\mathbb{E}} \left[\prod_{s \in \mathcal{V}'} X_s \right] - \mathbb{E} \left[\prod_{s \in \mathcal{V}'} X_s \right] \right| \leq 2|\mathcal{V}'| \mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{P}^{\dagger}) \right). \quad (48)$$

Theorem 7 guarantees small ssTV and in combination with (48) gives an upper bound on the higher order moment estimate (47). In Section 4.4, we provide further details and discussion about Theorem 10, Algorithm 2, that computes the sets $\mathcal{CP}_T(\mathcal{V}')$, $\mathcal{CP}_{T^{\text{CL}}}(\mathcal{V}')$, and the bound on the error of estimation (48).

So far we have studied the consistency of the estimator with respect to the $\mathcal{L}^{(2)}$ metric. We are also interested in sample complexity bounds for ϕ -divergences. While general divergences may be challenging, the most widely-used is the KL-divergence, particularly in testing Ising models (Daskalakis et al., 2018). The next result gives a bound for the sufficient number of samples to guarantee a small *symmetric* KL divergence $\mathcal{S}_{\text{KL}}(P||Q) \triangleq \mathcal{D}_{\text{KL}}(P||Q) + \mathcal{D}_{\text{KL}}(Q||P)$ with high probability. For any Ising model distributions P, Q of the form (10) with respective interaction parameters θ, θ' , we have

$$\mathcal{S}_{\text{KL}}(\theta||\theta') \triangleq \mathcal{S}_{\text{KL}}(P||Q) = \sum_{s,t \in \mathcal{E}} (\theta_{st} - \theta'_{st}) (\mu_{st} - \mu'_{st}). \quad (49)$$




Figure 3: Probability of incorrect structure recovery, The theoretical bound is given by Theorem 5. The top view of the figure is Figure 1 and provides a clear comparison between the experimental and theoretical results.

Theorem 11 (Upper Bounds for the Symmetric KL Divergence) *If the number of samples n_{\dagger} of \mathbf{Y} satisfies*

$$n_{\dagger} \geq 4 \frac{\beta^2(p-1)^2}{(1-2q)^4 \eta_s^2} \log \left(\frac{p^2}{\delta} \right), \quad (50)$$

then for $p(\cdot) \in \mathcal{P}_T(\alpha, \beta)$ we have

$$\mathbb{P} \left(\mathcal{S}_{KL} \left(p(\cdot) \parallel \Pi_{T_{\dagger}^{CL}}(\hat{P}^{\dagger}) \right) \leq \eta_s \right) \geq 1 - \delta, \quad (51)$$

where T_{\dagger}^{CL} is the Chow-Liu tree defined in (52) and the estimate $\Pi_{T_{\dagger}^{CL}}(\hat{P}^{\dagger})$ is given by (36).

The asymptotic behavior of the bound in (50) was recently studied by [Daskalakis et al. \(2018\)](#). In that work, a set of testing algorithms are proposed and analyzed under the assumption of an Ising model with respect to trees and arbitrary graphs. Theorem 11 gives rise to possible extensions of testing algorithms to the hidden model setting. We consider the latter as an interesting subject for future work.

3.4 Simulations

We provide empirical results based on synthetic data to illustrate the probability of error δ as function of the cross-over probability q and the number of samples n . For the simulations




Figure 4: Estimate of the probability of the ssTV to be greater than $\eta = 0.03$. The theoretical bound is given by Theorem 7. The top view of the figure is Figure 2 and provides a clear comparison between the experimental and theoretical results.




Figure 5: Estimate of the distribution error metric ssTV as a function of q and n .

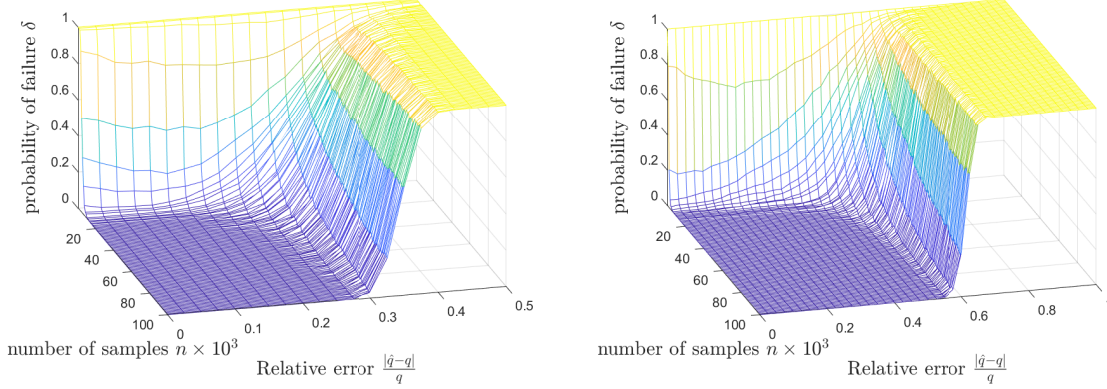


Figure 6: The probability of error in the predictive learning task for different values of n and estimation error of the parameter q . For both figures we consider $p = 31$, $\alpha = 0.2$, $\beta = 1$, $q = 0.1$ and averaging over 500 independent runs. Left: $\hat{q} \in [0.05, 0.15]$, $\eta = 0.1$, Right: $\hat{q} \in [0, 0.2]$, $\eta = 0.12$.

of this paper the original tree structure T is generated randomly where, starting from the root, we choose the parent of each new node uniformly at random among the nodes that are currently in the tree, in a sequential fashion. First, we estimate the probability of error $\mathbb{P}(T_{\dagger}^{\text{CL}} \neq T)$ (named as δ) of the structure learning problem, Figure 3. For the structure learning experiments, the number of nodes is 100, $\beta = \text{arctanh}(0.8)$, and $\alpha = \text{arctanh}(0.2)$. Further, we considering 100 Monte Carlo runs for averaging, and we plot the estimated probability of incorrect structure recovery while q and n vary. As a next step, we would like to see how well the theoretical bound of Theorem 5 matches with the experimental results. To do this we plot the top view of Figure 3 to get Figure 1. Quite remarkably, the theoretical and experimental bounds exactly match. The latter suggests that our theoretical bound that we derive, sample complexity of the Chow-Liu algorithm (Theorem 5), is indeed accurate. Second, we plot the probability of error for the predictive learning task, that is the probability of the ssTV to be greater than a positive number η (Figure 4). For the simulation part, we restrict our attention to the case that the first of three terms in the maximization of (41) is the dominant. In fact, $\eta = 0.03$, $p = 31$, while α and β are the same as the structure learning. Finally, Figure 5 presents the ssTV itself for different values of q and n . Finally, the top view of Figure 4 is Figure 2, the latter suggest that the bound of our main result, Theorem 7 is accurate.

Finally, we provide experimental results for the case of unknown q . Specifically, Figure 6 illustrates the relationship between the average probability of error and the relative error $|\hat{q} - q|/q$ for the predictive learning task. We notice that the distribution can be approximated by using an estimate \hat{q} of q even for relative error 30% or 60%.

4. Discussion

In this section, we present sketches of proofs, we compare our results with prior work, we further elaborate on Algorithm 1, Algorithm 2 and the error of higher order moment estimates. First, we discuss the convergence of the estimate T_{\dagger}^{CL} (Section 4.1). In section 4.2, we explain the connection between the hidden and noiseless settings on the tree structure learning problem. Later, in Section 4.3, we present the analysis and a sketch of proof for Theorem 7. Finally, in Section 4.4, we provide further details about Theorem 10, discussion about the Matching Pairs algorithm (Algorithm 2) and the accuracy of the proposed higher order moments estimator (47).

4.1 Estimating the Tree Structure T

In this work, the structure learning algorithm is based on the classical Chow-Liu algorithm, and is summarized in Algorithm 1. We can express its output as

$$T_{\dagger}^{\text{CL}} = \underset{T \in \mathcal{T}}{\operatorname{argmax}} \sum_{(i,j) \in \mathcal{E}_T} \left| \hat{\mu}_{i,j}^{\dagger} \right|, \quad (52)$$

(see also Section 2.6.). The difference between Algorithm 1 and the Chow-Liu algorithm of the noiseless scheme is the use of noisy observations as input, since we consider a hidden model, whereas [Bresler and Karzand \(2020\)](#) assume that observations directly from the tree-structured model are available. Further, (52) shows the consistency of the estimate T_{\dagger}^{CL} for sufficiently large n . The tree structure estimator T_{\dagger}^{CL} converges to T when $n \rightarrow \infty$, since

$$\lim_{n \rightarrow \infty} \hat{\mu}_{i,j}^{\dagger} \stackrel{\text{a.s.}}{=} c^2 \mu_{i,j}. \quad (53)$$

From (52) and (53) we have (under an appropriate metric)

$$\lim_{n \rightarrow \infty} T_{\dagger}^{\text{CL}} \stackrel{\text{a.s.}}{=} T. \quad (54)$$

Asymptotically, both T^{CL} and T_{\dagger}^{CL} converge to T , where T^{CL} denotes the structure estimate from noiseless data ($q = 0$). Most importantly, the Chow-Liu algorithm also returns the exact hidden structure with high probability given finite number of noisy samples. We provide the finite sample complexity bound in Theorem 5. For a fixed probability of exact structure recovery $1 - \delta$, more samples are required in the hidden model setting, compared to the noiseless one. Additionally, the difference of the sample complexity between the noisy and noiseless setting comes from Theorem 5 by comparing the bound for the values $q = 0$ and $q \neq 0$.

4.2 Hidden Structure Recovery and Comparison with Prior Results

Theorem 5 and Theorem 6 extend the noiseless setting ([Bresler and Karzand, 2020](#), Theorem 3.2, Theorem 3.1) to our hidden model; the noiseless results correspond to $q = 0$. In particular, in the presence of noise, *the dependence on p remains strictly logarithmic, that is, $\mathcal{O}(\log(p/\delta))$* . To make the connection between sufficient conditions more explicit, by setting $q = 0$ in (30) of Theorem 5, we retrieve the corresponding structure learning result

by [Bresler and Karzand \(2020, Theorem 3.2\)](#) *exactly*: Fix a number $\delta \in (0, 1)$. If the number of samples of \mathbf{X} satisfy the inequality

$$n \geq \frac{32}{\tanh^2 \alpha (1 - \tanh \beta)} \log \left(\frac{2p^2}{\delta} \right), \quad (55)$$

then the Chow-Liu algorithm returns $T^{\text{CL}} = T$ with probability at least $1 - \delta$. An equivalent condition of (55) is

$$\tanh \alpha \geq \frac{4\epsilon}{\sqrt{1 - \tanh \beta}} \triangleq \tau(\epsilon), \text{ and } \epsilon \triangleq \sqrt{2/n \log(2p^2/\delta)}, \quad (56)$$

the latter shows that the weight of weakest edge should satisfy the following inequality $\alpha > \text{arctanh}(4\epsilon/\sqrt{1 - \tanh \beta})$ ([Bresler and Karzand \(2020\)](#)). For the hidden model, the equivalent extended condition for the weakest edge is

$$\tanh \alpha \geq \frac{4\epsilon_{\dagger} \sqrt{1 - (1 - 2q)^4 \tanh \beta}}{(1 - 2q)^2 (1 - \tanh \beta)} \triangleq \tau^{\dagger}(\epsilon_{\dagger}), \text{ and } \epsilon_{\dagger} = \sqrt{\frac{2 \log(2p^2/\delta)}{n_{\dagger}}}, \quad (57)$$

(see Appendix C, Lemma 22) Condition (56) is retrieved through (57) for $q = 0$. Note that, for $q = 1/2$, the mutual information of the hidden and observable variables is zero, thus structure recovery is impossible.

Theorem 6 provides the necessary number of samples bound for exact structure recovery given noisy observations. In fact, it generalizes Theorem 3.1 by [Bresler and Karzand \(2020\)](#) to the hidden setting. By fixing $q = 0$, Theorem 6 recovers the noiseless case. Fix $\delta \in (0, 1)$. If the number of samples of \mathbf{X} satisfies the inequality

$$n < \frac{1}{16} e^{2\beta} [\alpha \tanh(\alpha)]^{-1} \log(p), \quad (58)$$

then for any algorithmic mapping (estimator) ψ , it is true that

$$\inf_{\psi} \sup_{\substack{T \in \mathcal{T} \\ P \in \mathcal{P}_T(\alpha, \beta)}} \mathbb{P}(\psi(X_{1:n}) \neq T) > \frac{1}{2}. \quad (59)$$

When there is no noise, $q = 0$, we retrieve the noiseless result, while for any $q \in (0, 1/2)$ the sample complexity increases since $[1 - (4q(1 - q))^p]^{-1} > 1$ in (32) and for $q \rightarrow 1/2$ the required number of samples $n_{\dagger} \rightarrow \infty$, which makes structure learning impossible. The ratio between the noiseless and noisy necessary conditions indicates the gap between the hidden model and the original (noiseless) setting, which reads

$$\frac{n_{\dagger}}{n} \leq [1 - (4q(1 - q))^p]^{-1} \leq \frac{1}{\eta_{\text{KL}}}, \quad (60)$$

(see Appendix E.). The right hand-side of (60) is the strong data processing inequality for the binary symmetric channel, which was recently developed by [Polyanskiy and Wu \(2017, Equation \(39\)\)](#). We continue by providing the main idea and the important steps of the proof of Theorem 7.

4.3 Theorem 7: A Sketch of the Proof

Recall that the indices $i, j \in \mathcal{V}$ of the quantities $\mu_{i,j}^\dagger$ and $\hat{\mu}_{i,j}^\dagger$ are pair of nodes, and in fact they can be considered as one (pair) index. For sake of space we introduce the notation μ_e^\dagger and $\hat{\mu}_e^\dagger$ for some $e \in \mathcal{E}_T$, that is consistent with our previous definition and e represents a pair of nodes. Theorem 7 guarantees that the estimated pairwise marginal distributions are close to the the original distributions by ensuring that the $\mathcal{L}^{(2)}$ is small. In this section we provide a sketch of the proof of the Theorem and we mention the main differences between the hidden model and the noiseless case (Bresler and Karzand, 2020). The intersection of three events is sufficient to guarantee that $\mathcal{L}^{(2)}$ is upper bound by $\eta > 0$:

$$E_{\dagger}^{\text{corr}}(\epsilon_{\dagger}) \triangleq \left\{ \sup_{i,j \in \mathcal{V}} |\mu_{i,j}^\dagger - \hat{\mu}_{i,j}^\dagger| \leq \epsilon_{\dagger} \right\}, \quad (61)$$

$$E_{\dagger}^{\text{strong}}(\epsilon_{\dagger}) \triangleq \left\{ \left\{ e \in \mathcal{E}_T : |\tanh \theta_e| \geq \frac{\tau^\dagger(\epsilon_{\dagger})}{(1-2q)^2} \right\} \in \mathcal{E}_{T_{\dagger}^{\text{CL}}} \right\}, \quad (62)$$

$$E_{\dagger}^{\text{cascade}}(\gamma_{\dagger}) \triangleq \left\{ \left| \prod_{e \in \text{path}_T(i,j)} \frac{\hat{\mu}_e^\dagger}{(1-2q)^2} - \prod_{e \in \text{path}_T(i,j)} \frac{\mu_e^\dagger}{(1-2q)^2} \right| \leq \gamma_{\dagger} : i, j \in \mathcal{V} \right\}, \quad (63)$$

where (57) gives the definition of $\tau^\dagger(\epsilon_{\dagger})$. The three events are equivalent events of the noiseless case, but they are modified accordingly to guarantee accurate estimation based on noisy data. The event $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$ guarantees that the error of the correlation estimates is not greater than ϵ_{\dagger} . Under the event $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$ all the strong edges are recovered by the Chow-Liu algorithm. Similarly to the noiseless setting, the event $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$ requires the Chow-Liu algorithm to recover all the strong edges, while the weak edges (those that do not satisfy the inequality in (62)) do not affect the accuracy of the predictive learning, even if the Chow-Liu algorithm fails to recover them. In contrast with structure learning, exact structure recovery is not necessary for the predictive learning task. In other words, even if α is extremely small, assume $\tanh(\alpha) \leq \tau^\dagger(\epsilon_{\dagger})/(1-2q)^2$ the required number of samples for accurate predictive learning will remain unaffected.

Under the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ the end-to-end error along paths is no greater than γ_{\dagger} . In fact, each path between two nodes of the tree can be considered a sequence of segments with strong and weak edges. The end-to-end path error is determined by the strong edge segments of the path through the parameter γ_{\dagger} for the $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ event, while the effect of weak edges parameters is controlled by the quantity $\tau^\dagger(\epsilon_{\dagger})$ (for the segmentation of the tree and the detailed proof see 26). Our goal is to find sufficient conditions on the parameters ϵ_{\dagger} and γ_{\dagger} that guarantee that the events $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$, $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ and $E^{\text{strong}}(\epsilon)$ occur with high probability.

Recall that our goal is to guarantee that the quantity $\mathcal{L}^{(2)}(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger}))$ is smaller than a fixed number $\eta > 0$ with probability at least $1 - \delta$. To do this, we follow the technique of prior work by Bresler and Karzand (2020), the triangle inequality gives

$$\mathcal{L}^{(2)}(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})) \leq \mathcal{L}^{(2)}(p(\cdot), \Pi_{T_{\dagger}^{\text{CL}}}(p(\cdot))) + \mathcal{L}^{(2)}(\Pi_{T_{\dagger}^{\text{CL}}}(p(\cdot)), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{p}_{\dagger})), \quad (64)$$

and we find the required number of samples such that each of the terms $\mathcal{L}^{(2)}(\mathbf{p}(\cdot), \Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)))$ and $\mathcal{L}^{(2)}(\Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)), \Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\hat{\mathbf{p}}_{\dagger}))$ is no greater than $\eta/2$ with probability at least $1 - \delta$. As we show the probability of the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ (Lemma 25, Appendix) and the $\mathcal{L}^{(2)}$ (between the true and estimated distribution) can be bounded by a constant uniformly over the set of all trees and is not affected by long paths. To prove these properties of the hidden model is non-trivial and ensures that the estimation error from noisy observations does not increase exponentially along paths as someone might expect. Specifically, the first quantity at the right hand-side of inequality (64) represents the loss due to graph estimation error, while the second term represents the loss due to parameter estimation error. Lemma 26 (Appendix) shows that under the event

$$E_{\dagger}(\epsilon_{\dagger}, \gamma_{\dagger}) \triangleq E_{\dagger}^{\text{corr}}(\epsilon_{\dagger}) \cap E_{\dagger}^{\text{cascade}}(\gamma_{\dagger}) \cap E_{\dagger}^{\text{strong}}(\epsilon_{\dagger}), \quad (65)$$

if

$$\gamma_{\dagger} \leq \frac{\eta}{3} \text{ and } \epsilon_{\dagger} \leq (1 - 2q)^2 e^{-\beta} \left[20 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right) \right]^{-1} \quad (66)$$

then $\mathcal{L}^{(2)}(\Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)), \Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\hat{\mathbf{p}}_{\dagger})) \leq \eta/2$. Further Lemma 27 (Appendix) shows that if

$$\epsilon_{\dagger} \leq \min \left\{ \frac{\eta}{16} (1 - 2q)^2, \frac{(1 - 2q)^2 e^{-\beta}}{24 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right)} \right\} \quad (67)$$

then $\mathcal{L}^{(2)}(\mathbf{p}(\cdot), \Pi_{\mathbf{T}_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot))) \leq \frac{\eta}{2}$ under the event $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger}) \cap E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$. Both conditions (66) and (67) should be satisfied, so it is necessary to have

$$\gamma_{\dagger} \leq \frac{\eta}{3} \text{ and } \epsilon_{\dagger} \leq \min \left\{ \frac{\eta}{16} (1 - 2q)^2, \frac{(1 - 2q)^2 e^{-\beta}}{24 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right)} \right\}. \quad (68)$$

To guarantee that the errors γ_{\dagger} and ϵ_{\dagger} are sufficient small such that (68) is satisfied, we need to make sure that the number of samples n is sufficiently large. In fact, the upper bounds on the errors translate into lower bounds on the number of samples through the concentration bounds for the events. Specifically, Lemma 18 gives a sufficient sample size to ensure that the event $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$ occurs with probability at least $1 - \delta$, Lemma 22 gives the concentration bound for the event $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$ and Lemma 25 gives the concentration bound of the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$. Lemma 18, Lemma 22 and Lemma 25 together with (68) give the final bound of the sample complexity (see the proof 28)

$$n \geq \max \left\{ \frac{512}{\eta^2 (1 - 2q)^4}, \frac{1152 e^{2\beta} B(\beta, q)}{(1 - 2q)^4}, \frac{48 e^{4\beta}}{\eta^2} \Gamma(\beta, q) \right\} \log \left(\frac{6p^3}{\delta} \right) \quad (69)$$

and its simplified but looser bound

$$n \geq \max \left\{ \frac{512}{\eta^2 (1 - 2q)^4}, \frac{1152 (1 + 3\sqrt{q})^2 e^{2\beta(1 + \mathbb{1}_{q \neq 0})}}{(1 - 2q)^4}, \frac{48 e^{4\beta}}{\eta^2} \mathbb{1}_{q \neq 0} \right\} \log \left(\frac{6p^3}{\delta} \right), \quad (70)$$

that provides the condition of Theorem 3. Although the general structure of our argument follows that of the noiseless case, the presence of noise introduces several technical challenges whose solution may be of independent interest. In the sequel, we highlight the most important aspects of our approach that do not appear in the noiseless case.

The proof of Theorem 3.3 is significantly different and includes additional steps and techniques compared with the approach by [Bresler and Karzand \(2020\)](#). Specifically, Lemma 23 is new and it is necessary for the hidden model and we use it later to prove (Lemma 25, Appendix). Lemma 24 is a non-trivial extension of the accurate estimation of edges' correlation. Although the resulting expression seems complicated is important for the proof of Lemma 25. In fact Lemma 25, the proof of the concentration bound for the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$, is significantly more complicated and longer than the noiseless model (see Appendix E by [Bresler and Karzand \(2020\)](#) for comparison). To show this result we have to consider a martingale difference sequence and evaluate upper bounds for the conditional variance and bias of that sequence. The bias is crucial for the final result because it introduces an extra term in the final bound that does not exist in the noiseless case. It is interesting that this term does not involve any parameter related to the noise and shows how the result is affected by the structural inconsistency between the hidden and the observable layer. As a consequence, the expression of the bound (170) in Lemma 25 involves two inequalities to guarantee the high-probability bound. The first inequality which introduces restrictions on the parameter Δ (see inequality 170) is an attribute of the noisy case. We continue by briefly explaining one of the main technical aspects of the proof.

To begin with, consider a path of length $d \geq 2$ in the original tree T , $X_1 - X_2 - \dots - X_{d+1}$ and we denote the edge $(k, k+1)$ as e_k , for some $k \in [d]$. Recall that $Y_k^{(i)}$ denotes the i^{th} sample of Y_k and $k \in [d+1]$. We would like derive a concentration bound of the probability of the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ (Lemma 25, Appendix). To do this, first we have to consider for all $\ell \in [n]$ and $k \in \{2, \dots, d\}$ the random variables

$$Z_k^{(\ell)} \triangleq \left(\frac{(X_k N_k X_{k+1} N_{k+1})^{(\ell)}}{(1-2q)^2} - \frac{\mu_{e_k}^{\dagger}}{(1-2q)^2} \right) \prod_{j=1}^{k-1} \frac{\hat{\mu}_{e_j}^{\dagger}}{(1-2q)^2} \prod_{j=k+1}^d \frac{\mu_{e_j}^{\dagger}}{(1-2q)^2}. \quad (71)$$

Define the martingale difference sequence (MDS) $\{\xi_k^{(i)}\}$ by setting $\xi_k^{(0)} \triangleq 0$, $\xi_k^{(1)} \triangleq Z_k^{(1)} - \mathbb{E}[Z_k^{(1)} | \hat{\mu}_{e_{k-1}}^{\dagger}, \dots, \hat{\mu}_{e_1}^{\dagger}]$, $\xi_k^{(i)} \triangleq Z_k^{(i)} - \mathbb{E}[Z_k^{(i)} | Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{e_{k-1}}^{\dagger}, \dots, \hat{\mu}_{e_1}^{\dagger}]$. Let \mathcal{F}_{i-1}^k be the σ -algebra generated by $Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{e_{k-1}}^{\dagger}, \dots, \hat{\mu}_{e_1}^{\dagger}$. Then the pair $(\xi_k^{(i)}, \mathcal{F}_i^k)_{i=1, \dots, n}$ is an MDS. In contrast with the noiseless case, *the conditional means are not zero*, which makes the problem significantly harder. To proceed, we apply a concentration bound for supermartingales (generalized Bennett's inequality) by [Fan et al. \(2012\)](#).

Secondly we have to evaluate the following expression

$$\begin{aligned} \mathbb{P}\left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = \pm 1 \middle| \hat{\mu}_{e_{k-1}}^{\dagger}, \dots, \hat{\mu}_{e_1}^{\dagger}\right) \\ = \frac{1 \pm \mu_{e_k}^{\dagger}}{2} \frac{1 - \mu_{e_{k-1}}^{\dagger} \hat{\mu}_{e_{k-1}}^{\dagger}}{1 - (\mu_{e_{k-1}}^{\dagger})^2} + \mu_{e_{k-1}}^{\dagger} \frac{1 \pm \mu_{e_k}^{\dagger} \hat{\mu}_{e_{k-1}}^{\dagger} - \mu_{e_{k-1}}^{\dagger}}{2} \frac{1 - \mu_{e_{k-1}}^{\dagger}}{1 - (\mu_{e_{k-1}}^{\dagger})^2}. \end{aligned} \quad (72)$$

In the noiseless case, the product variables $X_k^{(\ell)} X_{k+1}^{(\ell)}$ are independent, leading to a simple expression for this probability (see Lemma 13, Appendix). The closed form expression of

(72) is given by Lemma 23. Finally, the expectations $\mathbb{E} \left[Z_k^{(i)} | Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{e_{k-1}}^\dagger, \dots, \hat{\mu}_{e_1}^\dagger \right]$ are not zero, however when $n \rightarrow \infty$, they approach zero. As a consequence, a bias exists that affects the sample complexity by introducing an additional term in the bound that that does not appear in the noiseless case, the quantity $e^{4\beta}/\eta^2$ (see Equation 69 and Equation 70).

Finally, we continue by bounding the norm $\mathcal{L}^{(2)}$ between the true and estimated distribution in Appendix E. The proof of Lemma 26 shows that in the noisy setting as well, the $\mathcal{L}^{(2)}$ can be bounded by a constant uniformly over the set of all trees and it is not affected by long paths. This property of the hidden model is highly non-trivial and ensures that the estimation error from noisy observations does not increase along paths as someone might expect. Lemma 27 follows the corresponding approach of Lemma 6.1 by [Bresler and Karzand \(2020\)](#) and we provide only the required for the noisy setting differences. In Theorem 28, we combine the Lemmata of Appendices D and E, we find the appropriate choice of the parameter Δ that satisfies the necessary conditions of Lemma 25 and we derive the final sample complexity bound. For further details about the proof of the main result see Appendix, Section D and Section E.

4.4 Estimating Higher Order Moments

Our results also provide an analogue of Isserlis’ Theorem (Theorem 10) and the *Matching Pairs* algorithm, which returns the set $\mathcal{CP}_T(\mathcal{V}')$ in (45). We provide a short proof sketch for the bound on the error of estimation (48).

Proof sketch of Theorem 10: We prove that $\mathcal{C}_T(\mathcal{V}')$ always exists (when k is even) by induction (see Appendix A, Lemma 15). We define the set of edges $\mathcal{CP}_T(\mathcal{V}')$ as the union of the edge-disjoint paths⁴ $\mathcal{CP}_T(\mathcal{V}') = \cup_{w,w' \in \mathcal{C}_T(\mathcal{V}')} \text{path}(w, w')$. Combining the set $\mathcal{CP}_T(\mathcal{V}')$ together with the independent products property (see Lemma 13), we derive the final expression (see Appendix A, proof of Theorem 10). Given the tree structure T and the correlations μ_e for all $e \in \mathcal{E}$, we can calculate the higher order expectations. Notice that the collection of edge-disjoint paths \mathcal{CP}_T depends on the tree structure and as a consequence an algorithm is required to discover those paths. Different matching algorithms can be considered to find the set \mathcal{CP}_T . We propose Algorithm 2 which is simple and has low complexity of $\mathcal{O}(|\mathcal{E}|)$.

Matching Pairs Algorithm: Algorithm 2 requires as input the tree and the set of nodes $\mathcal{V}' \equiv \{i_1, \dots, i_k\} \subset \mathcal{V}$, and returns the set of edges $\mathcal{CP}_T(\mathcal{V}')$. For each node in the tree, a flag variable is assigned to each node and indicates if the corresponding node is a candidate for the final set $\mathcal{C}_T(\mathcal{V}')$ at the current step of the algorithm. The candidate nodes have to be matched with other nodes of the tree, such that the pairs generate edge-disjoint paths. Initially, the candidate nodes are the nodes of the set \mathcal{V}' . Starting from the nodes which appear in the deepest level of the tree, we “move” them to their ancestor. At each step, if two candidate nodes appear at the same point, we match them as pair, we store the pair in the set $\mathcal{CP}_T(\mathcal{V}')$ and we remove both of them from the set \mathcal{V}' . We continue until $\mathcal{V}' \equiv \emptyset$. The complexity of Algorithm 2 is $\mathcal{O}(|\mathcal{E}|)$. Finally, Theorem 10 can be extended to any forest F structure by considering the set $\mathcal{CP}_F(\mathcal{V}')$ instead of $\mathcal{CP}_T(\mathcal{V}')$, where $\mathcal{CP}_F(\mathcal{V}') \triangleq \cup_i \mathcal{CP}_{T_i}(\mathcal{V}')$ and T_i is the i^{th} connected tree of the forest.

4. By edge-disjoint paths we refer to paths with no common edges.

Estimation error of higher order moments: Inequality (48) bounds the error of estimation by the small set Total Variation (ssTV), that is guaranteed to be less than $\eta > 0$ by Theorem 7. Additionally, the bound on the error of the estimation in (48) can be found as follows

$$\left| \hat{\mathbb{E}} \left[\prod_{s \in \mathcal{V}'} X_s \right] - \mathbb{E} \left[\prod_{s \in \mathcal{V}'} X_s \right] \right| = \left| \prod_{e \in \mathcal{CP}_{\text{TCL}}(i_1, i_2, \dots, i_k)} \frac{\hat{\mu}_e^\dagger}{(1-2q)^2} - \prod_{e \in \mathcal{CP}_{\text{T}}(i_1, i_2, \dots, i_k)} \mu_e \right| \quad (73)$$

$$= \left| \prod_{e \in \mathcal{CP}_{\text{TCL}}(i_1, i_2, \dots, i_k)} \frac{\hat{\mu}_e^\dagger}{(1-2q)^2} - \prod_{e \in \mathcal{CP}_{\text{T}}(i_1, i_2, \dots, i_k)} \frac{\mu_e^\dagger}{(1-2q)^2} \right| = \left| \prod_{e \in \bigcup_{\{w, w'\} \in \mathcal{C}_{\text{TCL}}(\mathcal{V}')} \text{path}_{\text{TCL}}(w, w')} \frac{\hat{\mu}_e^\dagger}{(1-2q)^2} - \prod_{e \in \bigcup_{\{w, w'\} \in \mathcal{C}_{\text{T}}(\mathcal{V}')} \text{path}_{\text{T}}(w, w')} \frac{\mu_e^\dagger}{(1-2q)^2} \right| \quad (74)$$

$$= \left| \prod_{\{w, w'\} \in \mathcal{C}_{\text{TCL}}(\mathcal{V}')} \prod_{e \in \text{path}_{\text{TCL}}(w, w')} \frac{\hat{\mu}_e^\dagger}{(1-2q)^2} - \prod_{\{w, w'\} \in \mathcal{C}_{\text{T}}(\mathcal{V}')} \prod_{e \in \text{path}_{\text{T}}(w, w')} \frac{\mu_e^\dagger}{(1-2q)^2} \right| \leq 2|\mathcal{V}'| \mathcal{L}^{(2)} \left(p(\cdot), \Pi_{\text{TCL}}(\hat{p}_\dagger) \right), \quad (75)$$

where (73) holds due to (45) and (47), (74) comes from (44) and the last inequality (75) is being proved by [Bresler and Karzand \(2020, Lemma 1, supplementary material\)](#). Thus, if we can accurately estimate the distribution under the sense $\mathcal{L}^{(2)} \left(P, \Pi_{\text{TCL}}(\hat{p}_\dagger) \right) \leq \eta'$, for a sufficiently small positive number η' , then by using (47) and choosing $\eta' \leq \eta/(2|\mathcal{V}'|)$, Theorem 7 guarantees accurate estimates for higher order moments with probability at least $1 - \delta$.

5. Conclusion

We have considered and analyzed the problem of predictive learning on hidden tree-structures from noisy observations, using the well-known Chow-Liu algorithm. In particular, we derived sample complexity guarantees for exact structure learning and marginal distributions estimation. Our bounds extend prior work (see [Bresler and Karzand \(2020\)](#)) to the hidden model, by introducing the cross-over probability q of the $\text{BSC}(q)^p$. Our results exactly reduce to the noiseless setting when $q = 0$, and the explicit expressions of the bounds are also continuous functions of q . Additionally, by applying a graph property for tree structures and a probabilistic property for Ising models, we derived an equivalent of the well-known

Isserlis’ theorem for Gaussian distributions, which yields to a consistent high-order moments estimator for Ising models. Further, we considered simulations based on synthetic data to validate our theoretical results. Our theoretical bounds exactly match with the experiment, indicating that our results correctly characterize the dependence on the model parameters.

Our results show that the estimated structure statistic T_{\dagger}^{CL} is essential for successful statistical inference on the hidden (or observable) layer, while the sample complexity with respect to number of nodes and probability of error remains strictly logarithmic, as in the noiseless case. Our hidden setting constitutes a first step towards more technically challenging and potentially more realistic statistical models, such as, for instance, structure and distribution learning when the noise is generated by an erasure channel, or when the underlying hidden tree structured distribution has a larger, or even uncountable, support.

Appendix A. Preliminaries and Outline of Proof

The chart in Figure 7 shows the various dependencies of the Lemmata and intermediate results either considered or developed in this paper, and the resulting Theorems. The proofs can be found in the corresponding section of the Appendix.

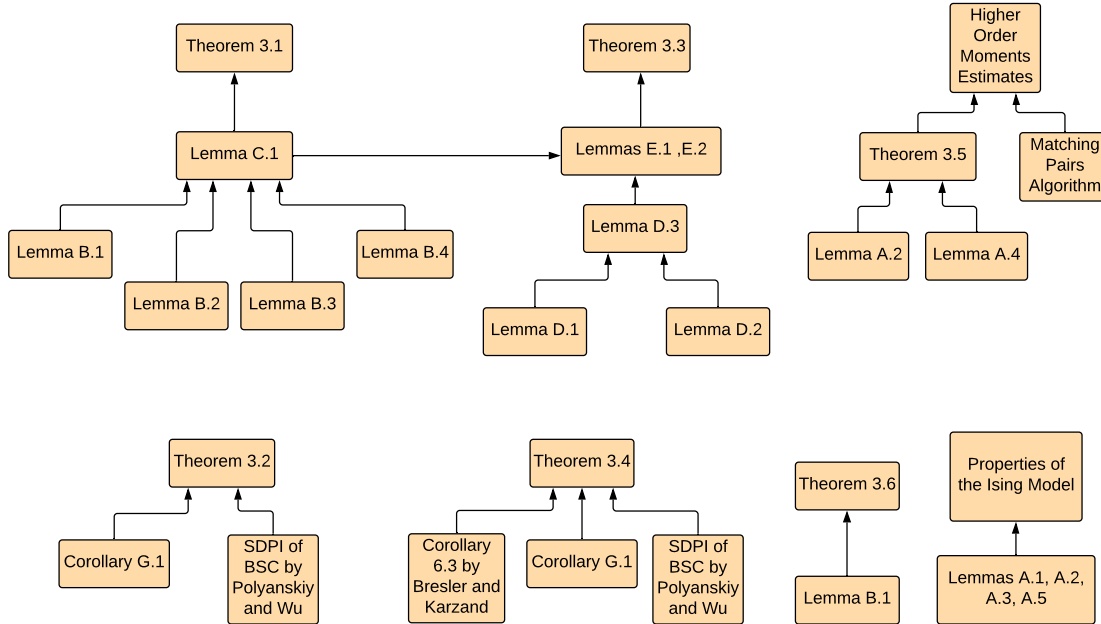


Figure 7: Stream mapping of the results

For completeness, we start with some properties that hold for any distribution with support $\{-1, +1\}^p$ and tree-structured graphical model (Lauritzen, 1996). Later we derive explicit formulas for the Ising model (10).

Lemma 12 Any distribution $p(\mathbf{x})$ with respect to a forest $F = (\mathcal{V}, \mathcal{E})$, where $\mathbf{x} \in \{-1, 1\}^p$ and uniform marginals $\mathbb{P}(X_i = \pm 1) = 1/2$, for all $i \in \mathcal{V}$ can be expressed as

$$p(\mathbf{x}) = \frac{1}{2} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{2}. \quad (76)$$

Proof We prove the result for an arbitrary tree $T = (\mathcal{V}, \mathcal{E})$ and then we extend it to any forest structure by applying cuts to \mathcal{E} . The distribution factorizes according to the tree structure T and under the assumption of no external field (uniform marginal distributions), we have

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) = \prod_{i \in V} p(x_i) \prod_{(i,j) \in \mathcal{E}} \frac{p(x_i, x_j)}{p(x_i)p(x_j)} = 2^{p-2} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{4} \quad (77)$$

$$= \frac{1}{2} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{2}. \quad (78)$$

(77) holds since the joint distribution of any pair (X_i, X_j) of distinct nodes $i, j \in \mathcal{V}$ is

$$p(x_i, x_j) = \mathbb{E}[\mathbf{1}_{X_i=x_i} \mathbf{1}_{X_j=x_j}] = \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{4}. \quad (79)$$

By setting $\mathbb{E}[X_i X_j] = 0$ for some $(i, j) \in \mathcal{E}$ we derive the distribution with respect to a forest generated by cutting the edge (i, j) of T . \blacksquare

In Lemma 13 we prove two fundamental properties of the model, the independence of the random variables $\{X_i X_j : (i, j) \in \mathcal{E}\}$ and the *correlation decay property* (CDP). To the best of our knowledge, these properties are known but there is no reference for the corresponding proofs in the literature.

Lemma 13 Let \mathbf{X} be a random binary vector in $\{-1, +1\}^p$ drawn according to a forest-structured distribution $p(\cdot)$ with uniform marginal distributions on each entry X_i for $i \in [p]$. Then the elements of the collection of $|\mathcal{E}|$ random variables $\{X_i X_j : (i, j) \in \mathcal{E}\}$, are independent. Furthermore, we have

$$\mathbb{E}[X_i X_j] = \prod_{e \in \text{path}(i,j)} \mu_e, \quad (80)$$

so the Correlation Decay Property (CDP) holds since $|\mu_e| \leq 1$ for all $e \in \mathcal{E}$.

Proof Let $(i_r)_{r=1}^p$ be an arbitrary permutation of $\mathcal{L} = \{1, 2, \dots, p\}$. Notice that the singletons $\{i_r\}$, $r = 1, \dots, p$ form a partition of \mathcal{L} . Then, the set of edges \mathcal{E} is defined as

$$\mathcal{E} = (i_r, j_r)_{r=2}^p, \quad \text{and } j_1 = \emptyset, \quad j_r \in \{i_1, \dots, i_{r-1}\} \subset \mathcal{L}. \quad (81)$$

(81) defines a tree $T = (\mathcal{V}, \mathcal{E})$ with root the node i_1 (since $j_1 = \emptyset$). For the first part, it is sufficient to show that for any $\{c_r : r = 2, 3, \dots, p\} \in \{-1, +1\}^{p-1}$, the following holds

$$\mathbb{P}\left(\bigcap_{r=2}^p \{x_{i_r} x_{j_r} = c_r\}\right) = \prod_{r=2}^p \mathbb{P}(x_{i_r} x_{j_r} = c_r). \quad (82)$$

We have

$$\begin{aligned}
 \mathbb{P}\left(\bigcap_{r=2}^p \{X_{i_r} X_{j_r} = c_r\}\right) &= \sum_{\mathbf{x}: x_{i_r} x_{j_r} = c_r |_{r=2}^p} p(\mathbf{x}) \\
 &= \sum_{\mathbf{x}: x_{i_r} x_{j_r} = c_r |_{r=2}^p} \frac{1}{2} \prod_{r=2}^p \frac{1 + x_{i_r} x_{j_r} \mathbb{E}[X_{i_r} X_{j_r}]}{2} \quad (83) \\
 &= \sum_{\mathbf{x}: x_{i_r} = c_r x_{j_r} |_{r=2}^p} \frac{1}{2} \prod_{r=2}^p \frac{1 + x_{i_r} x_{j_r} \mathbb{E}[X_{i_r} X_{j_r}]}{2} \\
 &= \sum_{x_{i_1} \in \{-1, +1\}} \frac{1}{2} \prod_{r=2}^p \frac{1 + c_r \mathbb{E}[X_{i_r} X_{j_r}]}{2} = \prod_{r=2}^p \mathbb{P}(X_{i_r} X_{j_r} = c_r), \quad (84)
 \end{aligned}$$

(83) comes from (81) and Lemma 12 and the last from (79). For the second part of the statement note that for all $i, j \in \mathcal{V}$ there exists a unique path $\{i, k_1, k_2, \dots, k_\ell, j\}$ from i to j . Define the variable $\mathbf{1}_{(i,j)} \triangleq (X_{k_1} X_{k_1})(X_{k_2} X_{k_2}) \dots (X_{k_\ell} X_{k_\ell})$, which is equal to 1 almost surely, since $\mathbf{X} \in \{-1, +1\}^p$.⁵ Then, we have

$$\begin{aligned}
 \mathbb{E}[X_i X_j] &= \mathbb{E}[X_i \mathbf{1}_{(i,j)} X_j] \\
 &= \mathbb{E}[X_i (X_{k_1} X_{k_1})(X_{k_2} X_{k_2}) \dots (X_{k_\ell} X_{k_\ell}) X_j] \\
 &= \mathbb{E}[X_i X_{k_1}] \left(\prod_{m=1}^{\ell-1} \mathbb{E}[X_{k_m} X_{k_{m+1}}] \right) \mathbb{E}[X_\ell X_j] = \prod_{e \in \text{path}(i,j)} \mu_e, \quad (85)
 \end{aligned}$$

and (85) comes from (82) and completes the proof. \blacksquare

The next lemma relates the pairwise correlations to the parameters of the Ising model.

Lemma 14 *An equivalent expression of (10) is the following*

$$p(\mathbf{x}) = \frac{\prod_{(i,j) \in \mathcal{E}} [1 + x_i x_j \tanh(\theta_{ij})]}{\sum_{\mathbf{x}} \prod_{(i,j) \in \mathcal{E}} [1 + x_i x_j \tanh(\theta_{ij})]} \quad \mathbf{x} \in \{-1, 1\}^p. \quad (86)$$

Further, for a tree-structure Ising model $\mathbb{E}[X_i, X_j] = \tanh(\theta_{ij})$, for all $(i, j) \in \mathcal{E}$.

Proof We can write $\exp(\theta_{ij} x_i x_j)$ as

$$\begin{aligned}
 \exp(\theta_{ij} x_i x_j) &= \frac{\exp(\theta_{ij} x_i x_j) + \exp(-\theta_{ij} x_i x_j)}{2} + \frac{\exp(\theta_{ij} x_i x_j) - \exp(-\theta_{ij} x_i x_j)}{2} \\
 &= \frac{\exp(\theta_{ij}) + \exp(-\theta_{ij})}{2} + x_i x_j \frac{\exp(\theta_{ij}) - \exp(-\theta_{ij})}{2} \quad (87)
 \end{aligned}$$

$$= \cosh(\theta_{ij}) [1 + x_i x_j \tanh(\theta_{ij})], \quad (88)$$

5. $\mathbf{1}_{(\cdot)}$ should not be confused with $\mathbb{1}_A$, where the last denotes the indicator function of a set A .

(87) holds because $x_i x_j \in \{-1, +1\}$. The partition function can be written as

$$\begin{aligned} Z(\theta) &= \sum_{\mathbf{x}} \prod_{(i,j) \in \mathcal{E}} \exp(\theta_{ij} x_i x_j) \\ &= \sum_{\mathbf{x}} \prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij}) [1 + x_i x_j \tanh(\theta_{ij})] \\ &= \prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij}) \sum_{\mathbf{x}} \prod_{(i,j) \in \mathcal{E}} [1 + x_i x_j \tanh(\theta_{ij})] = 2^p \prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij}). \end{aligned} \quad (89)$$

Notice that $\sum_{\mathbf{x}} \prod_{(i,j) \in \mathcal{E}} [1 + x_i x_j \tanh(\theta_{ij})] = 2^p$ under the tree-structure assumption. Then

$$\mathbb{P}(\mathbf{X} = \mathbf{x}) = \frac{\prod_{(i,j) \in \mathcal{E}} \exp(\theta_{ij} x_i x_j)}{Z(\theta)} = \frac{\prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij}) [1 + x_i x_j \tanh(\theta_{ij})]}{2^p \prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij})} \quad (90)$$

$$= \frac{1}{2} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \tanh(\theta_{ij})}{2}, \quad (91)$$

(88) and (89) give (90) and $|\mathcal{E}| = p - 1$ gives (91). Finally

$$\mathbb{E}[X_i X_j] \stackrel{(10)}{=} \frac{\partial \ln Z(\theta)}{\partial \theta_{ij}} \stackrel{(89)}{=} \frac{\partial \ln \left[2^p \prod_{(i,j) \in \mathcal{E}} \cosh(\theta_{ij}) \right]}{\partial \theta_{ij}} = \tanh(\theta_{ij}), \quad \forall (i, j) \in \mathcal{E}, \quad (92)$$

and the latter gives the second part of the Lemma. \blacksquare


Lemma 15 *Let \mathcal{V}' be a set of nodes such that $\mathcal{V}' \subset \mathcal{V}$ and $|\mathcal{V}'| \in 2\mathbb{N}$. Then it exists a set $\mathcal{C}_T(\mathcal{V}')$ of $|\mathcal{V}'|/2$ pairs of nodes of \mathcal{V}' , such that any two distinct pairs $(w, w'), (v, v')$ in $\mathcal{C}_T(\mathcal{V}')$ are pairwise disjoint (their paths have no common edge), that is,*

$$\text{path}_T(w, w') \cap \text{path}_T(v, v') = \emptyset, \quad \forall (w, w'), (v, v') \in \mathcal{C}_T(\mathcal{V}') : (w, w') \neq (v, v'). \quad (93)$$

Proof We prove the existence of $\mathcal{C}_T(\mathcal{V}')$ by contradiction. Assume that the two distinct paths $\text{path}_T(w, w'), \text{path}_T(u, u')$ share at least one edge. Let their common sub-path be $\text{path}_T(z, z')$, Figure 8 and note that z and z' do not necessarily differ from w, w', u, u' . Notice that the common sub-path is unique (acyclic graph). Then we can always consider the permutation of the endpoints which gives the edge-disjoint paths $\text{path}_T(w, u)$ and $\text{path}_T(w', u')$. Now the paths $\text{path}_T(w, u)$ and $\text{path}_T(w', u')$ are disjoint, however it is possible that one of them or both, contain sub-paths with common edges. Then, we similarly proceed by removing the common sub-paths as previously. The set of common edges strictly decreases through the process, which terminates when there are only paths with no common edge. \blacksquare

Theorem 16 (Theorem 10) *Assume $\mathbf{X} \sim p(\mathbf{x}) \in \mathcal{P}_T(\alpha, \beta)$, $\{i_1, i_2, \dots, i_k\} \subset \mathcal{V}$, then*

$$\mathbb{E}[X_{i_1} X_{i_2} \dots X_{i_k}] = \begin{cases} \prod_{e \in \mathcal{C}_T(i_1, i_2, \dots, i_k)} \mu_e, & \forall k \in 2\mathbb{N} \\ 0, & \forall k \in 2\mathbb{N} + 1. \end{cases} \quad (94)$$


 Figure 8: Proof of the existence of $\mathcal{C}_T(\mathcal{V}')$, Lemma 15

Recall that the set of edges $\mathcal{CP}_T(i_1, \dots, i_k)$ is a collection of $k/2$ edge-disjoint paths with endpoints pairs of the nodes i_1, \dots, i_k for each path. Given a tree structure T , $\mathcal{CP}_T(i_1, \dots, i_k)$ is found by running Algorithm 2 on T .

Proof Even k . We proceed by showing that the Algorithm 2 returns the unique set \mathcal{CP}_T . When $k = 2$ the expression is proved in Lemma 13. For $k > 2$ we proceed by using Lemmas 13 and 15. For all $i, j \in \mathcal{V}$ there exists a unique path $\{i, k_1, k_2, \dots, k_\ell, j\}$ from i to j . Define as previously the variable $\mathbf{1}_{(i,j)} \triangleq (X_{k_1} X_{k_1})(X_{k_2} X_{k_2}) \dots (X_{k_\ell} X_{k_\ell})$, which is equal to 1 almost surely, and define the set of nodes $\mathcal{V}' \triangleq \{i_1, i_2, \dots, i_k\}$. Without loss of generality we assume that the variables in the product $X_{i_1} X_{i_2} \dots X_{i_k}$ are ordered such such that the pairs $X_{i_j}, X_{i_{j+1}}$ for all $j \in \{1, 3, 5, \dots, k-1\} \triangleq [k-1]^{\text{odd}}$ form edge-disjoint paths (Lemma 15), in other words

$$\text{path}(i_j, i_{j+1}) \cap \text{path}(i_{j'}, i_{j'+1}) = \emptyset, \forall j \neq j' \in [k-1]^{\text{odd}}. \quad (95)$$

Then, we have

$$\mathbb{E}[X_{i_1} X_{i_2} \dots X_{i_k}] = \mathbb{E}[X_{i_1} \mathbb{1}_{(i_1, i_2)} X_{i_2} X_{i_3} \mathbb{1}_{(i_3, i_4)} X_{i_4} \dots X_{i_{k-1}} \mathbb{1}_{(i_{k-1}, i_k)} X_{i_k}] \quad (96)$$

$$= \prod_{j \in [k-1]^{\text{odd}}} \mathbb{E}[X_{i_j} \mathbb{1}_{(i_j, i_{j+1})} X_{i_{j+1}}] \quad (97)$$

$$= \prod_{j \in [k-1]^{\text{odd}}} \prod_{e \in \text{path}(i_j, i_{j+1})} \mu_e \quad (98)$$

$$= \prod_{e \in \mathcal{CP}_T(i_1, i_2, \dots, i_k)} \mu_e, \quad (99)$$

where (96) and (97) come from (80), and (98) holds because of (95).

Odd k . Lemma 12 gives $p(\mathbf{x}) = 2^{-p} \prod_{(i,j) \in \mathcal{E}} (1 + x_i x_j \mathbb{E}[X_i X_j])$. Then

$$\mathbb{E}[X_{i_1} X_{i_2} \dots X_{i_k}] = \frac{1}{2} \sum_{\mathbf{x} \in \{-1, +1\}^{i_k}} x_{i_1} x_{i_2} \dots x_{i_k} \prod_{(i,j) \in \mathcal{E}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{2} = 0, \quad (100)$$

gives the second part of (94). ■

Lemma 17 *The mutual information of $X_i, X_j \in \{-1, +1\}$ is symmetric function of the correlation $\mathbb{E}[X_i X_j]$ and increasing with respect to $|\mathbb{E}[X_i X_j]|$,*

$$I(X_i, X_j) = \frac{1}{2} \log_2 \left((1 - \mathbb{E}[X_i X_j])^{1 - \mathbb{E}[X_i X_j]} (1 + \mathbb{E}[X_i X_j])^{1 + \mathbb{E}[X_i X_j]} \right). \quad (101)$$

The proof can be derived through the definition of $I(X_i, X_j)$ and the expression (79), under the assumption of uniform marginal distributions.

Appendix B. Bounding the Probability of Mis-Estimating Correlations

The following lemma bounds the probability that the estimated pairwise correlations in the graph deviate from their true values. This follows from standard concentration of measure arguments.

Lemma 18 *Fix $\delta > 0$. Then for any $\epsilon_{\dagger} > 0$, if*

$$n_{\dagger} \geq 2 \log(p^2/\delta) / \epsilon_{\dagger}^2, \quad (102)$$

then the event $E_{\dagger}^{corr}(\epsilon_{\dagger})$ defined in (61) holds with high probability:

$$\mathbb{P}(E_{\dagger}^{corr}(\epsilon_{\dagger})) \geq 1 - \delta = 1 - p^2 \exp\left(\frac{-n_{\dagger} \epsilon_{\dagger}^2}{2}\right). \quad (103)$$

Proof Let $Z_{\dagger}^{(i)}$ be the i th sample of $Z_{\dagger} = Y_w Y_{\bar{w}} = N_w X_w N_{\bar{w}} X_{\bar{w}}$. Then $\hat{\mu}_{i,j}^{\dagger} = \frac{1}{n_{\dagger}} \sum_{i=1}^{n_{\dagger}} Z_{\dagger}^{(i)} = \frac{1}{n_{\dagger}} \sum_{i=1}^{n_{\dagger}} N_w^{(i)} X_w^{(i)} N_{\bar{w}}^{(i)} X_{\bar{w}}^{(i)}$ for all $i \neq j \in \mathcal{V}$. Then Hoeffding's inequality and union bound over all pairs of nodes $\binom{p}{2} < p^2/2$ give (103). \blacksquare

For the rest of the paper we consider $\epsilon_{\dagger} = \sqrt{2 \log(2p^2/\delta) / n_{\dagger}}$, which satisfies Lemma 18. We apply Lemmata 19, 20, 21 to Lemma 22 to bound the required number of samples for exact structure recovery using noisy observations of the hidden model. To analyze the error event we use the ‘‘Two trees lemma’’ of [Bresler and Karzand \(2020, Appendix F, supplementary material\)](#). Informally, if two maximum spanning trees T, T' differ in how a pair of nodes are connected then there exists at least one edge in \mathcal{E}_T which does not exist in $\mathcal{E}_{T'}$ and vice versa. Lemma 19 characterizes errors in the Chow-Liu in terms of correlations.

Lemma 19 *Suppose the error event $\{T \neq T_{\dagger}^{CL}\}$ holds and let $f \triangleq (w, \bar{w})$ be an edge such that $f \in T$ and $f \notin T_{\dagger}^{CL}$. Then there exists an edge $g \triangleq (u, \bar{u}) \in T_{\dagger}^{CL}$ and $g \notin T$ such that $f \in \text{path}_T(u, \bar{u})$ and $g \in \text{path}_{T_{\dagger}^{CL}}(w, \bar{w})$ and*

$$\left(\sum_{i=1}^{n_{\dagger}} Z_{f,u,\bar{u}}^{(i)} \right) \left(\sum_{i=1}^{n_{\dagger}} M_{f,u,\bar{u}}^{(i)} \right) < 0, \quad (104)$$

where $Z_{f,u,\bar{u}} \triangleq Y_w Y_{\bar{w}} - Y_u Y_{\bar{u}}$ and $M_{f,u,\bar{u}} \triangleq Y_w Y_{\bar{w}} + Y_u Y_{\bar{u}}$.

Proof Using similar approaches to the procedures as in (Bresler and Karzand, 2020, Lemmata 8.2, 8.3) we have that the condition $|\hat{\mu}_f^\dagger| \leq |\hat{\mu}_g^\dagger|$ implies

$$\begin{aligned}
 0 &\geq \left| \hat{\mu}_f^\dagger \right|^2 - \left| \hat{\mu}_g^\dagger \right|^2 \\
 &= \left(\hat{\mu}_f^\dagger - \hat{\mu}_g^\dagger \right) \left(\hat{\mu}_f^\dagger + \hat{\mu}_g^\dagger \right) \\
 &= \frac{1}{n_\dagger^2} \left(\sum_{i=1}^{n_\dagger} N_w^{(i)} X_w^{(i)} N_{\bar{w}}^{(i)} X_{\bar{w}}^{(i)} - N_u^{(i)} X_u^{(i)} N_{\bar{u}}^{(i)} X_{\bar{u}}^{(i)} \right) \\
 &\quad \times \left(\sum_{i=1}^{n_\dagger} N_w^{(i)} X_w^{(i)} N_{\bar{w}}^{(i)} X_{\bar{w}}^{(i)} + N_w^{(i)} X_u^{(i)} N_{\bar{u}}^{(i)} X_{\bar{u}}^{(i)} \right) \\
 &= \frac{1}{n_\dagger^2} \left(\sum_{i=1}^{n_\dagger} Z_{f,u,\bar{u}}^{(i)} \right) \left(\sum_{i=1}^{n_\dagger} M_{f,u,\bar{u}}^{(i)} \right). \tag{105}
 \end{aligned}$$

■

Setting

$$\epsilon_\dagger \triangleq \sqrt{\frac{2 \log(2p^2/\delta)}{n_\dagger}} \tag{106}$$

$$\tau^\dagger \triangleq \frac{4\epsilon_\dagger \sqrt{1 - (1 - 2q)^4 \tanh \beta}}{1 - \tanh \beta} \tag{107}$$

$$\mu_e \triangleq \mathbb{E}[X_w X_{\bar{w}}], \tag{108}$$

we have that μ_A is defined through the following relationship

$$\mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}] = \mu_e(1 - \mu_A), \tag{109}$$

and

$$\mu_A^\dagger \triangleq (1 - 2q)^4 \mu_A. \tag{110}$$

In Lemmata 20, 21 we derive two concentration of measure inequalities for the variables $Z_{f,u,\bar{u}}^{(i)}$, $M_{f,u,\bar{u}}^{(i)}$. In fact, we have that the event

$$\mathbb{E}_Z \triangleq \left\{ \left| \sum_{i=1}^{n_\dagger} Z_{e,u,\bar{u}}^{(i)} - n_\dagger \mathbb{E}[Z_{e,u,\bar{u}}] \right| \leq n_\dagger \max \left\{ 8\epsilon_\dagger^2, 4\epsilon_\dagger \sqrt{1 - \mu_A^\dagger} \right\} : \forall e \in \mathcal{E} \text{ and } \forall u, \bar{u} \in \mathcal{V} \right\} \tag{111}$$

happens with probability at least $1 - \frac{\delta'}{2}$ and the event

$$\mathbb{E}_M \triangleq \left\{ \left| \sum_{i=1}^{n_\dagger} M_{e,u,\bar{u}}^{(i)} - n_\dagger \mathbb{E}[M_{e,u,\bar{u}}] \right| \leq n_\dagger \max \left\{ 8\epsilon_\dagger^2, 4\epsilon_\dagger \sqrt{1 + \mu_A^\dagger} \right\} : \forall e \in \mathcal{E} \text{ and } \forall u, \bar{u} \in \mathcal{V} \right\} \tag{112}$$

happens with probability at least $1 - \frac{\delta''}{2}$. The parameters ϵ_{\dagger} and μ_A , defined below, are decreasing functions of n_{\dagger} . Finally, we apply the union bound to guarantee that the event $E_Z \cup E_M$ happens with probability at least $1 - \delta$, where $\frac{\delta'}{2} + \frac{\delta''}{2} \leq 2 \max\{\frac{\delta'}{2}, \frac{\delta''}{2}\} \triangleq \delta$. The union bound is first applied over all tuples (w, \bar{w}, u, \bar{u}) in Lemmata 20 and 21 and then for the events E_Z and E_M .

Lemma 20 *Fix $\delta > 0$ and let ϵ_{\dagger} be given by (106). For all pairs of vertices $u, \bar{u} \in V$ and edges $e = (w, \bar{w})$ in the path $\text{path}_T(u, \bar{u})$ from u to \bar{u} , given n_{\dagger} samples $Z_{e,u,\bar{u}}^{(1)}, Z_{e,u,\bar{u}}^{(2)}, \dots, Z_{e,u,\bar{u}}^{(n)}$ of $Z_{e,u,\bar{u}} = Y_w Y_{\bar{w}} - Y_u Y_{\bar{u}}$, it is true that*

$$\mathbb{P} \left(\left| \sum_{i=1}^{n_{\dagger}} Z_{e,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E}[Z_{e,u,\bar{u}}] \right| \leq n_{\dagger} \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1 - \mu_A^{\dagger}} \right\} \right) \geq 1 - \frac{\delta}{2}, \quad (113)$$

where $A = \text{path}_T(u, \bar{u}) \setminus \{e\}$.

Proof The proof is an application of Bernstein's inequality. First, it is true that

$$\begin{aligned} Z_{e,u,\bar{u}} &= X_w N_w X_{\bar{w}} N_{\bar{w}} - N_u X_u N_{\bar{u}} X_{\bar{u}} \\ &= N_w X_w N_{\bar{w}} X_{\bar{w}} (1 - N_w X_w N_{\bar{w}} X_{\bar{w}} N_u X_u N_{\bar{u}} X_{\bar{u}}). \end{aligned} \quad (114)$$

Then,

$$\begin{aligned} \mathbb{E}[Z_{e,u,\bar{u}}] &= (1 - 2q)^2 \mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}] \\ &= (1 - 2q)^2 \mu_e (1 - \mu_A) \end{aligned} \quad (115)$$

$$\begin{aligned} \text{Var}(Z_{e,u,\bar{u}}) &= \mathbb{E}[(Z_{e,u,\bar{u}})^2] - \mathbb{E}[Z_{e,u,\bar{u}}]^2 \\ &= \mathbb{E}[(X_w N_w X_{\bar{w}} N_{\bar{w}} - N_u X_u N_{\bar{u}} X_{\bar{u}})^2] - \left[(1 - 2q)^2 \mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}] \right]^2 \\ &= \mathbb{E}[1 + 1 - 2X_w N_w X_{\bar{w}} N_{\bar{w}} N_u X_u N_{\bar{u}} X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}]^2 \\ &= 2 - 2\mathbb{E}[X_w N_w X_{\bar{w}} N_{\bar{w}} N_u X_u N_{\bar{u}} X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}]^2 \\ &= 2 - 2(1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} X_u X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} - X_u X_{\bar{u}}]^2 \\ &= 2 - 2(1 - 2q)^4 \mu_A - (1 - 2q)^4 (\mu_e (1 - \mu_A))^2 \\ &= 2 - (1 - 2q)^4 \left[2\mu_A + \mu_e^2 (1 - \mu_A)^2 \right]. \end{aligned} \quad (116)$$

Using the expressions for the mean and the variance, we apply Bernstein's inequality (Bennett, 1962) for the noisy setting: for all $i \in [n_{\dagger}]$ we have $|Z_{e,u,\bar{u}}^{(i)} - \mathbb{E}[Z_{e,u,\bar{u}}]| \leq M$ almost surely. Then, Bernstein's inequality gives, for all $t > 0$

$$\mathbb{P} \left[\left| \sum_{i=1}^{n_{\dagger}} Z_{e,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E}[Z_{e,u,\bar{u}}] \right| \leq t \right] \geq 1 - 2 \exp \left(- \frac{t^2}{2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) + \frac{2}{3} M t} \right), \quad \forall t > 0. \quad (117)$$

Choose a $\delta > 0$ and find t such that

$$\delta/2 = 2 \exp \left(- \frac{t^2}{2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) + \frac{2}{3} M t} \right).$$

After some algebra, we have

$$\log \frac{4}{\delta} = \frac{t^2}{2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) + \frac{2}{3}Mt}$$

From this we can solve for t :

$$\begin{aligned} 0 &= t^2 - \frac{2}{3}Mt \log \frac{4}{\delta} - 2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) \log \frac{4}{\delta} \\ t_{1,2} &= \frac{\frac{2}{3}M \log \frac{4}{\delta} \pm \sqrt{\left(\frac{2}{3}M \log \frac{4}{\delta}\right)^2 + 8n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) \log \frac{4}{\delta}}}{2} \\ &= \frac{1}{3}M \log \frac{4}{\delta} \pm \sqrt{\left(\frac{1}{3}M \log \frac{4}{\delta}\right)^2 + 2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) \log \frac{4}{\delta}}. \end{aligned} \quad (118)$$

Since $t > 0$, we have, setting $M = 4$:

$$t = \frac{4}{3} \log \frac{4}{\delta} + \sqrt{\left(\frac{4}{3} \log \frac{4}{\delta}\right)^2 + 2n_{\dagger} \text{Var}(Z_{e,u,\bar{u}}) \log \frac{4}{\delta}}. \quad (119)$$

If the probability of the union

$$\bigcup_{\forall u, \bar{u}, w, \bar{w}: (w, \bar{w}) \in \text{path}_T(u, \bar{u})} \left\{ \left| \sum_{i=1}^{n_{\dagger}} Z_{e,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E}[Z_{e,u,\bar{u}}] \right| \geq t \right\}$$

is at most $\frac{\delta}{2p^3}$, then the union bound gives probability at most $\frac{\delta}{2}$. Also,

$$\begin{aligned} \text{Var}(Z_{e,u,\bar{u}}) &= 2 - (1 - 2q)^4 \left[2\mu_A + \mu_e^2 (1 - \mu_A)^2 \right] \\ &= 2 - (1 - 2q)^4 2\mu_A - (1 - 2q)^4 \mu_e^2 (1 - \mu_A)^2 \\ &\leq 2 - (1 - 2q)^4 2\mu_A + 0 \\ &= 2 \left(1 - (1 - 2q)^4 \mu_A \right) \\ &= 2 \left(1 - \mu_A^{\dagger} \right). \end{aligned} \quad (120)$$

From (119) and (120), we have

$$\begin{aligned} t &= \frac{4}{3} \log \frac{4p^3}{\delta} + \sqrt{\left(\frac{4}{3} \log \frac{4p^3}{\delta}\right)^2 + 4n_{\dagger} \left(1 - \mu_A^{\dagger}\right) \log \frac{4p^3}{\delta}} \\ &\leq \frac{8}{3} \log \frac{4p^3}{\delta} + \sqrt{4n_{\dagger} \left(1 - \mu_A^{\dagger}\right) \log \frac{4p^3}{\delta}}, \end{aligned} \quad (121)$$

which implies that

$$t = n_{\dagger} \left(\frac{4}{3n_{\dagger}} \log \frac{4p^3}{\delta} + \sqrt{\left(\frac{4}{3n_{\dagger}} \log \frac{4p^3}{\delta}\right)^2 + \frac{4}{n_{\dagger}} \left(1 - \mu_A^{\dagger}\right) \log \frac{4p^3}{\delta}} \right)$$

$$\leq n_{\dagger} \left(\frac{8}{3n_{\dagger}} \log \frac{4p^3}{\delta} + \sqrt{\frac{4}{n_{\dagger}} (1 - \mu_A^{\dagger}) \log \frac{4p^3}{\delta}} \right). \quad (122)$$

Define $\epsilon_{\dagger} = \sqrt{\log(2p^2/\delta) 2/n_{\dagger}}$ (as it is defined in [Bresler and Karzand \(2020\)](#)), then we get

$$t \leq n_{\dagger} \left(4\epsilon_{\dagger}^2 + 2\epsilon_{\dagger} \sqrt{1 - \mu_A^{\dagger}} \right) \leq n_{\dagger} \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1 - \mu_A^{\dagger}} \right\}. \quad (123)$$

■

Lemma 21 gives the concentration of measure bound for the event E_M defined in (112).

Lemma 21 *Fix $\delta > 0$ and let ϵ_{\dagger} be given by (106). For all pairs of vertices $u, \bar{u} \in V$ and edges $e = (w, \bar{w})$ in the path $\text{path}_{\text{T}}(u, \bar{u})$ from u to \bar{u} , given n_{\dagger} samples $M_{e,u,\bar{u}}^{(1)}, M_{e,u,\bar{u}}^{(2)}, \dots, M_{e,u,\bar{u}}^{(n)}$ of $M_{e,u,\bar{u}} = Y_w Y_{\bar{w}} + Y_u Y_{\bar{u}}$, it is true that*

$$\mathbb{P} \left(\left| \sum_{i=1}^{n_{\dagger}} M_{e,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E}[M_{e,u,\bar{u}}] \right| \leq n_{\dagger} \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1 + \mu_A^{\dagger}} \right\} \right) \geq 1 - \frac{\delta}{2}, \quad (124)$$

$$A = \text{path}_{\text{T}}(u, \bar{u}) \setminus \{e\}.$$

Proof Similarly to the prior Lemma, we calculate the mean and the variance as

$$\mathbb{E}[M_{e,u,\bar{u}}] = (1 - 2q)^2 \mathbb{E}[X_w X_{\bar{w}} + X_u X_{\bar{u}}] = (1 - 2q)^2 \mu_e (1 + \mu_A) \quad (125)$$

$$\begin{aligned} \text{Var}(M_{e,u,\bar{u}}) &= \mathbb{E}[(M_{e,u,\bar{u}})^2] - \mathbb{E}[M_{e,u,\bar{u}}]^2 \\ &= \mathbb{E}[(X_w N_w X_{\bar{w}} N_{\bar{w}} + N_u X_u N_{\bar{u}} X_{\bar{u}})^2] - [(1 - 2q)^2 \mathbb{E}[X_w X_{\bar{w}} + X_u X_{\bar{u}}]]^2 \\ &= \mathbb{E}[1 + 1 + 2X_w N_w X_{\bar{w}} N_{\bar{w}} N_u X_u N_{\bar{u}} X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} + X_u X_{\bar{u}}]^2 \\ &= 2 + 2\mathbb{E}[X_w N_w X_{\bar{w}} N_{\bar{w}} N_u X_u N_{\bar{u}} X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} + X_u X_{\bar{u}}]^2 \\ &= 2 + 2(1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} X_u X_{\bar{u}}] - (1 - 2q)^4 \mathbb{E}[X_w X_{\bar{w}} + X_u X_{\bar{u}}]^2 \\ &= 2 + 2(1 - 2q)^4 \mu_A - (1 - 2q)^4 (\mu_e (1 + \mu_A))^2 \\ &= 2 + (1 - 2q)^4 [2\mu_A - \mu_e^2 (1 + \mu_A)^2]. \end{aligned} \quad (126)$$

By applying Bernstein's inequality and we get that for any $t > 0$

$$\mathbb{P} \left[\left| \sum_{i=1}^{n_{\dagger}} M_{e,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E}[M_{e,u,\bar{u}}] \right| \leq t \right] \geq 1 - 2 \exp \left(- \frac{t^2}{2n_{\dagger} \text{Var}(M_{e,u,\bar{u}}) + \frac{2}{3} M t} \right).$$

Similarly, we find

$$t \leq n_{\dagger} \left(\frac{8}{3n_{\dagger}} \log \frac{4p^3}{\delta} + \sqrt{\frac{2}{n_{\dagger}} \text{Var}(M_{e,u,\bar{u}}) \log \frac{4p^3}{\delta}} \right) \quad (127)$$

and

$$\begin{aligned}
 \text{Var}(M_{e,u,\bar{u}}) &= 2 + (1 - 2q)^4 \left[2\mu_A - \mu_e^2 (1 + \mu_A)^2 \right] \\
 &\leq 2 + (1 - 2q)^4 2\mu_A \\
 &= 2 \left(1 + \mu_A^\dagger \right).
 \end{aligned} \tag{128}$$

We define $\epsilon_\dagger \triangleq \sqrt{\log(2p^2/\delta) 2/n_\dagger}$, then

$$t \leq n_\dagger \left(4\epsilon_\dagger^2 + 2\epsilon_\dagger \sqrt{1 + \mu_A^\dagger} \right) \leq n_\dagger \max \left\{ 8\epsilon_\dagger^2, 4\epsilon_\dagger \sqrt{1 + \mu_A^\dagger} \right\}, \tag{129}$$

which completes the proof. \blacksquare

Appendix C. Recovering Strong Edges

In Lemma 22, we define the set of strong edges for the hidden model and show that the event $E_\dagger^{\text{strong}}(\epsilon_\dagger)$ defined in (62) occurs with high probability. That is, only the strong edges are guaranteed to exist in the estimated structure T_\dagger^{CL} . We also find a lower bound for the necessary number of samples for exact structure recovery. In fact we have $n_\dagger \geq n$, as expected. Our bounds coincide with the noiseless case (Bresler and Karzand, 2020) by setting the noise level $q = 0$.

Lemma 22 Fix $\delta \in (0, 1)$, and let $\epsilon_\dagger = \sqrt{2 \log(2p^2/\delta) / n_\dagger}$, for any $n_\dagger > 0$. Consider the set of strong edges

$$\tau^\dagger \triangleq \frac{4\epsilon_\dagger \sqrt{1 - (1 - 2q)^4 \tanh \beta}}{(1 - \tanh \beta)} \text{ and } \left\{ (i, j) \in \mathcal{E}_T : |\tanh \theta_{ij}| \geq \frac{\tau^\dagger}{(1 - 2q)^2} \right\}. \tag{130}$$

Then, the Chow-Liu algorithm recovers the strong edges with probability at least $1 - \delta$. In other words, it is true that

$$\mathbb{P} \left[E_\dagger^{\text{strong}}(\epsilon_\dagger) \right] \geq 1 - 2p^2 \exp \left(-\frac{n_\dagger \epsilon_\dagger^2}{2} \right). \tag{131}$$

Proof From Lemma 19, if there is an error then for an edge f not recovered in the tree T_\dagger^{CL} , we have

$$\left(\sum_{i=1}^{n_\dagger} Z_{f,u,\bar{u}}^{(i)} \right) \left(\sum_{i=1}^{n_\dagger} M_{f,u,\bar{u}}^{(i)} \right) < 0$$

Therefore one of the sums must be negative. Expanding, one of the two following inequalities must hold:

$$\left| \sum_{i=1}^{n_\dagger} Z_{f,u,\bar{u}}^{(i)} - n_\dagger \mathbb{E} \left[Z_{f,u,\bar{u}}^{(i)} \right] \right| \geq n_\dagger \mathbb{E} \left[Z_{f,u,\bar{u}}^{(i)} \right]$$

$$\left| \sum_{i=1}^{n_{\dagger}} Y_{f,u,\bar{u}}^{(i)} - n_{\dagger} \mathbb{E} \left[Y_{f,u,\bar{u}}^{(i)} \right] \right| \geq n_{\dagger} \mathbb{E} \left[M_{f,u,\bar{u}}^{(i)} \right].$$

In addition, (115), (125), Lemma 20 and Lemma 21 give the following pairs of inequalities:

$$\begin{aligned} (1-2q)^2 \mu_f (1-\mu_A) &\leq \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}} \right\} \\ (1-2q)^2 \mu_f (1+\mu_A) &\leq \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1+\mu_A^{\dagger}} \right\} \\ |\mu_f^{\dagger}| &\leq (1-\mu_A)^{-1} \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}} \right\} \\ |\mu_f^{\dagger}| &\leq (1+\mu_A)^{-1} \max \left\{ 8\epsilon_{\dagger}^2, 4\epsilon_{\dagger} \sqrt{1+\mu_A^{\dagger}} \right\}. \end{aligned}$$

Putting these together:

$$\begin{aligned} |\mu_f^{\dagger}| &\leq \max \left\{ \frac{8\epsilon_{\dagger}^2}{(1-\mu_A)}, \frac{8\epsilon_{\dagger}^2}{(1+\mu_A)}, \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)}, \frac{4\epsilon_{\dagger} \sqrt{1+\mu_A^{\dagger}}}{(1+\mu_A)} \right\} \\ &\leq \max \left\{ \frac{8\epsilon_{\dagger}^2}{(1-\mu_A)}, \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)} \right\} \\ &\leq \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)}. \end{aligned} \tag{132}$$

We get the last inequality for non trivial values of the bound $\frac{8\epsilon_{\dagger}^2}{(1-\mu_A^{\dagger})} \leq 1$ and by using the following bound

$$\frac{8\epsilon_{\dagger}^2}{(1-\mu_A)} \leq \frac{16\epsilon_{\dagger}^2}{(1-\mu_A)} \leq \frac{4\epsilon_{\dagger}}{\sqrt{1-\mu_A}} = \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A}}{(1-\mu_A)} \leq \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)}. \tag{133}$$

Finally the function $f(\mu_A) = \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)} = \frac{4\epsilon_{\dagger} \sqrt{1-(1-2q)^4 \mu_A}}{(1-\mu_A)}$ is increasing with respect to μ_A (for all $\mu_A \leq 1$) and $\mu_A \leq \tanh \beta < 1$, so we have

$$|\mu_f^{\dagger}| \leq \frac{4\epsilon_{\dagger} \sqrt{1-\mu_A^{\dagger}}}{(1-\mu_A)} \leq \frac{4\epsilon_{\dagger} \sqrt{1-(1-2q)^4 \tanh \beta}}{(1-\tanh \beta)} \triangleq \tau^{\dagger}. \tag{134}$$

The weakest edge should satisfy $|\mu_f^{\dagger}| \geq \tau^{\dagger}$ to guarantee the correct recovery of the tree under the event $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$. This yields a condition on the edge strengths:

$$|\mu_f^{\dagger}| \geq \tau^{\dagger} \implies$$

$$\begin{aligned}
 (1-2q)^2 \tanh \alpha &\geq \frac{4\epsilon_{\dagger} \sqrt{1 - (1-2q)^4 \tanh \beta}}{(1 - \tanh \beta)} \\
 \tanh \alpha &\geq \frac{4\epsilon_{\dagger} \sqrt{1 - (1-2q)^4 \tanh \beta}}{(1-2q)^2 (1 - \tanh \beta)}, \quad q \in [0, \frac{1}{2}].
 \end{aligned} \tag{135}$$

The last inequality gives the definition of the strong edges in the noisy scheme. \blacksquare

Based on the definition (134) we derive the following bound on τ^{\dagger}

$$\begin{aligned}
 \tau^{\dagger} &= \frac{4\epsilon_{\dagger} \sqrt{1 - (1-2q)^4 \tanh \beta}}{(1 - \tanh \beta)} \\
 &= \frac{4\epsilon_{\dagger} \sqrt{1 - (1-8q+24q^2-32q^3+16q^4) \tanh \beta}}{(1 - \tanh \beta)} \\
 &\leq 4\epsilon_{\dagger} \frac{\sqrt{1 - \tanh \beta} + \sqrt{(1-3q+4q^2-2q^3) 8q \tanh \beta}}{(1 - \tanh \beta)} \\
 &\leq 4\epsilon_{\dagger} e^{\beta} \left(1 + e^{\beta} \sqrt{(1-q)(2q^2-2q+1) 8q \tanh \beta} \right) \\
 &< 4\epsilon_{\dagger} e^{\beta} \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right).
 \end{aligned} \tag{136}$$

$$< 4\epsilon_{\dagger} e^{\beta} \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right). \tag{137}$$

(136) holds because $1 - \tanh(\beta) \geq e^{-2\beta}$. We will later use (137) in Lemma 25.

In comparison to the noiseless setting (see [Bresler and Karzand \(2020\)](#)), we can guarantee exact recover with high probability under the event $E^{\text{strong}}(\epsilon)$ when the weakest edge satisfies the inequality

$$\tanh \alpha \geq \frac{4\epsilon}{\sqrt{1 - \tanh \beta}}. \tag{138}$$

Notice that (138) can be obtained by (135) when $q = 0$ and $n = n_{\dagger}$. When $q > 0$ and $n = n_{\dagger}$ it is clear that the set of trees that can be recovered from noisy observations is a subset of the set of trees that can be recovered from the original observations. Also, we have

$$\begin{aligned}
 \epsilon = \sqrt{2 \log(2p^2/\delta)/n} &\implies n = \frac{2}{\epsilon^2} \log(2p^2/\delta) \quad \text{and} \\
 \epsilon_{\dagger} = \sqrt{2 \log(2p^2/\delta)/n_{\dagger}} &\implies n_{\dagger} = \frac{2}{\epsilon_{\dagger}^2} \log(2p^2/\delta).
 \end{aligned} \tag{139}$$

By combining (135) with (139) we found the number of samples that we need to recover the tree with probability at $1 - \delta$ (Theorem 5):

$$n_{\dagger} > \frac{32 \left[1 - (1-2q)^4 \tanh \beta \right]}{(1 - \tanh \beta)^2 (1-2q)^4 \tanh^2 \alpha} \log \frac{2p^2}{\delta}. \tag{140}$$

On the other hand when there is no noise ([Bresler and Karzand, 2020](#)) we need

$$n > \frac{32}{\tanh^2 \alpha (1 - \tanh \beta)} \log \frac{2p^2}{\delta}. \tag{141}$$

Appendix D. Analysis of the Event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$

Lemma 23 Consider a path of length $d \geq 2$ in the original tree T , and without loss of generality assume that path is $X_1 - X_2 - \dots - X_{d+1}$. Recall that $Y_m^{(i)}$ is the i^{th} sample of Y_m and $m \in [d+1]$ and $\hat{\mu}_k^{\dagger} \triangleq \frac{1}{n} \sum_{i=1}^n Y_k^{(i)} Y_{k+1}^{(i)}$, $k \in [d]$. Then

$$\begin{aligned} & \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^{\dagger}, \dots, \hat{\mu}_1^{\dagger} \right) \\ &= \frac{1 \pm (1-2q)^2 \mu_k}{2} \frac{1 - \mu_{k-1}^{\dagger} \hat{\mu}_{k-1}^{\dagger}}{1 - (\mu_{k-1}^{\dagger})^2} + \mu_{k-1}^{\dagger} \frac{1 \pm \mu_k \hat{\mu}_{k-1}^{\dagger} - \mu_{k-1}^{\dagger}}{2} \frac{1 - \mu_{k-1}^{\dagger}}{1 - (\mu_{k-1}^{\dagger})^2}. \end{aligned} \quad (142)$$

Proof Note that

$$\hat{\mu}_k^{\dagger} = \frac{1}{n} \sum_{i=1}^n Y_k^{(i)} Y_{k+1}^{(i)} = \frac{1}{n} \sum_{i=1}^n (X_k N_k X_{k+1} N_{k+1})^{(i)}, \quad (143)$$

where each term

$$(X_k N_k X_{k+1} N_{k+1})^{(\ell)} \perp \hat{\mu}_r^{\dagger} \quad \forall r \in [1, 2, \dots, k-2], \forall \ell \in [1, 2, \dots, n] \quad (144)$$

thus

$$\begin{aligned} & \mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^{\dagger}, \dots, \hat{\mu}_1^{\dagger} \right) \\ &= \mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^{\dagger} \right) \\ &= \mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right) \\ &= \frac{\mathbb{P} \left(\hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \middle| (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right)}{\mathbb{P} \left(\hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right)} \\ &\quad \times \mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right). \end{aligned} \quad (145)$$

First we compute the probability $\mathbb{P} \left(\hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right)$. Define the Bernoulli random variable Z_{k-1} as

$$Z_{k-1} \triangleq \frac{X_{k-1} N_{k-1} X_k N_k + 1}{2} = \begin{cases} 0, & \text{w.p. } \frac{1-(1-2q)^2 \mu_{k-1}}{2} \\ 1, & \text{w.p. } \frac{1+(1-2q)^2 \mu_{k-1}}{2}. \end{cases} \quad (146)$$

Then

$$\begin{aligned} & \mathbb{P} \left(\hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right) \\ &= \mathbb{P} \left(\hat{\mu}_{k-1}^{\dagger} = \frac{1}{n} \sum_{i=1}^n (2Z_{k-1} - 1)^{(i)} \right) \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{P} \left(\sum_{i=1}^n Z_{k-1}^{(i)} = n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right) \\
 &= \binom{n}{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 - (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n - n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 + (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}}. \quad (147)
 \end{aligned}$$

As a second step we compute the probability

$$\begin{aligned}
 &\mathbb{P} \left(\hat{\mu}_{k-1}^\dagger = \frac{1}{n} \sum_{i=1}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \quad (148) \\
 &= \mathbb{P} \left(\hat{\mu}_{k-1}^\dagger = \frac{1}{n} \sum_{i=1, i \neq \ell}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right. \\
 &\quad \left. + \frac{(X_{k-1} N_{k-1} X_k N_k)^{(\ell)}}{n} \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right).
 \end{aligned}$$

Note that

$$(X_{k-1} N_{k-1} X_k N_k)^{(i)} \perp (X_k N_k X_{k+1} N_{k+1})^{(\ell)}, \quad \forall i \neq \ell, \quad (149)$$

and we would like to find the conditional distribution of $(X_{k-1} N_{k-1} X_k N_k)^{(\ell)}$ under the event $\{(X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1\}$. We have

$$\begin{aligned}
 &\mathbb{P} \left((X_{k-1} N_{k-1} X_k N_k)^{(\ell)} = c \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \\
 &= \frac{\mathbb{P} \left((X_{k-1} N_{k-1} X_k N_k)^{(\ell)} = c, (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right)}{\mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right)} \\
 &= \frac{\frac{1 \pm c \mathbb{E}[X_{k-1} N_{k-1} X_{k+1} N_{k+1}] + c(1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{4}}{\frac{1 \pm (1-2q)^2 \mu_k}{2}} \\
 &= \frac{1 \pm c(1-2q)^2 \mu_{k-1} \mu_k + c(1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{2(1 \pm (1-2q)^2 \mu_k)}, \quad c \in \{-1, +1\}. \quad (150)
 \end{aligned}$$

Define

$$P_1 \triangleq \mathbb{P} \left((X_{k-1} N_{k-1} X_k N_k)^{(\ell)} = +1 \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \quad (151)$$

$$P_2 \triangleq \mathbb{P} \left((X_{k-1} N_{k-1} X_k N_k)^{(\ell)} = -1 \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right), \quad (152)$$

then

$$\begin{aligned}
 &\mathbb{P} \left(\hat{\mu}_{k-1}^\dagger = \frac{1}{n} \sum_{i=1, i \neq \ell}^n (X_{k-1} N_{k-1} X_k N_k)^{(i)} \right. \\
 &\quad \left. + \frac{(X_{k-1} N_{k-1} X_k N_k)^{(\ell)}}{n} \mid (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right)
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{P} \left(\hat{\mu}_{k-1}^\dagger = \frac{1}{n} \sum_{i=1, i \neq \ell}^n (2Z_{k-1} - 1)^{(i)} + \frac{(2Z_{k-1} - 1)^{(\ell)}}{n} \middle| (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \\
 &= \mathbb{P} \left(\sum_{i=1, i \neq \ell}^n (Z_{k-1})^{(i)} + (Z_{k-1})^{(\ell)} = n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \middle| (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \\
 &= \mathbb{P} \left((Z_{k-1})^{(\ell)} = 0 \middle| (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \mathbb{P} \left(\sum_{i=1, i \neq \ell}^n (Z_{k-1})^{(i)} = n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right) \\
 &\quad + \mathbb{P} \left((Z_{k-1})^{(\ell)} = 1 \middle| (X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \mathbb{P} \left(\sum_{i=1, i \neq \ell}^n (Z_{k-1})^{(i)} = n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} - 1 \right) \\
 &= P_2 \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right) \left(\frac{1 - (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n-1-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 + (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \\
 &\quad + P_1 \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} - 1 \right) \left(\frac{1 - (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n-1-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} + 1} \\
 &\quad \left(\frac{1 + (1 - 2q)^2 \mu_{k-1}}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} - 1}. \tag{153}
 \end{aligned}$$

Finally

$$\mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) = \frac{1 \pm (1 - 2q)^2 \mu_k}{2}, \tag{154}$$

and (145), (147), (148), (153) give

$$\begin{aligned}
 &\frac{\mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right)}{\mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right)} \\
 &= P_2 \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right) \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right)^{-1} \frac{\left(\frac{1 - \mu_{k-1}^\dagger}{2} \right)^{n-1-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 + \mu_{k-1}^\dagger}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}}}{\left(\frac{1 - \mu_{k-1}^\dagger}{2} \right)^{n-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 + \mu_{k-1}^\dagger}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}}} \\
 &\quad + P_1 \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} - 1 \right) \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right)^{-1} \frac{\left(\frac{1 - \mu_{k-1}^\dagger}{2} \right)^{n-1-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} + 1} \left(\frac{1 + \mu_{k-1}^\dagger}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} - 1}}{\left(\frac{1 - \mu_{k-1}^\dagger}{2} \right)^{n-n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}} \left(\frac{1 + \mu_{k-1}^\dagger}{2} \right)^{n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2}}} \\
 &= P_2 \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right) \left(n \frac{\hat{\mu}_{k-1}^\dagger + 1}{2} \right)^{-1} \left(\frac{1 - \mu_{k-1}^\dagger}{2} \right)^{-1}
 \end{aligned}$$

$$\begin{aligned}
 & + P_1 \binom{n-1}{n^{\frac{\hat{\mu}_{k-1}^\dagger+1}{2}}-1} \binom{n}{n^{\frac{\hat{\mu}_{k-1}^\dagger+1}{2}}}^{-1} \left(\frac{1+\mu_{k-1}^\dagger}{2} \right)^{-1} \\
 & = P_2 \frac{n - n^{\frac{\hat{\mu}_{k-1}^\dagger+1}{2}}}{n} \left(\frac{1-\mu_{k-1}^\dagger}{2} \right)^{-1} + P_1 \frac{n^{\frac{\hat{\mu}_{k-1}^\dagger+1}{2}}}{n} \left(\frac{1+\mu_{k-1}^\dagger}{2} \right)^{-1} \\
 & = P_2 \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} + P_1 \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger}.
 \end{aligned} \tag{155}$$

The latter and the definition of P_1, P_2 (see Equations 151 and 152) give

$$\begin{aligned}
 & \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} \pm 1 \mid \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right) \\
 & = \left[P_2 \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} + P_1 \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} \right] \mathbb{P} \left((X_k N_k X_{k+1} N_{k+1})^{(\ell)} = \pm 1 \right) \\
 & = \left[P_2 \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} + P_1 \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} \right] \frac{1 \pm (1-2q)^2 \mu_k}{2} \\
 & = \left[\frac{1 \mp (1-2q)^2 \mu_{k-1} \mu_k - (1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{2(1 \pm (1-2q)^2 \mu_k)} \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} \right. \\
 & \quad \left. + \frac{1 \pm (1-2q)^2 \mu_{k-1} \mu_k + (1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{2(1 \pm (1-2q)^2 \mu_k)} \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} \right] \frac{1 \pm (1-2q)^2 \mu_k}{2} \\
 & = \frac{1 \mp (1-2q)^2 \mu_{k-1} \mu_k - (1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{4} \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} \\
 & \quad + \frac{1 \pm (1-2q)^2 \mu_{k-1} \mu_k + (1-2q)^2 \mu_{k-1} \pm (1-2q)^2 \mu_k}{4} \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} \\
 & = \frac{1 \pm (1-2q)^2 \mu_k}{4} \left(\frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} + \frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} \right) \\
 & \quad + \frac{\pm(1-2q)^2 \mu_{k-1} \mu_k + (1-2q)^2 \mu_{k-1}}{4} \left(\frac{1+\hat{\mu}_{k-1}^\dagger}{1+\mu_{k-1}^\dagger} - \frac{1-\hat{\mu}_{k-1}^\dagger}{1-\mu_{k-1}^\dagger} \right) \\
 & = \frac{1 \pm (1-2q)^2 \mu_k}{2} \frac{1-\mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1-(\mu_{k-1}^\dagger)^2} + (1-2q)^2 \mu_{k-1} \frac{1 \pm \mu_k}{2} \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1-(\mu_{k-1}^\dagger)^2}.
 \end{aligned} \tag{156}$$

Note that $(1-2q)^2 \mu_{k-1} = \mu_{k-1}^\dagger$, and the proof is completed. ■

Lemma 24 Define the function $K(\beta, q)$ as

$$K(\beta, q) \triangleq \frac{10(1 - \tanh^2(\beta))}{9 + (1 - 2q)^2 - \tanh^2(\beta)(1 - 2q)^2(9(1 - 2q)^2 + 1)} \tag{157}$$

and the event $E_{e,\dagger}^{\text{edge}}$ as

$$E_{e,\dagger}^{\text{edge}} \triangleq \left\{ \left| \hat{\mu}_e^\dagger - \mu_e^\dagger \right| \leq \gamma_e \right\}, \quad e \in \mathcal{E}_T, \quad \gamma_e > 0, \quad (158)$$

and $E_\dagger^{\text{edge}}(\mathcal{E}_T) \triangleq \cap_{e \in \mathcal{E}_T} E_{e,\dagger}^{\text{edge}}$. If

$$n \geq \frac{108e^{2\beta} \log(2p/\delta)}{(1-2q)^4 K(\beta, q)} \text{ and } \gamma_e = \sqrt{3 \frac{1-\mu_e^2}{nK(\beta, q)} \log(2p/\delta)} \quad (159)$$

then $\mathbb{P} \left[\left(E_\dagger^{\text{edge}}(\mathcal{E}_T) \right)^c \right] \leq \delta$.

Proof The variance of $\hat{\mu}_e^\dagger$ is $(1 - (\mu_e^\dagger)^2)/n$ and by applying Bernstein's inequality

$$\mathbb{P} \left[\left(E_{e,\dagger}^{\text{edge}} \right)^c \right] \leq 2 \exp \left(- \frac{n\gamma_e^2}{2 \left(1 - (\mu_e^\dagger)^2 \right) + \frac{4}{3}\gamma_e} \right), \quad \forall \gamma_e > 0. \quad (160)$$

We choose $\gamma_e = \sqrt{3 \frac{1-\mu_e^2}{nK(\beta, q)} \log(2p/\delta)}$ (because the parameter γ_e is free, that is, Bernstein's inequality holds for all $\gamma_e > 0$). If n satisfies (159) then

$$\gamma_e \leq \sqrt{3 \frac{1-\mu_e^2}{108e^{2\beta}(1-2q)^4}} \leq \frac{(1-2q)^2}{6} (1-\mu_e^2), \quad (161)$$

and the last is true because $e^{-2\beta} \leq 1 - \tanh(\beta) \leq 1 - |\mu_e| \leq 1 - \mu_e^2$. By applying (159) and (161) on (160) we get

$$\begin{aligned} & \mathbb{P} \left[\left(E_{e,\dagger}^{\text{edge}} \right)^c \right] \\ & \leq 2 \exp \left(- \frac{n\gamma_e^2}{2 \left(1 - (\mu_e^\dagger)^2 \right) + \frac{4}{3}\gamma_e} \right) \\ & \leq 2 \exp \left(- \frac{3 \frac{1-\mu_e^2}{K(\beta, q)} \log(2p/\delta)}{2 \left(1 - (1-2q)^4 \mu_e^2 \right) + \frac{4}{3} \frac{(1-2q)^2}{6} (1-\mu_e^2)} \right) \\ & = 2 \exp \left(- \frac{3}{K(\beta, q)} \frac{1-\mu_e^2}{2 + \frac{2}{9}(1-2q)^2 - \mu_e^2(1-2q)^2(2(1-2q)^2 + \frac{2}{9})} \log(2p/\delta) \right) \\ & = 2 \exp \left(- \frac{3}{2K(\beta, q)} \frac{1-\mu_e^2}{1 + \frac{1}{9}(1-2q)^2 - \mu_e^2(1-2q)^2((1-2q)^2 + \frac{1}{9})} \log(2p/\delta) \right) \\ & \leq 2 \exp \left(- \frac{10}{9K(\beta, q)} \frac{1-\mu_e^2}{1 + \frac{1}{9}(1-2q)^2 - \mu_e^2(1-2q)^2((1-2q)^2 + \frac{1}{9})} \log(2p/\delta) \right). \end{aligned} \quad (162)$$

The following function

$$f(x) = \frac{10}{9} \frac{1-x}{1 + \frac{1}{9}(1-2q)^2 - x(1-2q)^2((1-2q)^2 + \frac{1}{9})}, \quad x \in [\tanh^2(\alpha), \tanh^2(\beta)] \quad (163)$$

is strictly decreasing, thus we have $f(\tanh^2(\beta)) \leq f(x)$ for all $x \in [\tanh^2(\alpha), \tanh^2(\beta)]$. Also $K(\beta, q) \equiv f(\tanh^2(\beta))$, the latter together with (162) give

$$\begin{aligned}
 & \mathbb{P} \left[\left(E_{e, \dagger}^{\text{edge}} \right)^c \right] \\
 & \leq 2 \exp \left(-\frac{1}{K(\beta, q)} f(\mu_e^2) \log(2p/\delta) \right) \\
 & \leq 2 \exp \left(-\frac{1}{K(\beta, q)} f(\tanh^2(\beta)) \log(2p/\delta) \right) \\
 & = 2 \exp \left(-\frac{1}{K(\beta, q)} K(\beta, q) \log(2p/\delta) \right) \\
 & = \frac{\delta}{p}.
 \end{aligned} \tag{164}$$

Finally, by applying union over the $p - 1$ edges of the tree we get $\mathbb{P} \left[\left(E_{\dagger}^{\text{edge}}(\mathcal{E}_T) \right)^c \right] \leq \delta$. ■

The next Lemma is the extension of Lemma 8.7 by [Bresler and Karzand \(2020\)](#). The sample complexity bound exactly recovers the noiseless case and its expression is continuous at $q = 0$. Further, the bound is independent of the length of the longest path d , similarly to the noiseless setting. Finally, we provide upper bounds on the functions that appear in the bound. The latter give a more tractable version of the result and a clear representation of the required number of samples as a function of the parameters.

Lemma 25 (Concentration bound for the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$) For $\beta > 0$ and $q \in [0, 1/2)$ we define the functions $S(\cdot)$, $G(\cdot)$, $K(\cdot)$, $A(\cdot)$, $\Delta(\cdot)$

$$S(\beta, q) \triangleq 2 + \frac{(1 - 2q)^2}{6} (1 - (1 - 2q)^2) \tanh^2(\beta) \leq 3 - (1 - 2q)^2 \triangleq S \tag{165}$$

$$A(\beta, q) \triangleq A \triangleq (1 - 2q)^2 [1 - \tanh(\beta)(1 - (1 - 2q)^2)] \tag{166}$$

$$G(\beta, q) \triangleq \frac{3}{4(1 - 2q)^2} \left[d(1 - A) \left(\frac{A + 2}{3} \right)^d + 1 \right] \leq \frac{3(3e^{-1} \mathbb{1}_{q \neq 0} + 1)}{4(1 - 2q)^2} \triangleq G, \tag{167}$$

and the inequality in (167) holds because the function $G(\beta, q)$ is bounded for all $d \in \mathcal{N} \setminus \{1\}$.

$$K(\beta, q) \triangleq \frac{10(1 - \tanh^2(\beta))}{9 + (1 - 2q)^2 - \tanh^2(\beta)(1 - 2q)^2(9(1 - 2q)^2 + 1)} \geq e^{-2\beta \mathbb{1}_{q=0}} \triangleq K \tag{168}$$

$$\Delta \triangleq \frac{1 - (1 - 2q)^2}{1 - (1 - 2q)^4 \tanh^2(\beta)} \sqrt{\frac{3 \log(2p^3/\delta)}{n}} \tanh^2(\beta) e^{2\beta}. \tag{169}$$

If $\Delta < \gamma_{\dagger} \leq S(\beta, q)G(\beta, q)/3 + \Delta$ and

$$n \geq \max \left\{ \frac{S^2(\beta, q)G^2(\beta, q)}{0.3^2 (\gamma_{\dagger} - \Delta)^2} \log(4p^2/\delta), \frac{108e^{2\beta}}{(1 - 2q)^4 K(\beta, q)} \log(2p^3/\delta) \right\} \tag{170}$$

then for any path $\mathcal{A}_d = \{e_1, e_2, \dots, e_d\}$ of \mathbb{T} with d edges, it is true that

$$\mathbb{P} \left(\left| \prod_{e \in \mathcal{A}_d} \frac{\hat{\mu}_e^{\dagger}}{(1 - 2q)^2} - \prod_{e \in \mathcal{A}_d} \frac{\mu_e^{\dagger}}{(1 - 2q)^2} \right| \geq \gamma_{\dagger} \right) \leq \frac{2\delta}{p^2}, \quad d > 2. \tag{171}$$

Proof For sake of space we proceed by using the notation μ_k^\dagger and $\hat{\mu}_k^\dagger$ instead of $\mu_{e_k}^\dagger$ and $\hat{\mu}_{e_k}^\dagger$ for $k \in [d]$. Define the random variable

$$M_i^\dagger \triangleq \left(\frac{\hat{\mu}_i^\dagger}{(1-2q)^2} - \frac{\mu_i^\dagger}{(1-2q)^2} \right) \prod_{j=1}^{i-1} \frac{\hat{\mu}_j^\dagger}{(1-2q)^2} \prod_{j=i+1}^d \frac{\mu_j^\dagger}{(1-2q)^2}. \quad (172)$$

Then $\sum_{i=1}^d M_i^\dagger = \prod_{i=1}^d \frac{\hat{\mu}_i^\dagger}{(1-2q)^2} - \prod_{i=1}^d \frac{\mu_i^\dagger}{(1-2q)^2}$, and define the sequence of paths with length k as $\mathcal{A}_k \triangleq \{e_1, e_2, \dots, e_k\} \subset \mathcal{A}_d$, for $2 \leq k \leq d$. Although we provided the definition of the event $E_\dagger^{\text{edge}}(\cdot)$ in Lemma 24, we restate it below for completeness. For some $\gamma_e > 0$ the definition follows

$$E_\dagger^{\text{edge}}(\mathcal{A}_k) \triangleq \bigcap_{e \in \mathcal{A}_k} \left\{ \left| \hat{\mu}_e^\dagger - \mu_e^\dagger \right| \leq \gamma_e \right\}. \quad (173)$$

The law of total probability gives

$$\mathbb{P} \left[\left| \sum_{i=1}^d M_i^\dagger \right| > \gamma \right] \leq \mathbb{P} \left[\left| \sum_{i=1}^d M_i^\dagger \right| > \gamma \mid E_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right] + \mathbb{P} \left[\left(E_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right)^c \right]. \quad (174)$$

For second term, Lemma 24 gives that $\mathbb{P} \left[\left(E_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right)^c \right] \leq \delta/p^2$ if

$$n \geq \frac{108e^{2\beta} \log(2p^3/\delta)}{(1-2q)^4 K(\beta, q)}, \quad (175)$$

we define the function $K(\beta, q)$ in Lemma 24 (157). Here we will find an upper bound for the first term of the right hand-side of (174). Note that M_k^\dagger is written as

$$\begin{aligned} M_k^\dagger &= \left(\frac{\frac{1}{n} \sum_{\ell=1}^n (Y_k Y_{k+1})^{(\ell)}}{(1-2q)^2} - \frac{\mu_k^\dagger}{(1-2q)^2} \right) \prod_{j=1}^{k-1} \frac{\hat{\mu}_j^\dagger}{(1-2q)^2} \prod_{j=k+1}^d \frac{\mu_j^\dagger}{(1-2q)^2} \\ &= \frac{1}{n} \sum_{\ell=1}^n \left(\frac{(X_k N_k X_{k+1} N_{k+1})^{(\ell)}}{(1-2q)^2} - \frac{\mu_k^\dagger}{(1-2q)^2} \right) \prod_{j=1}^{k-1} \frac{\hat{\mu}_j^\dagger}{(1-2q)^2} \prod_{j=k+1}^d \frac{\mu_j^\dagger}{(1-2q)^2}. \end{aligned} \quad (176)$$

and we define

$$Z_k^{(\ell)} \triangleq \left(\frac{(X_k N_k X_{k+1} N_{k+1})^{(\ell)}}{(1-2q)^2} - \frac{\mu_k^\dagger}{(1-2q)^2} \right) \prod_{j=1}^{k-1} \frac{\hat{\mu}_j^\dagger}{(1-2q)^2} \prod_{j=k+1}^d \frac{\mu_j^\dagger}{(1-2q)^2}. \quad (177)$$

The random variables $Z_k^{(\ell)}$ for $\ell \in [n]$ and fixed $k \in [d]$ are independent conditioned on the event $E_\dagger^{\text{edge}}(\mathcal{A}_{k-1})$. However the conditional expectation $\mathbb{E}[Z_k^{(i)} | Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger]$ is not zero. To apply a concentration of measure result on $Z_k^{(\ell)}$ we use the extended Bennet's inequality for supermartingales (Fan et al., 2012).

Martingale Differences: Define $\xi_k^{(0)} \triangleq 0$, $\xi_k^{(1)} \triangleq Z_k^{(1)} - \mathbb{E} \left[Z_k^{(1)} | \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right]$, $\xi_k^{(i)} \triangleq Z_k^{(i)} - \mathbb{E} \left[Z_k^{(i)} | Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right]$. Also, define as \mathcal{F}_{i-1}^k the σ -algebra generated by $Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger$, then $(\xi_k^{(i)}, \mathcal{F}_i^k)_{i=1, \dots, n}$ is a Martingale Difference Sequence (MDS).

Additionally, conditioned on $Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger$ we have

$$Z_k^{(i)} = \begin{cases} \frac{1}{(1-2q)^{2d}} \left(1 - \mu_k^\dagger \right) \prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger, & \text{w.p. } \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = +1 \middle| \hat{\mu}_{k-1}^\dagger \right) \\ -\frac{1}{(1-2q)^{2d}} \left(1 + \mu_k^\dagger \right) \prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger, & \text{w.p. } \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = -1 \middle| \hat{\mu}_{k-1}^\dagger \right), \end{cases} \quad (178)$$

and we have proved (Lemma 23) that

$$\begin{aligned} & \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right) \\ &= \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = \pm 1 \middle| \hat{\mu}_{k-1}^\dagger \right) \\ &= \frac{1 \pm \mu_k^\dagger}{2} \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} + \mu_{k-1}^\dagger \frac{1 \pm \mu_k}{2} \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2}. \end{aligned} \quad (179)$$

Thus we have

$$\begin{aligned} \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] &= \mathbb{E} \left[Z_k^{(i)} | Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right] \\ &= \left[\left(1 - \mu_k^\dagger \right) \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = +1 \middle| \hat{\mu}_{k-1}^\dagger \right) - \left(1 + \mu_k^\dagger \right) \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = -1 \middle| \hat{\mu}_{k-1}^\dagger \right) \right] \\ &\quad \times \frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}}. \end{aligned} \quad (180)$$

Note that (179) gives

$$\begin{aligned} & \left[\left(1 - \mu_k^\dagger \right) \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = +1 \middle| \hat{\mu}_{k-1}^\dagger \right) - \left(1 + \mu_k^\dagger \right) \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = -1 \middle| \hat{\mu}_{k-1}^\dagger \right) \right] \\ &= \left(\left(1 - \mu_k^\dagger \right) \frac{1 + \mu_k^\dagger}{2} - \left(1 + \mu_k^\dagger \right) \frac{1 - \mu_k^\dagger}{2} \right) \mu_{k-1}^\dagger \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &\quad + \left(\left(1 - \mu_k^\dagger \right) \frac{1 + \mu_k}{2} - \left(1 + \mu_k^\dagger \right) \frac{1 - \mu_k}{2} \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= \left(\left(1 - \mu_k^\dagger \right) \frac{1 + \mu_k}{2} - \left(1 + \mu_k^\dagger \right) \frac{1 - \mu_k}{2} \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= \frac{1}{2} \left(1 - \mu_k^\dagger + \mu_k - \mu_k^\dagger \mu_k - 1 + \mu_k - \mu_k^\dagger + \mu_k^\dagger \mu_k \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= (\mu_k - \mu_k^\dagger) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2}. \end{aligned} \quad (181)$$

Combine the latter with (180) to get

$$\mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] = \mu_{k-1}^\dagger (\mu_k - \mu_k^\dagger) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1 - 2q)^{2d}}, \quad i \in [n]. \quad (182)$$

If $q = 0$ then $\mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] = 0$. Also $\lim_{n \rightarrow \infty} \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \rightarrow 0$ for all $q \in [0, 1/2)$ because $\lim_{n \rightarrow \infty} \hat{\mu}_{k-1}^\dagger \rightarrow \mu_{k-1}^\dagger$. Note that

$$\begin{aligned} \mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] &= \mathbb{E} \left[\left(Z_k^{(i)} - \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right)^2 \middle| \mathcal{F}_{i-1}^k \right] \\ &= \mathbb{E} \left[\left(Z_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] - \mathbb{E}^2 \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right]. \end{aligned} \quad (183)$$

We compute $\mathbb{E} \left[\left(Z_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right]$:

$$\begin{aligned} \mathbb{E} \left[\left(Z_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] &= \mathbb{E} \left[\left(Z_k^{(i)} \right)^2 \middle| Z_k^{(i-1)}, \dots, Z_k^{(1)}, \hat{\mu}_{k-1}^\dagger, \dots, \hat{\mu}_1^\dagger \right] \\ &= \left[\left(1 - \mu_k^\dagger \right)^2 \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = +1 \middle| \hat{\mu}_{k-1}^\dagger \right) + \left(1 + \mu_k^\dagger \right)^2 \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = -1 \middle| \hat{\mu}_{k-1}^\dagger \right) \right] \\ &\quad \times \left(\frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1 - 2q)^{2d}} \right)^2. \end{aligned} \quad (184)$$

We use (179) to find

$$\begin{aligned} &\left[\left(1 - \mu_k^\dagger \right)^2 \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = +1 \middle| \hat{\mu}_{k-1}^\dagger \right) + \left(1 + \mu_k^\dagger \right)^2 \mathbb{P} \left(Y_k^{(\ell)} Y_{k+1}^{(\ell)} = -1 \middle| \hat{\mu}_{k-1}^\dagger \right) \right] \\ &= \left(\left(1 - \mu_k^\dagger \right)^2 \frac{1 + \mu_k^\dagger}{2} + \left(1 + \mu_k^\dagger \right)^2 \frac{1 - \mu_k^\dagger}{2} \right) \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &\quad + \left(\left(1 - \mu_k^\dagger \right)^2 \frac{1 + \mu_k}{2} + \left(1 + \mu_k^\dagger \right)^2 \frac{1 - \mu_k}{2} \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= (1 - (\mu_k^\dagger)^2) \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} + \left(1 + (\mu_k^\dagger)^2 - 2\mu_k \mu_k^\dagger \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= (1 - (\mu_k^\dagger)^2) \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} + \left(1 + (\mu_k^\dagger)^2 - 2(\mu_k^\dagger)^2 + 2(\mu_k^\dagger)^2 - 2\mu_k \mu_k^\dagger \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &= (1 - (\mu_k^\dagger)^2) \frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} + \left(1 - (\mu_k^\dagger)^2 \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\ &\quad + 2\mu_k^\dagger \left(\mu_k^\dagger - \mu_k \right) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \end{aligned}$$

$$\begin{aligned}
 &= \left(1 - (\mu_k^\dagger)^2\right) \left[\frac{1 - \mu_{k-1}^\dagger \hat{\mu}_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} + \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \right] + 2\mu_k^\dagger (\mu_k^\dagger - \mu_k) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \\
 &= \left(1 - (\mu_k^\dagger)^2\right) + 2\mu_k^\dagger (\mu_k^\dagger - \mu_k) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2}. \tag{185}
 \end{aligned}$$

Now we combine (182), (183), (184) and (185) to get

$$\begin{aligned}
 &\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] \\
 &= \mathbb{E} \left[\left(Z_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] - \mathbb{E}^2 \left[Z_k^{(i)} \middle| \mathcal{F}_{i-1}^k \right] \\
 &= \left[\left(1 - (\mu_k^\dagger)^2\right) + 2\mu_k^\dagger (\mu_k^\dagger - \mu_k) \mu_{k-1}^\dagger \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} - \left(\mu_{k-1}^\dagger (\mu_k - \mu_k^\dagger) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \right)^2 \right] \\
 &\quad \times \left(\frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}} \right)^2 \\
 &= \left[1 - \left(\mu_k^\dagger + \mu_{k-1}^\dagger (\mu_k - \mu_k^\dagger) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \right)^2 \right] \left(\frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}} \right)^2. \tag{186}
 \end{aligned}$$

For sake of space we define the function

$$\hat{f}_{\mu_{k-1}}^{\mu_k}(q) \equiv f(\mu_k^\dagger, \mu_{k-1}^\dagger, \hat{\mu}_{k-1}^\dagger, q) \triangleq \mu_{k-1}^\dagger (\mu_k - \mu_k^\dagger) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2}, \tag{187}$$

and then (182) and (186) can be written as

$$\mathbb{E} \left[Z_k^{(i)} \middle| \mathcal{F}_{i-1}^k \right] = \hat{f}_{\mu_{k-1}}^{\mu_k}(q) \frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}}, \quad i \in [n], \tag{188}$$

$$\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] = \left[1 - \left(\mu_k^\dagger + \hat{f}_{\mu_{k-1}}^{\mu_k}(q) \right)^2 \right] \left(\frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}} \right)^2, \quad i \in [n]. \tag{189}$$

We would like to find an upper bound on the summation $\sum_{k=1}^d \mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right]$. Define $A \equiv A(\beta, q) \triangleq (1-2q)^2 [1 - \tanh(\beta)(1 - (1-2q)^2)]$, for all $\beta > 0$ and $q \in [0, 1)$. Then

$$\begin{aligned}
 &\left[1 - \left(\mu_k^\dagger + \hat{f}_{\mu_{k-1}}^{\mu_k}(q) \right)^2 \right] \\
 &= \left[1 - \left(\mu_k^\dagger + \mu_{k-1}^\dagger (\mu_k - \mu_k^\dagger) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \right)^2 \right] \\
 &= \left[1 - \mu_k^2 \left((1-2q)^2 + \mu_{k-1}^\dagger (1 - (1-2q)^2) \frac{\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger}{1 - (\mu_{k-1}^\dagger)^2} \right)^2 \right]
 \end{aligned}$$

$$\begin{aligned}
 &\leq \left[1 - \mu_k^2 \left((1-2q)^2 - |\mu_{k-1}^\dagger| (1 - (1-2q)^2) \frac{|\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger|}{1 - (\mu_{k-1}^\dagger)^2} \right)^2 \right] \\
 &\leq \left[1 - \mu_k^2 ((1-2q)^2 - (1-2q)^2 \tanh(\beta)(1 - (1-2q)^2))^2 \right] \tag{190} \\
 &\leq [1 - \mu_k^2 A(\beta, q)], \tag{191}
 \end{aligned}$$

and (190) holds because $|\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger| \leq \gamma_j^\dagger \leq (1-2q)^2(1 - \mu_j^2)/6 \leq (1-2q)^2(1 - (\mu_j^\dagger)^2)/6$. Then (189) and (191) give

$$\begin{aligned}
 &\sum_{k=1}^d \mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] \\
 &\leq \sum_{k=1}^d [1 - \mu_k^2 A(\beta, q)] \left(\frac{\prod_{j=1}^{k-1} \hat{\mu}_j^\dagger \prod_{j=k+1}^d \mu_j^\dagger}{(1-2q)^{2d}} \right)^2 \\
 &\leq \frac{1}{(1-2q)^2} \sum_{k=1}^d [1 - \mu_k^2 A(\beta, q)] \prod_{j=1, j \neq k}^d \left[\mu_j^2 + 2 \frac{\gamma_j^\dagger}{(1-2q)^2} \right]. \tag{192}
 \end{aligned}$$

The inequality (192) holds under the event $E_\dagger^{\text{edge}}(\mathcal{E}_T)$ defined in (158) (see Lemma 24) because

$$\left(\frac{\hat{\mu}_j^\dagger}{(1-2q)^2} \right)^2 \leq \left(\frac{\mu_j^\dagger}{(1-2q)^2} \right)^2 + 2 \frac{\gamma_j^\dagger}{(1-2q)^2}, \tag{193}$$

since $|\mu_j^\dagger| \leq (1-2q)^2$, $|\hat{\mu}_j^\dagger| \leq (1-2q)^2$ under the assumption of known q . Next we define $x_j \triangleq \mu_j^2 + 2 \frac{\gamma_j^\dagger}{(1-2q)^2}$, then $3(1 - A(\beta, q)x_j)/2 + (A(\beta, q) - 1)/2 \geq 1 - \mu_k^2 A(\beta, q)$ and (192) gives

$$\begin{aligned}
 &\sum_{k=1}^d \mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right] \\
 &\leq \frac{1}{(1-2q)^2} \sum_{k=1}^d \left[\frac{3}{2}(1 - Ax_j) + \frac{A-1}{2} \right] \prod_{j=1, j \neq k}^d x_j \\
 &\leq \frac{d}{(1-2q)^2} \left[\frac{3}{2}(1 - Ax) + \frac{A-1}{2} \right] x^{d-1}. \tag{194}
 \end{aligned}$$

The latter is maximized at $x^* = (A+2)(1 - \frac{1}{d})/3$ and $A \in (0, 1]$, thus we have

$$\begin{aligned}
 &\frac{d}{(1-2q)^2} \left[\frac{3}{2}(1 - Ax) + \frac{A-1}{2} \right] x^{d-1} \\
 &\leq \frac{d}{(1-2q)^2} \left[\frac{3}{2}(1 - A \frac{A+2}{3} + A \frac{A+2}{3d}) + \frac{A-1}{2} \right] \left(\frac{A+2}{3} \right)^{d-1} \left(1 - \frac{1}{d} \right)^{d-1} \\
 &= \frac{d}{(1-2q)^2} \left[\frac{3}{2} - \frac{A(A+2)}{2} + \frac{A-1}{2} \right] (x^*)^{d-1} + \frac{A(A+2)}{2(1-2q)^2} (x^*)^{d-1}
 \end{aligned}$$

$$\begin{aligned}
 &= d \frac{2 - A^2 - A}{2(1 - 2q)^2} \left(\frac{A + 2}{3} \right)^{d-1} \left(1 - \frac{1}{d} \right)^{d-1} + \frac{A(A + 2)}{2(1 - 2q)^2} \left(\frac{A + 2}{3} \right)^{d-1} \left(1 - \frac{1}{d} \right)^{d-1} \\
 &\leq d \frac{2 - A^2 - A}{2(1 - 2q)^2} \left(\frac{A + 2}{3} \right)^{d-1} \left(1 - \frac{1}{2} \right) + \frac{A(A + 2)}{2(1 - 2q)^2} \left(\frac{A + 2}{3} \right) \left(1 - \frac{1}{2} \right) \\
 &= d \frac{(A + 2)(1 - A)}{4(1 - 2q)^2} \left(\frac{A + 2}{3} \right)^{d-1} + \frac{A(A + 2)^2}{12(1 - 2q)^2} \\
 &\leq d \frac{3(1 - A)}{4(1 - 2q)^2} \left(\frac{A + 2}{3} \right)^d + \frac{3}{4(1 - 2q)^2} \triangleq G(\beta, q) \tag{195} \\
 &\leq \frac{(1 - A)}{4(1 - 2q)^2} \frac{3}{e \log \frac{3}{A+2}} + \frac{3}{4(1 - 2q)^2} \\
 &= \frac{3(1 - A)}{e4(1 - 2q)^2 \log \frac{3}{A+2}} + \frac{3}{4(1 - 2q)^2} \\
 &\leq \frac{3(3e^{-1} + 1)}{4(1 - 2q)^2}.
 \end{aligned}$$

In (195) we define the function $G(\beta, q)$ and we proved that it has an upper bound independent of $d \in [p - 1]$,

$$G(\beta, q) \triangleq d \frac{3(1 - A(\beta, q))}{4(1 - 2q)^2} \left(\frac{A(\beta, q) + 2}{3} \right)^d + \frac{3}{4(1 - 2q)^2} \leq \frac{3(3e^{-1} \mathbf{1}_{q \neq 0} + 1)}{4(1 - 2q)^2} \equiv G. \tag{196}$$

For the rest of the proof and the final result $G(\beta, q)$ can be replaced by its upper bound in (196), however the definition of $G(\beta, q)$ shows the continuity of the result for $q \rightarrow 0$.

The following inequality holds with probability 1 for all $i \in [n]$ and $k \in [d]$ under the event $E_{\dagger}^{\text{edge}}(\mathcal{E}_T)$ (Lemma 24),

$$\begin{aligned}
 |\xi_k^{(i)}| &\leq 2 + \left| \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right| \\
 &= 2 + \frac{|\hat{f}_{\mu_{k-1}}^{\mu_k}(q)|}{(1 - 2q)^2} \frac{\prod_{j=1}^{k-1} |\hat{\mu}_j^{\dagger}| \prod_{j=k+1}^d |\mu_j^{\dagger}|}{(1 - 2q)^{2d-2}} \\
 &\leq 2 + \frac{|\hat{f}_{\mu_{k-1}}^{\mu_k}(q)|}{(1 - 2q)^2} \\
 &= 2 + \frac{1}{(1 - 2q)^2} |\mu_{k-1}^{\dagger}| |(\mu_k - \mu_k^{\dagger})| \frac{|\hat{\mu}_{k-1}^{\dagger} - \mu_{k-1}^{\dagger}|}{1 - (\mu_{k-1}^{\dagger})^2} \\
 &= 2 + \tanh^2(\beta) (1 - (1 - 2q)^2) \frac{\gamma_{k-1}^{\dagger}}{1 - (\mu_{k-1}^{\dagger})^2} \\
 &\leq 2 + \frac{(1 - 2q)^2}{6} (1 - (1 - 2q)^2) \tanh^2(\beta) \triangleq S(\beta, q). \tag{197}
 \end{aligned}$$

the last step comes from the inequality $\gamma_e^\dagger \leq \frac{(1-2q)^2}{6}(1 - (\mu_e)^2)$ (161), which holds if the inequality $n > \frac{108e^{2\beta}}{K(\beta, q)(1-2q)^4} \log(4p)$ holds (see 159). Recall that Lemma 24 gives

$$\gamma_e = \sqrt{3 \frac{1 - \mu_e^2}{nK(\beta, q)} \log(2p^3/\delta)} \leq \sqrt{3 \frac{1 - \tanh^2(\beta)}{nK(\beta, q)} \log(2p^3/\delta)} \leq \sqrt{\frac{3 \log(2p^3/\delta)}{n}} \equiv \tilde{\gamma}. \quad (198)$$

Also, for all $i \in [n]$ we have

$$\begin{aligned} & \left| \sum_{k=1}^d \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right| \\ &= \left| \sum_{k=2}^d \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right| \\ &\leq \sum_{k=2}^d \left| \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right| \\ &= \sum_{k=2}^d \frac{|\hat{f}_{\mu_{k-1}}^{\mu_k}(q)| \prod_{j=1}^{k-1} |\hat{\mu}_j^\dagger| \prod_{j=k+1}^d |\mu_j^\dagger|}{(1-2q)^2 (1-2q)^{2d-2}} \\ &\leq \frac{1}{(1-2q)^2} \sum_{k=2}^d |\mu_{k-1}^\dagger| (\mu_k - \mu_k^\dagger) \frac{|\hat{\mu}_{k-1}^\dagger - \mu_{k-1}^\dagger|}{1 - (\mu_{k-1}^\dagger)^2} \prod_{j=1}^{k-1} \frac{|\hat{\mu}_j^\dagger|}{(1-2q)^2} \prod_{j=k+1}^d \frac{|\mu_j^\dagger|}{(1-2q)^2} \\ &\leq \tanh^2(\beta) (1 - (1-2q)^2) \frac{\gamma_e}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1}^{k-1} \frac{|\hat{\mu}_j^\dagger|}{(1-2q)^2} \prod_{j=k+1}^d \frac{|\mu_j^\dagger|}{(1-2q)^2} \quad (199) \end{aligned}$$

$$\leq \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1}^{k-1} \frac{|\hat{\mu}_j^\dagger|}{(1-2q)^2} \prod_{j=k+1}^d \frac{|\mu_j^\dagger|}{(1-2q)^2} \quad (200)$$

$$\leq \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1}^{k-1} \frac{|\mu_j^\dagger| + \gamma_j}{(1-2q)^2} \prod_{j=k+1}^d \frac{|\mu_j^\dagger|}{(1-2q)^2} \quad (201)$$

$$= \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1}^{k-1} \tanh(\beta) + \frac{\gamma_j}{(1-2q)^2} \prod_{j=k+1}^d \tanh(\beta)$$

$$\leq \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1, j \neq k}^d \tanh(\beta) + \frac{\gamma_j}{(1-2q)^2}$$

$$\leq \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \sum_{k=2}^d \prod_{j=1, j \neq k}^d \left(\tanh(\beta) + \frac{1}{6} (1 - \tanh^2(\beta)) \right)$$

$$= \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} (d-1) \left(\frac{5}{3} - \frac{1}{6} (\tanh(\beta) - 3)^2 \right)^{d-1}$$

$$\leq \tanh^2(\beta) \frac{(1 - (1-2q)^2) \tilde{\gamma}}{1 - (1-2q)^4 \tanh^2(\beta)} \frac{1}{-e \log \left(\frac{5}{3} - \frac{1}{6} (\tanh(\beta) - 3)^2 \right)}$$

$$\begin{aligned}
 &\leq \frac{(1 - (1 - 2q)^2) \tanh^2(\beta) \tilde{\gamma}}{1 - (1 - 2q)^4 \tanh^2(\beta)} e^{2\beta-1}. \\
 &\leq \frac{(1 - (1 - 2q)^2) \tanh^2(\beta) e^{2\beta}}{1 - (1 - 2q)^4 \tanh^2(\beta)} \sqrt{\frac{3 \log(2p^3/\delta)}{n}} \triangleq \Delta,
 \end{aligned} \tag{202}$$

where (199), (200), (201) come from Lemma 24 and (198). Finally, $0 < \frac{5}{3} - \frac{1}{6}(\tanh(\beta) - 3)^2 < 1$ for all $\beta > 0$ and $-1/\log(\frac{5}{3} - \frac{1}{6}(\tanh(\beta) - 3)^2) \leq e^{2\beta}$.

We use the symbol $\mathbb{E}_{\mathcal{A}_{k-1}}[\cdot]$ to denote the conditional expectation given the event $\mathbb{E}_{\dagger}^{\text{edge}}(\mathcal{A}_{k-1})$, for instance

$$\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k M_i^{\dagger} \right) \middle| \mathbb{E}_{\dagger}^{\text{edge}}(\mathcal{A}_{k-1}) \right] \equiv \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^k M_i^{\dagger} \right) \right]. \tag{203}$$

Further we define the function $F(\cdot, \cdot)$ as

$$F(t, \lambda) = \log \left(\frac{1}{1+t} e^{-\lambda t} + \frac{t}{1+t} e^{\lambda} \right). \tag{204}$$

For any $k \leq d$ we have

$$\begin{aligned}
 &\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k M_i^{\dagger} \right) \middle| \mathbb{E}_{\dagger}^{\text{edge}}(\mathcal{A}_{k-1}) \right] \\
 &= \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^{\dagger} \right) \mathbb{E} \left[\exp \left(\lambda M_k^{\dagger} \right) \middle| \hat{\mu}_1^{\dagger}, \dots, \hat{\mu}_{k-1}^{\dagger} \right] \right]
 \end{aligned} \tag{205}$$

$$= \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^{\dagger} \right) \mathbb{E} \left[\exp \left(\lambda \frac{1}{n} \sum_{i=1}^n Z_k^{(i)} \right) \middle| \hat{\mu}_1^{\dagger}, \dots, \hat{\mu}_{k-1}^{\dagger} \right] \right] \tag{206}$$

$$\begin{aligned}
 &= \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^{\dagger} \right) \right. \\
 &\quad \times \mathbb{E} \left[\exp \left(\lambda \frac{1}{n} \sum_{i=1}^n \xi_k^{(i)} \right) \exp \left(\lambda \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right) \middle| \hat{\mu}_1^{\dagger}, \dots, \hat{\mu}_{k-1}^{\dagger} \right] \left. \right]
 \end{aligned} \tag{207}$$

$$\begin{aligned}
 &\leq \exp \left(\lambda \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right) \\
 &\quad \times \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^{\dagger} \right) \mathbb{E} \left[\exp \left(\lambda \frac{1}{n} \sum_{i=1}^n \xi_k^{(i)} \right) \middle| \hat{\mu}_1^{\dagger}, \dots, \hat{\mu}_{k-1}^{\dagger} \right] \right]
 \end{aligned} \tag{208}$$

$$\begin{aligned}
 &\leq \exp \left(\lambda \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right) \\
 &\quad \times \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^{\dagger} \right) \exp \left\{ n F \left(\frac{\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right]}{S(\beta, q)^2}, |\lambda| \frac{S(\beta, q)}{n} \right) \right\} \right]
 \end{aligned} \tag{209}$$

$$\begin{aligned}
 &\leq \exp \left(\lambda \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right) \\
 &\quad \times \exp \left\{ n F \left(\frac{\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 | \mathcal{F}_{i-1}^k \right]}{S(\beta, q)^2}, |\lambda| \frac{S(\beta, q)}{n} \right) \right\} \mathbb{E}_{\mathcal{A}_{k-1}} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \right]. \quad (210)
 \end{aligned}$$

The equation (205) comes from change of measure and tower property, the definitions (176) and (177) of M_k^\dagger and $\xi_k^{(i)}$ respectively give (206) and (207). The (208) is derived by upper bounding the quantity $\left| \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right|$ similarly to (197), (209) is the upper bound on the moment generating function of the supermartingale [Fan et al. \(2012\)](#). To get a recurrence we proceed as follows:

$$\begin{aligned}
 &\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \right] \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \right) \\
 &= \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \cap \mathbb{E}_{\text{ek-1}, \dagger}^{\text{edge}} \right] \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \right) \\
 &= \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \cap \mathbb{E}_{\text{ek-1}, \dagger}^{\text{edge}} \right] \\
 &\quad \times \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right) \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right) \\
 &= \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \cap \mathbb{E}_{\text{ek-1}, \dagger}^{\text{edge}} \right] \\
 &\quad \times \mathbb{P} \left(\mathbb{E}_{\text{ek-1}, \dagger}^{\text{edge}} \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right) \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right) \\
 &= \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \mathbf{1}_{\mathbb{E}_{\text{ek-1}, \dagger}^{\text{edge}}} \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right] \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right) \\
 &\leq \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^{k-1} M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right] \mathbb{P} \left(\mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-2}) \right). \quad (211)
 \end{aligned}$$

By applying the recurrence d times, we derive the following bound

$$\begin{aligned}
 &\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k M_i^\dagger \right) \middle| \mathbb{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \right] \\
 &\leq \exp \left(\lambda \sum_{k=1}^d \mathbb{E} \left[Z_k^{(i)} | \mathcal{F}_{i-1}^k \right] \right)
 \end{aligned}$$

$$\begin{aligned}
 & \times \exp \left\{ n \sum_{k=1}^d F \left(\frac{\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right]}{S(\beta, q)^2}, |\lambda| \frac{S(\beta, q)}{n} \right) \right\} \\
 & \leq \exp \left(\lambda \sum_{k=1}^d \mathbb{E} \left[Z_k^{(i)} \middle| \mathcal{F}_{i-1}^k \right] \right) \\
 & \times \exp \left\{ n d F \left(\frac{1}{d} \sum_{k=1}^d \frac{\mathbb{E} \left[\left(\xi_k^{(i)} \right)^2 \middle| \mathcal{F}_{i-1}^k \right]}{S(\beta, q)^2}, |\lambda| \frac{S(\beta, q)}{n} \right) \right\}. \tag{212}
 \end{aligned}$$

Further (196), (197), (202) and (212) give

$$\begin{aligned}
 & \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k M_i^\dagger \right) \middle| \mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{k-1}) \right] \\
 & \leq \exp(\lambda \Delta(\beta, q)) \exp \left\{ nd F \left(\frac{1}{d} \frac{G(\beta, q)}{S_\xi^2(\beta, q)}, |\lambda| \frac{S(\beta, q)}{n} \right) \right\} \\
 & \leq \exp \left\{ nd F \left(\frac{G(\beta, q)}{d}, |\lambda| \frac{S(\beta, q)}{n} \right) + \lambda \Delta(\beta, q) \right\}. \tag{213}
 \end{aligned}$$

For sake of space, we denote the functions $G(\beta, q)$, $S(\beta, q)$, and $\Delta(\beta, q)$ as G , S , and Δ respectively. It is true that

$$\begin{aligned}
 & \mathbb{E} \left[\exp \left(\lambda \sum_{k=1}^d M_k^\dagger \right) \middle| \mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right] \mathbb{P} \left[\mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \cap \mathcal{E}_q \right] \\
 & \leq \exp \left\{ nd F \left(\frac{G}{d}, |\lambda| \frac{S}{n} \right) + \lambda \Delta \right\}. \tag{214}
 \end{aligned}$$

Under the assumption $n > \frac{108e^{2\beta}}{(1-2q)^4 K(\beta, q)} \log(4p)$, we have

$$\mathbb{P} \left[\left(\mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right)^c \right] \leq \frac{1}{2}. \tag{215}$$

The latter gives

$$\mathbb{E} \left[\exp \left(\lambda \sum_{k=1}^d M_k^\dagger \right) \middle| \mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right] \leq 2 \exp \left\{ nd F \left(\frac{G}{d}, |\lambda| \frac{S}{n} \right) + \lambda \Delta \right\}, \tag{216}$$

which implies that

$$\mathbb{P} \left[\sum_{k=1}^d M_k^\dagger \geq \gamma \middle| \mathcal{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right] \leq 2 \min_{\lambda > 0} \exp \left\{ nd F \left(\frac{G}{d}, \lambda \frac{S}{n} \right) + \lambda \Delta - \lambda \gamma \right\}$$

$$= 2 \min_{\lambda > 0} \exp \left\{ nd F \left(\frac{G}{d}, \lambda \frac{S}{n} \right) + \lambda (\Delta - \gamma) \right\}, \quad (217)$$

and we define $\gamma' \triangleq \gamma - \Delta$. The minimum value is attained at

$$\lambda^* = \frac{n/S}{1 + G/d} \log \frac{1 + \frac{\gamma'}{GS}}{1 - \frac{\gamma'}{dS}} \quad (218)$$

and by substituting the optimal value we get

$$\begin{aligned} & \exp \left\{ nd F \left(\frac{G}{d}, \lambda^* \frac{S}{n} \right) - \lambda^* \gamma' \right\} \\ &= \left[\frac{1}{1 + \frac{G}{d}} \left(\frac{1 + \frac{\gamma'}{GS}}{1 - \frac{\gamma'}{dS}} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} + \frac{\frac{G}{d}}{1 + \frac{G}{d}} \left(\frac{1 + \frac{\gamma'}{GS}}{1 - \frac{\gamma'}{dS}} \right)^{\frac{1}{1+G/d} \left(1 - \frac{\gamma'}{dS} \right)} \right]^{nd} \\ &= \left[\frac{1}{1 + \frac{G}{d}} \left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)}} \right. \\ & \quad \left. + \frac{\frac{G}{d}}{1 + \frac{G}{d}} \left(1 + \frac{\gamma'}{GS} \right)^{\frac{1}{1+G/d} \left(1 - \frac{\gamma'}{dS} \right)} \left(1 - \frac{\gamma'}{dS} \right)^{-\frac{1}{1+G/d} \left(1 - \frac{\gamma'}{dS} \right)} \right]^{nd} \\ &= \left[\frac{1}{1 + \frac{G}{d}} \left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)} - 1} \left(1 - \frac{\gamma'}{dS} \right) \right. \\ & \quad \left. + \frac{\frac{G}{d}}{1 + \frac{G}{d}} \left(1 + \frac{\gamma'}{GS} \right)^{\frac{1}{1+G/d} \left(1 - \frac{\gamma'}{dS} \right) - 1} \left(1 + \frac{\gamma'}{GS} \right) \left(1 - \frac{\gamma'}{dS} \right)^{-\frac{1}{1+G/d} \left(1 - \frac{\gamma'}{dS} \right)} \right]^{nd} \\ &= \left[\frac{1}{1 + \frac{G}{d}} \left(1 - \frac{\gamma'}{dS} \right) + \frac{\frac{G}{d}}{1 + \frac{G}{d}} \left(1 + \frac{\gamma'}{GS} \right) \right]^{nd} \\ & \quad \times \left[\left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)} - 1} \right]^{nd} \\ &= 1^{nd} \times \left[\left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)} - 1} \right]^{nd} \\ &= \left[\left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{-1}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)}} \right]^{nd}. \end{aligned} \quad (219)$$

Then (217) and (219) give

$$\mathbb{P} \left[\sum_{k=1}^d M_k^\dagger \geq \gamma \middle| \mathbf{E}_\dagger^{\text{edge}}(\mathcal{A}_{d-1}) \right]$$

$$\leq 2 \left[\left(1 + \frac{\gamma'}{GS} \right)^{-\frac{G/d}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)}} \left(1 - \frac{\gamma'}{dS} \right)^{\frac{-1}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)}} \right]^{nd}. \quad (220)$$

As a final step we want to express the upper bound as an exponential function of γ , we define $\zeta \triangleq \gamma'/(SG)$ and we proceed as follows:

$$\begin{aligned} & d \left[\left(\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)} \right) \log \left(1 + \frac{\gamma'}{GS} \right) + \left(\frac{1}{1+G/4d} - \frac{\gamma'/(dy)}{(1+G/4d)} \right) \log \left(1 - \frac{\gamma'}{dy} \right) \right] \\ & \geq d \left[\left(\frac{G/d}{1+G/d} + \frac{\gamma'/(dS)}{(1+G/d)} \right) \left(\frac{\gamma'}{GS} - \frac{1}{2} \left[\frac{\gamma'}{GS} \right]^2 \right) \right. \\ & \quad \left. - \left(\frac{1}{1+G/d} - \frac{\gamma'/(dS)}{(1+G/d)} \right) \log \left(1 + \frac{\gamma'/dS}{1-\gamma'/dS} \right) \right] \\ & \geq \left[\left(\frac{dG}{d+G} + \frac{d\gamma'/S}{(d+G)} \right) \left(\frac{\gamma'}{GS} - \frac{1}{2} \left[\frac{\gamma'}{GS} \right]^2 \right) \right. \\ & \quad \left. - \left(\frac{d^2}{d+G} - \frac{d\gamma'/S}{(d+G)} \right) \left(\frac{\gamma'/dS}{1-\gamma'/dS} - \frac{1}{2} \left(\frac{\gamma'/dS}{1-\gamma'/dS} \right)^2 \right) \right] \\ & = \frac{d}{d+G} \left[(G + \gamma'/S) \left(\frac{\gamma'}{GS} - \frac{1}{2} \left[\frac{\gamma'}{GS} \right]^2 \right) \right. \\ & \quad \left. - (d - \gamma'/S) \left(\frac{\gamma'/dS}{1-\gamma'/dS} - \frac{1}{2} \left(\frac{\gamma'/dS}{1-\gamma'/dS} \right)^2 \right) \right] \\ & \geq \frac{d}{d+G} \left[(G + \gamma'/S) \left(\frac{\gamma'}{GS} - \frac{1}{2} \left[\frac{\gamma'}{GS} \right]^2 \right) - \gamma'/S \right] \\ & = \frac{dG}{d+G} \left[(1 + \gamma'/(SG)) \left(\frac{\gamma'}{GS} - \frac{1}{2} \left[\frac{\gamma'}{GS} \right]^2 \right) - \gamma'/(SG) \right] \\ & = \frac{dG}{d+G} \left[(1 + \zeta) \left(\zeta - \frac{1}{2} \zeta^2 \right) - \zeta \right] \\ & \geq \frac{2G}{2+G} \left[(1 + \zeta) \left(\zeta - \frac{1}{2} \zeta^2 \right) - \zeta \right] \\ & \geq \left(\frac{3\gamma'}{10GS} \right)^2, \quad \forall \gamma' \in (0, \frac{SG}{3}). \end{aligned} \quad (221)$$

Recall that $\zeta \triangleq \gamma'/(SG)$, $\gamma' = \gamma - \Delta$. If $\Delta < \gamma \leq S(\beta, q)G(\beta, q)/3 + \Delta$ then (220) and (221) give

$$\mathbb{P} \left[\sum_{k=1}^d M_k^\dagger \geq \gamma \middle| \mathbf{E}_+^{\text{edge}}(\mathcal{A}_{d-1}) \right] \leq 2 \exp \left(-0.3^2 n \frac{(\gamma - \Delta)^2}{S^2(\beta, q)G^2(\beta, q)} \right). \quad (222)$$

In a similar way we derive the bound

$$\mathbb{P} \left[\sum_{k=1}^d M_k^\dagger \leq -\gamma |E_{\dagger}^{\text{edge}}(\mathcal{A}_{d-1})| \right] \leq 2 \exp \left(-0.3^2 n \frac{(\gamma - \Delta)^2}{S^2(\beta, q) G^2(\beta, q)} \right). \quad (223)$$

Finally, we combine (174), Lemma 24, (222) and (223) to derive the bound (170) which guarantees that

$$\mathbb{P} \left[\left| \prod_{i=1}^d \frac{\hat{\mu}_i^\dagger}{(1-2q)^2} - \prod_{i=1}^d \frac{\mu_i^\dagger}{(1-2q)^2} \right| > \gamma \right] \leq 2\delta/p^2, \quad \forall d \geq 2. \quad (224)$$

To summarize we proved that the event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ happens with probability at least $1 - 2\delta/p^2$ by combining Bresler's and Karzand's technique, the Corollary 2.3 by [Fan et al. \(2012\)](#) and Lemma 23. \blacksquare

Appendix E. Predictive Learning, Proof of Theorem 3 and Theorem 7

Recall that our goal is to guarantee that the quantity $\mathcal{L}^{(2)}(p(\cdot), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger}))$ is smaller than a number $\eta > 0$ with probability at least $1 - \delta$. To do this, we use the triangle inequality as

$$\mathcal{L}^{(2)}(p(\cdot), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger})) \leq \mathcal{L}^{(2)}(p(\cdot), \Pi_{\dagger}^{\text{CL}}(p(\cdot))) + \mathcal{L}^{(2)}(\Pi_{\dagger}^{\text{CL}}(p(\cdot)), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger})) \quad (225)$$

and we find the required number of samples such that each of the terms $\mathcal{L}^{(2)}(p(\cdot), \Pi_{\dagger}^{\text{CL}}(p(\cdot)))$ and $\mathcal{L}^{(2)}(\Pi_{\dagger}^{\text{CL}}(p(\cdot)), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger}))$ in (225) is less than $\eta/2$ with probability at least $1 - \delta$. The next Lemma provides the necessary bounds on γ_{\dagger} and ϵ_{\dagger} that guarantee

$$\mathcal{L}^{(2)}(\Pi_{\dagger}^{\text{CL}}(p(\cdot)), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger})) \leq \eta/2.$$

Lemma 26 *If $\gamma_{\dagger} \leq \frac{\eta}{3}$ and*

$$\epsilon_{\dagger} \leq (1-2q)^2 e^{-\beta} \left[20 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right) \right]^{-1},$$

then $\mathcal{L}^{(2)}(\Pi_{\dagger}^{\text{CL}}(p(\cdot)), \Pi_{\dagger}^{\text{CL}}(\hat{p}_{\dagger})) \leq \eta/2$ under the event $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger}) \cap E_{\dagger}^{\text{cascade}}(\gamma_{\dagger}) \cap E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$.

Proof The derivation of the bound is similar to the approach by [Bresler and Karzand \(2020, Section 6.1\)](#) but with different calculations. In the hidden model, we consider the path between two nodes i, j in the estimated structure T_{\dagger}^{CL} , namely $\text{path}_{T_{\dagger}^{\text{CL}}}(i, j)$, to be $(\mathcal{F}_0, e_1, \mathcal{F}_1, e_1, \dots, \mathcal{F}_{t-1}, e_t, \mathcal{F}_t)$, and \mathcal{F}_i are segments with all strong edges and e_i are all weak edges. We consider the case of at least one weak edge to exist in the path. If there is no weak edge the bound reduces to the case of Lemma 25. The length of each sub-path \mathcal{F}_i is denoted as d_i , for all $i \in \{0, 1, \dots, t\}$. Each segment (sub-path) \mathcal{F}_i has exactly d_i edges,

and the total number of edges in the path are d ; thus $d = \sum_{i=0}^t d_i + t$. Note that $t \geq 1$ and $d_i \geq 0$ for all $i \in \{0, 1, \dots, t\}$. Recall that

$$\Pi_{T_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)) = \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_{T_{\dagger}^{\text{CL}}}} \frac{1 + x_i x_j \mathbb{E}[X_i X_j]}{2} = \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_{T_{\dagger}^{\text{CL}}}} \frac{1 + x_i x_j \frac{\mathbb{E}[Y_i Y_j]}{(1-2q)^2}}{2} \quad (226)$$

(the latter comes from (19)), and

$$\Pi_{T_{\dagger}^{\text{CL}}}(\hat{\mathbf{p}}_{\dagger}) \triangleq \frac{1}{2} \prod_{(i,j) \in \mathcal{E}_{T_{\dagger}^{\text{CL}}}} \frac{1 + x_i x_j \frac{\hat{\mathbb{E}}[Y_i Y_j]}{(1-2q)^2}}{2}. \quad (227)$$

Further, for any tree-structured Ising model distributions P, \tilde{P} with structures $T = (\mathcal{V}, \mathcal{E})$ and $\tilde{T} = (\mathcal{V}, \tilde{\mathcal{E}})$ respectively, we have

$$\mathcal{L}^{(2)}(P, \tilde{P}) \triangleq \sup_{i,j \in \mathcal{V}} \frac{1}{2} \sum_{x_i, x_j \in \{-1, +1\}^2} \left| P(x_i, x_j) - \tilde{P}(x_i, x_j) \right| \quad (228)$$

$$= \sup_{i,j \in \mathcal{V}} \frac{1}{2} \left| \prod_{e \in \text{path}_T(i,j)} \mu_e - \prod_{e' \in \text{path}_{\tilde{T}}(i,j)} \tilde{\mu}_{e'} \right|. \quad (229)$$

To upper bound the quantity $\mathcal{L}^{(2)}(\Pi_{T_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{\mathbf{p}}_{\dagger}))$ we have

$$\begin{aligned} & 2\mathcal{L}^{(2)}(\Pi_{T_{\dagger}^{\text{CL}}}(\mathbf{p}(\cdot)), \Pi_{T_{\dagger}^{\text{CL}}}(\hat{\mathbf{p}}_{\dagger})) \\ &= \left| \prod_{e \in \text{path}_{T_{\dagger}^{\text{CL}}}(i,j)} \frac{\mu_e^{\dagger}}{(1-2q)^2} - \prod_{e \in \text{path}_{T_{\dagger}^{\text{CL}}}(i,j)} \frac{\hat{\mu}_e^{\dagger}}{(1-2q)^2} \right| \\ &= \frac{1}{(1-2q)^{2d}} \left| \hat{\mu}_{\mathcal{F}_0}^{\dagger} \prod_{i=1}^t \hat{\mu}_{\mathcal{F}_i}^{\dagger} \hat{\mu}_{e_i}^{\dagger} - \mu_{\mathcal{F}_0}^{\dagger} \prod_{i=1}^t \mu_{\mathcal{F}_i}^{\dagger} \mu_{e_i}^{\dagger} \right| \end{aligned} \quad (230)$$

$$\begin{aligned} &\leq \frac{1}{(1-2q)^{2d}} \left[\left| \hat{\mu}_{\mathcal{F}_0}^{\dagger} - \mu_{\mathcal{F}_0}^{\dagger} \right| \prod_{j=1}^t \left| \mu_{\mathcal{F}_j}^{\dagger} \mu_{e_j}^{\dagger} \right| \right. \\ &\quad \left. + \sum_{i=1}^t \left| \hat{\mu}_{\mathcal{F}_i}^{\dagger} \hat{\mu}_{e_i}^{\dagger} - \mu_{\mathcal{F}_i}^{\dagger} \mu_{e_i}^{\dagger} \right| \left| \hat{\mu}_{\mathcal{F}_0}^{\dagger} \right| \prod_{j=1}^{i-1} \left| \hat{\mu}_{\mathcal{F}_j}^{\dagger} \hat{\mu}_{e_j}^{\dagger} \right| \prod_{k=i+1}^t \left| \mu_{\mathcal{F}_k}^{\dagger} \mu_{e_k}^{\dagger} \right| \right] \end{aligned} \quad (231)$$

$$\begin{aligned} &= \frac{\left| \hat{\mu}_{\mathcal{F}_0}^{\dagger} - \mu_{\mathcal{F}_0}^{\dagger} \right|}{(1-2q)^{2d_0}} \prod_{j=1}^t \left| \frac{\mu_{\mathcal{F}_j}^{\dagger}}{(1-2q)^{2d_j}} \frac{\mu_{e_j}^{\dagger}}{(1-2q)^{2t}} \right| \\ &\quad + \sum_{i=1}^t \frac{\left| \hat{\mu}_{\mathcal{F}_i}^{\dagger} \hat{\mu}_{e_i}^{\dagger} - \mu_{\mathcal{F}_i}^{\dagger} \mu_{e_i}^{\dagger} \right|}{(1-2q)^{2(d_i+1)}} \frac{\left| \hat{\mu}_{\mathcal{F}_0}^{\dagger} \right|}{(1-2q)^{2d_0}} \frac{\prod_{j=1}^{i-1} \left| \hat{\mu}_{\mathcal{F}_j}^{\dagger} \hat{\mu}_{e_j}^{\dagger} \right| \prod_{k=i+1}^t \left| \mu_{\mathcal{F}_k}^{\dagger} \mu_{e_k}^{\dagger} \right|}{(1-2q)^{2(d-d_i-d_0-1)}} \end{aligned} \quad (232)$$

$$\leq \gamma_{\dagger} \left(\frac{\tau^{\dagger}}{(1-2q)^2} \right)^{t-1} + \left(\frac{\tau^{\dagger} + \epsilon_{\dagger}}{(1-2q)^2} \right)^{t-1} \sum_{i=1}^t \frac{|\hat{\mu}_{\mathcal{F}_i}^{\dagger} \hat{\mu}_{e_i}^{\dagger} - \mu_{\mathcal{F}_i}^{\dagger} \mu_{e_i}^{\dagger}|}{(1-2q)^{2(d_i+1)}} \quad (233)$$

$$\leq \gamma_{\dagger} \left(\frac{\tau^{\dagger}}{(1-2q)^2} \right)^{t-1} + \left(\frac{\tau^{\dagger} + \epsilon_{\dagger}}{(1-2q)^2} \right)^{t-1} \sum_{i=1}^t \left(\frac{|\hat{\mu}_{\mathcal{F}_i}^{\dagger} - \mu_{\mathcal{F}_i}^{\dagger}| |\hat{\mu}_{e_i}^{\dagger}|}{(1-2q)^{2(d_i+1)}} + \frac{|\mu_{\mathcal{F}_i}^{\dagger} (\hat{\mu}_{e_i}^{\dagger} - \mu_{e_i}^{\dagger})|}{(1-2q)^{2(d_i+1)}} \right) \quad (234)$$

$$\leq \left(\frac{\tau^{\dagger} + \epsilon_{\dagger}}{(1-2q)^2} \right)^{t-1} (2t+1) \max \left\{ \gamma_{\dagger}, \frac{\epsilon_{\dagger}}{(1-2q)^2} \right\} \leq \left(\frac{4\epsilon_{\dagger} e^{\beta} (1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta}) + \epsilon_{\dagger}}{(1-2q)^2} \right)^{t-1} (2t+1) \max \left\{ \gamma_{\dagger}, \frac{\epsilon_{\dagger}}{(1-2q)^2} \right\} \quad (235)$$

$$\leq \left(\frac{5\epsilon_{\dagger} e^{\beta}}{(1-2q)^2} \right)^{t-1} (1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta})^{t-1} (2t+1) \max \left\{ \gamma_{\dagger}, \frac{\epsilon_{\dagger}}{(1-2q)^2} \right\} \leq \frac{2t+1}{4^{t-1}} \frac{\eta}{3} \quad (236)$$

$$\leq \eta. \quad (237)$$

Telescoping summation and triangle inequality give (230) and (231). We use the definition of d , $d = \sum_{i=0}^t d_i + t$ to get (232). The inequalities $|\mu_{\mathcal{F}_i}^{\dagger}| \leq (1-2q)^{2d_i}$, $|\hat{\mu}_{\mathcal{F}_i}^{\dagger}| \leq (1-2q)^{2d_i}$, $|\mu_{e_i}^{\dagger}| \leq \tau^{\dagger}$, $|\hat{\mu}_{e_i}^{\dagger}| \leq \tau^{\dagger} + \epsilon_{\dagger}$ hold under $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$, $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$. Further, under event $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ (Lemma 25) it is true that $|\hat{\mu}_{\mathcal{F}_i}^{\dagger} - \mu_{\mathcal{F}_i}^{\dagger}| \leq \gamma_{\dagger}$, the latter give (233) and (234). The bound $\tau^{\dagger} \leq 4\epsilon_{\dagger} e^{\beta} (1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta})$ gives (235) (see inequality 137). Inequality (236) requires

$$\max \left\{ \frac{\epsilon_{\dagger}}{(1-2q)^2}, \gamma_{\dagger} \right\} \leq \frac{\eta}{3} \quad (238)$$

and

$$\epsilon_{\dagger} \leq (1-2q)^2 e^{-\beta} \left[20 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right) \right]^{-1}. \quad (239)$$

Finally (237) holds for all $t \in \mathbb{N}$. The latter completes the proof. \blacksquare

The next Lemma provides the set of values of ϵ_{\dagger} that guarantee $\mathcal{L}^{(2)}(\mathbf{p}(\cdot), \Pi_{\text{TCL}}(\mathbf{p}(\cdot))) \leq \frac{\eta}{2}$ with high probability.

Lemma 27 *If*

$$\epsilon_{\dagger} \leq \min \left\{ \frac{\eta}{16} (1-2q)^2, \frac{(1-2q)^2 e^{-\beta}}{24 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right)} \right\}. \quad (240)$$

then $\mathcal{L}^{(2)}\left(\mathbf{p}(\cdot), \Pi_{\mathbf{T}^{\text{CL}}}(\mathbf{p}(\cdot))\right) \leq \frac{\eta}{2}$ under the event $\mathbf{E}_{\dagger}^{\text{corr}}(\epsilon_{\dagger}) \cap \mathbf{E}_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$.

Proof Recall that

$$\begin{aligned} \mathcal{L}^{(2)}\left(\mathbf{p}(\cdot), \Pi_{\mathbf{T}^{\text{CL}}}(\mathbf{p}(\cdot))\right) &= \frac{1}{2} \left| \prod_{e \in \text{path}_{\mathbf{T}}(w, \tilde{w})} \frac{\mu_e^{\dagger}}{(1-2q)^2} - \prod_{e \in \text{path}_{\mathbf{T}^{\text{CL}}}(w, \tilde{w})} \frac{\mu_e^{\dagger}}{(1-2q)^2} \right| \\ &= \frac{1}{2} \left| \prod_{e \in \text{path}_{\mathbf{T}}(w, \tilde{w})} \mu_e - \prod_{e \in \text{path}_{\mathbf{T}^{\text{CL}}}(w, \tilde{w})} \mu_e \right|. \end{aligned} \quad (241)$$

We follow Bresler's and Karzand's technique "Loss due to graph estimation" (Bresler and Karzand, 2020, Section 6.2) and highlight the difference that appears in our setting. For the noisy case/hidden model, the argument changes slightly in the following manner:

$$\begin{aligned} 2\mathcal{L}^{(2)}\left(\mathbf{p}(\cdot), \Pi_{\mathbf{T}^{\text{CL}}}(\mathbf{p}(\cdot))\right) &\leq |\mu_f \mu_A \mu_{\tilde{A}} \mu_B \mu_{\tilde{B}}| |\mu_C^2 \mu_{\tilde{C}}^2 - 1| + |\mu_f| \left(\Delta(k) + \Delta(\tilde{k}) + \Delta(\tilde{k}) \Delta(k) \right) \\ &= \left| \frac{\mu_f^{\dagger}}{(1-2q)^2} \mu_A \mu_{\tilde{A}} \mu_B \mu_{\tilde{B}} \right| |\mu_C^2 \mu_{\tilde{C}}^2 - 1| \\ &\quad + \left| \frac{\mu_f^{\dagger}}{(1-2q)^2} \right| \left(\Delta(k) + \Delta(\tilde{k}) + \Delta(\tilde{k}) \Delta(k) \right) \\ &\leq 8 \frac{\epsilon_{\dagger}}{(1-2q)^2} + \frac{\tau^{\dagger}}{(1-2q)^2} (2\eta + \eta^2) \end{aligned} \quad (242)$$

$$\leq \eta. \quad (243)$$

(242) holds since $|\mu_f^{\dagger}| - |\mu_g^{\dagger}| \leq 4\epsilon_{\dagger}$, $|\mu_f| \left(1 - \mu_C^2 \mu_{\tilde{C}}^2\right) \leq 2|\mu_f| - 2|\mu_g|$, $|\mu_f^{\dagger}| \leq \tau^{\dagger}$, and (243) holds for all the values of ϵ_{\dagger} that satisfy

$$\epsilon_{\dagger} \leq \min \left\{ \frac{\eta}{16} (1-2q)^2, \frac{(1-2q)^2 e^{-\beta}}{24 \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta}\right)} \right\}. \quad (244)$$

The latter provides the statement of the Lemma. \blacksquare

The next Theorem provides the sufficient number of samples for predictive learning that recovers exactly the noiseless setting for $q = 0$. Note that the dependence on β changes from $e^{2\beta}$ to $e^{4\beta}$ when the data are noisy. A key component of the bound is the following function

$$\Gamma(\beta, q) \triangleq \left(\frac{1 - (1-2q)^2}{1 - (1-2q)^4 \tanh^2(\beta)} \right)^2, \quad \beta > 0 \text{ and } q \in [0, 1/2]. \quad (245)$$

Note that $\Gamma(\beta, q) \in [0, 1]$ for all $\beta > 0$ and $q \in [0, 1/2)$, and $\Gamma(\beta, 0) = 0$ for all $\beta > 0$. Further we define

$$B(\beta, q) \triangleq \max \left\{ \frac{1}{K(\beta, q)}, \left(1 + 2e^{\beta} \sqrt{2(1-q)q \tanh \beta} \right)^2 \right\}, \quad (246)$$

and the expression of $K(\beta, q)$ is given by (168).

Theorem 28 Fix $\delta \in (0, 1)$. Choose $\eta > 0$ (independent of δ). If

$$n \geq \max \left\{ \frac{512}{\eta^2(1-2q)^4}, \frac{1152e^{2\beta}B(\beta, q)}{(1-2q)^4}, \frac{48e^{4\beta}}{\eta^2}\Gamma(\beta, q) \right\} \log \left(\frac{6p^3}{\delta} \right), \quad (247)$$

then

$$\mathbb{P} \left(\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger}) \right) \leq \eta \right) \geq 1 - \delta. \quad (248)$$

Additionally, as a consequence of (247), if

$$n \geq \max \left\{ \frac{512}{\eta^2(1-2q)^4}, \frac{1152(1+3\sqrt{q})^2 e^{2\beta(1+\mathbb{1}_{q \neq 0})}}{(1-2q)^4}, \frac{48e^{4\beta}}{\eta^2} \mathbb{1}_{q \neq 0} \right\} \log \left(\frac{6p^3}{\delta} \right), \quad (249)$$

then

$$\mathbb{P} \left(\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger}) \right) \leq \eta \right) \geq 1 - \delta. \quad (250)$$

Proof Recall that

$$\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(p(\cdot)) \right) = \frac{1}{2} \left| \prod_{e \in \text{path}_T(w, \tilde{w})} \mu_e - \prod_{e \in \text{path}_{T_{\dagger}^{CL}}(w, \tilde{w})} \mu_e \right|. \quad (251)$$

We combine the triangle inequality

$$\mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger}) \right) \leq \mathcal{L}^{(2)} \left(p(\cdot), \Pi_{T_{\dagger}^{CL}}(p(\cdot)) \right) + \mathcal{L}^{(2)} \left(\Pi_{T_{\dagger}^{CL}}(p(\cdot)), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger}) \right), \quad (252)$$

Lemma 26, and Lemma 27 to get that $\mathcal{L}^{(2)}(p(\cdot), \Pi_{T_{\dagger}^{CL}}(\hat{p}_{\dagger})) \leq \eta$ with probability at least $1 - \delta$ if

$$\gamma_{\dagger} \leq \frac{\eta}{3} \text{ and } \epsilon_{\dagger} \leq \min \left\{ \frac{\eta}{16}(1-2q)^2, \frac{(1-2q)^2 e^{-\beta}}{24(1+2e^{\beta}\sqrt{2(1-q)q \tanh \beta})} \right\}. \quad (253)$$

First, we find the necessary number of samples such that for $\gamma_{\dagger} \leq \eta/3$ the probability of the complement of $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ is not greater than $\delta/3$. Recall that

$$G \triangleq \frac{3(3e^{-1}\mathbb{1}_{q \neq 0} + 1)}{4(1-2q)^2}, \quad (254)$$

$$S \triangleq 3 - (1-2q)^2. \quad (255)$$

Recall that

$$\Gamma(\beta, q) \triangleq \left(\frac{1 - (1-2q)^2}{1 - (1-2q)^4 \tanh^2(\beta)} \right)^2, \quad \beta > 0 \text{ and } q \in [0, 1/2). \quad (256)$$

Lemma 25 gives that for any $\Delta > 0$ and $\eta > \Delta$ if

$$n \geq \max \left\{ \frac{0.3^{-2} S^2 G^2}{(\eta - \Delta)^2}, \frac{108 e^{2\beta}}{(1 - 2q)^4 K(\beta, q)}, \frac{3e^{4\beta}}{\Delta^2} \Gamma(\beta, q) \right\} \log \left(\frac{6p^3}{\delta} \right), \quad (257)$$

then the probability of the complement of $E_{\dagger}^{\text{cascade}}(\gamma_{\dagger})$ is upper bounded by $\delta/3$ and we write

$$\mathbb{P} \left((E_{\dagger}^{\text{cascade}}(\gamma_{\dagger}))^c \right) \leq \frac{\delta}{3}. \quad (258)$$

Second, we find the necessary number of samples such that the complements of the events $E_{\dagger}^{\text{strong}}(\epsilon_{\dagger})$ and $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$ occur with probability not greater than $\delta/3$ each. In fact the upper bound on ϵ_{\dagger} (253) and Lemma 18 gives that if

$$n \geq \max \left\{ \frac{512}{\eta^2(1 - 2q)^4}, \frac{1152e^{2\beta}}{(1 - 2q)^4} \left(1 + 2e^{\beta} \sqrt{2(1 - q)q \tanh \beta} \right)^2 \right\} \log \left(\frac{6p^3}{\delta} \right), \quad (259)$$

then ϵ_{\dagger} satisfies the inequality in (253) with probability at least $1 - \delta/3$. Note that (257) holds for any $\Delta \in (0, \eta)$ and we will choose $\Delta = \eta/4$. Under the choice $\Delta = \eta/4$

$$\frac{0.3^{-2} S^2 G^2}{(\eta - \Delta)^2} = \frac{0.3^{-2} S^2 G^2}{(\eta - \eta/4)^2} < \frac{512}{\eta^2}, \quad \forall \eta > 0, q \in [0, 1/2). \quad (260)$$

Recall that

$$B(\beta, q) \triangleq \max \left\{ \frac{1}{K(\beta, q)}, \left(1 + 2e^{\beta} \sqrt{2(1 - q)q \tanh \beta} \right)^2 \right\}. \quad (261)$$

Combining (257), (259), (260) and (261) yields

$$n \geq \max \left\{ \frac{512}{\eta^2(1 - 2q)^4}, \frac{1152e^{2\beta} B(\beta, q)}{(1 - 2q)^4}, \frac{48e^{4\beta}}{\eta^2} \Gamma(\beta, q) \right\} \log \left(\frac{6p^3}{\delta} \right). \quad (262)$$

The latter gives the sample complexity for accurate predictive learning, it reduces exactly to the noiseless setting of prior work by [Bresler and Karzand \(2020\)](#) and it is continuous because

$$\lim_{q \rightarrow 0^+} \Gamma(\beta, q) = \Gamma(\beta, 0) = 0, \quad \forall \beta > 0 \quad (263)$$

and

$$\lim_{q \rightarrow 0^+} K(\beta, q) = K(\beta, 0) = 1, \quad \forall \beta > 0 \quad (264)$$

$$\lim_{q \rightarrow 0^+} \left(1 + 2e^{\beta} \sqrt{2(1 - q)q \tanh \beta} \right)^2 = 1 \quad (265)$$

thus

$$\lim_{q \rightarrow 0^+} B(\beta, q) = \Gamma(\beta, 0) = 1, \quad \forall \beta > 0. \quad (266)$$

To derive a simplified version of (262) note that

$$\frac{1}{K(\beta, q)} \leq e^{2\beta \mathbb{1}_{q \neq 0}} \quad (267)$$

by the definition (168) of $K(\beta, q)$ and

$$\begin{aligned} \left(1 + 2e^\beta \sqrt{2(1-q)q \tanh \beta}\right)^2 &\leq \left(e^{\beta \mathbb{1}_{q \neq 0}} + 2e^{\beta \mathbb{1}_{q \neq 0}} \sqrt{2(1-q)q \tanh \beta}\right)^2 \\ &\leq (1 + 3\sqrt{q})^2 e^{2\beta \mathbb{1}_{q \neq 0}}. \end{aligned} \quad (268)$$

Then (261), (267) and (268) give

$$B(\beta, q) \leq (1 + 3\sqrt{q})^2 e^{2\beta \mathbb{1}_{q \neq 0}} \quad (269)$$

and by the definition (256) $\Gamma(\beta, q) \in [0, 1)$ and $\Gamma(\beta, 0) = 0$ for all $\beta > 0$, thus

$$\Gamma(\beta, q) \leq \mathbb{1}_{q \neq 0}. \quad (270)$$

Finally, we combine (262), (269), (270) to get

$$n \geq \max \left\{ \frac{512}{\eta^2(1-2q)^4}, \frac{1152(1+3\sqrt{q})^2 e^{2\beta(1+\mathbb{1}_{q \neq 0})}}{(1-2q)^4}, \frac{48e^{4\beta}}{\eta^2} \mathbb{1}_{q \neq 0} \right\} \log \left(\frac{6p^3}{\delta} \right). \quad (271)$$

This completes the proof. ■

Appendix F. Theorem 11: KL-Divergence Loss

Assume the Ising model tree distributions P_θ according to a tree $T_\theta = (\mathcal{V}, \mathcal{E}_\theta)$ and the estimate $P_{\theta'}$ according a tree $T_{\theta'} = (\mathcal{V}, \mathcal{E}_{\theta'})$. The goal is to upper bound the symmetric KL divergence

$$S_{\text{KL}}(\theta || \theta') = \sum_{s, t \in \mathcal{E}} (\theta_{st} - \theta'_{st}) (\mu_{st} - \mu'_{st})$$

with high probability. Under the event $E^{\text{corr}}(\epsilon)$ we can upper bound the quantity $|\mu_{st} - \mu'_{st}|$ for all $(s, t) \in \mathcal{E}$ with high probability.

By using bounds $|\theta_{st} - \theta'_{st}| \leq 2\beta$ and $|\mu_{st} - \mu'_{st}| \leq \epsilon$ for all $(s, t) \in \mathcal{E}$ under the event $E^{\text{corr}}(\epsilon)$, we have

$$\begin{aligned} S_{\text{KL}}(\theta || \theta') &= |S_{\text{KL}}(\theta || \theta')| \\ &= \left| \sum_{s, t \in \mathcal{E}} (\theta_{st} - \theta'_{st}) (\mu_{st} - \mu'_{st}) \right| \\ &\leq \sum_{s, t \in \mathcal{E}} |\theta_{st} - \theta'_{st}| |\mu_{st} - \mu'_{st}| \end{aligned}$$

$$\begin{aligned} &\leq (p-1) |\beta - (-\beta)| \epsilon \\ &\leq \eta_S, \end{aligned} \tag{272}$$

by assuming $\epsilon \leq \frac{\eta_S}{2\beta(p-1)}$. The sufficient number of samples satisfies the inequality

$$n \geq 4 \log(p^2/\delta) \frac{\beta^2(p-1)^2}{\eta_S^2}. \tag{273}$$

Now assume that n_{\dagger} samples of \mathbf{Y} are given, by using the estimate $P_{\theta'} = \Pi_{\text{TCL}}(\hat{p}_{\dagger})$ defined in (227) under the event $E_{\dagger}^{\text{corr}}(\epsilon_{\dagger})$ we have $\left| \mu_{st} - \frac{\hat{\mu}_{st}^{\dagger}}{(1-2q)^2} \right| \leq \frac{\epsilon_{\dagger}}{(1-2q)^2}$ from Lemma 18. In the same way by assuming $\epsilon_{\dagger} \leq \frac{\eta_S(1-2q)^2}{2\beta(p-1)}$, we get

$$n_{\dagger} \geq 4 \log(p^2/\delta) \frac{\beta^2(p-1)^2}{(1-2q)^4 \eta_S^2}. \tag{274}$$

Appendix G. Theorem 6 and Theorem 8: Proofs

We combine Fano's inequality and a Strong Data Processing Inequality to prove the necessary number of samples in the hidden model setting, first for structure learning (Theorem 6) and then for inference (Theorem 8). We use the following variation of Fano's inequality.

Corollary 29 (*Tsybakov, 2009, Corollary 2.6*): Assume that Θ is a family of $M+1$ distributions $\theta_0, \theta_1, \dots, \theta_M$ such that $M \geq 2$. Let P_{θ_i} be the distribution of the variable X under the model θ_i , if

$$\frac{1}{M+1} \sum_{i=1}^M D_{KL}(P_{\theta_i} || P_{\theta_0}) \leq \gamma \log M, \quad \text{for } \gamma \in (0, 1) \tag{275}$$

then for the probability of error p_e the following inequality holds: $p_e \geq \frac{\log(M+1) - \log(2)}{\log(M)} - \gamma$.

The construction from the noiseless case, with Corollary 29 and the Strong Data Processing Inequality for the BSC yield the bound of Theorem 6. We start by presenting Bresler's and Karzand's construction, which gives a sufficiently tight upper bound on symmetric KL divergence.

Proof of Theorem 6: Consider a family of $M+1$ different Ising model distributions $\{P_{\theta^i} : i \in \{0, \dots, M\}\}$. This family of the structured distributions is chosen such that the structure recovery task (through The Chow-Liu algorithm) is sufficiently hard. First, we define P_{θ^0} to be an Ising model distribution with underlying structure a chain with p nodes and parameters $\theta_{j,j+1}^0 = \alpha$, when j is odd and $\theta_{j,j+1}^0 = \beta$ when j is even. The rest of family is constructed as follows: the elements of each θ^i , $i \in [M]$ are equal to the elements of θ^0 apart from two elements $\theta_{i,i+1}^i = 0$ and $\theta_{i,i+2}^i = \text{arctanh}(\tanh(\alpha) \tanh(\beta))$, for each odd value of j . There are $(p+1)/2$ distinct distributions in the constructed family. Through the expression (49), we derive the following upper bound on the $S_{KL}(P_{\theta^0} || P_{\theta^i})$, for all $i \in [M]$, (*Bresler and Karzand, 2020, Section 7.1*),

$$S_{KL}(P_{\theta^0} || P_{\theta^i}) = \alpha (\tanh(\alpha) - \tanh(\alpha) \tanh^2(\beta)) \leq 4\alpha \tanh(\alpha) e^{-2\beta}. \tag{276}$$

Strong Data Processing Inequality: For each distribution P_{θ^i} and $i \in \{0, \dots, M\}$ we consider the distribution of the noisy variable in the hidden model $P_{\theta^i}^\dagger \triangleq P_{\mathbf{Y}|\mathbf{X}} \circ P_{\theta^i}$. We would like to find an upper bound for the quantities $\mathcal{S}_{\text{KL}}(P_{\theta^0}^\dagger || P_{\theta^i}^\dagger)$, $i \in \{0, \dots, M\}$. For that purpose, we use a strong data processing inequality result for the BSC by [Polyanskiy and Wu \(2017\)](#). The input random variable \mathbf{X} is considered to have correlated binary elements, while the noise variables N_i are i.i.d Rademacher(q). This scheme is equivalent to the hidden model that we consider in this paper. In fact we have the following bound

$$\eta_{\text{KL}} \leq 1 - (4q(1-q))^p, \quad (277)$$

that is proved by Polyanski ([Polyanskiy and Wu, 2017](#), ‘‘Evaluation for the BSC’’, equation (39)), where the quantity η_{KL} is defined as

$$\eta_{\text{KL}} \triangleq \sup_Q \sup_{P: 0 < \mathcal{D}_{\text{KL}}(P||Q) < \infty} \frac{\mathcal{D}_{\text{KL}}(P_{\mathbf{Y}|\mathbf{X}} \circ P || P_{\mathbf{Y}|\mathbf{X}} \circ Q)}{\mathcal{D}_{\text{KL}}(P||Q)}, \quad (278)$$

$P_{\mathbf{Y}|\mathbf{X}}$ is the distribution of the BSC and P, Q are any distributions of the input variable \mathbf{X} .

Since the supremum in (278) is with respect to all possible distributions, it covers any pair of distributions in the desired family $\{P_{\theta^j} : j \in \{0, \dots, M\}\}$. Thus, for all $k, \ell \in \{0, 1, \dots, M\}$ and $k \neq \ell$, it is true that

$$\frac{\mathcal{D}_{\text{KL}}(P_{\theta^k}^\dagger || P_{\theta^\ell}^\dagger)}{\mathcal{D}_{\text{KL}}(P_{\theta^k} || P_{\theta^\ell})} \leq 1 - (4q(1-q))^p, \quad (279)$$

which comes from (277),(278) and implies the following

$$\mathcal{S}_{\text{KL}}(P_{\theta^k}^\dagger || P_{\theta^\ell}^\dagger) \leq [1 - (4q(1-q))^p] \mathcal{S}_{\text{KL}}(P_{\theta^k} || P_{\theta^\ell}), \quad \forall k \neq \ell \in \{0, 1, \dots, M\}. \quad (280)$$

We combine (276) and (280) to get

$$\mathcal{S}_{\text{KL}}(P_{\theta^k}^\dagger || P_{\theta^i}^\dagger) \leq [1 - (4q(1-q))^p] 4\alpha \tanh(\alpha) e^{-2\beta} \leq [1 - (4q(1-q))^p] 4\alpha^2 e^{-2\beta}. \quad (281)$$

Finally, from (281) and Corollary 29 we derive the first part of Theorem 6.

Proof of Theorem 8: Theorem 8 is the extended version of Theorem 3.4 by [Bresler and Karzand \(2020\)](#) to the hidden model. Following a similar technique, we consider chain structured Ising models with parameters θ^j for $j \in [M]$ such that $\theta_{j,j+1}^j = \alpha$ and $\theta_{i,i+1}^j = \text{arctanh}(\tanh(\alpha) + 2\eta)$, for all $i \neq j$. Then

$$\mathcal{L}^{(2)}(P_{\theta^j}, P_{\theta^{j'}}) = \max_{s,t} |\mathbb{E}_{\theta^j}[X_s X_t] - \mathbb{E}_{\theta^{j'}}[X_s X_t]| \geq 2\eta \quad (282)$$

and

$$\mathcal{S}_{\text{KL}}(P_{\theta^j}, P_{\theta^{j'}}) \leq 2\eta [\text{arctanh}(\tanh(\alpha) + 2\eta) - \alpha] \leq 2\eta \frac{2\eta}{1 - [\tanh(\alpha) + 2\eta]^2}, \quad (283)$$

where the last inequality is a consequence of Mean Value Theorem (see ([Bresler and Karzand, 2020](#), Section 6.3) for the original statement). We derive the bound of Theorem 8 by combining the strong data processing inequality (278) with (277), (283), and Corollary 29.

Appendix H. Supplementary Discussion

In this section we provide supplementary material that supports the discussion in Sections 2.4 and 2.6. First, we present one marginal case for which perfect denoising is possible before applying the Chow-Liu algorithm. Then we show a structure-preserving case.

H.1 The Gap between the Upper and Lower Bounds

We continue by analyzing the gap that appears between the upper and lower bounds for an example where perfect denoising can be applied on a specific class of tree models in the high-dimensional regime. This shows why the effect of noise vanishes in Theorems 1.2 and 1.4 for $p \rightarrow \infty$. Further, while it seems counter-intuitive that when $p \rightarrow \infty$ the problem becomes easier, we show below one example that this is the case. Our lower bound is directly affected by marginal cases like this, for instance see Proposition 1.4.

The gap is introduced by the terms $(1 - 2q)^4$ and $1 - (4q(1 - 2q))^p$ in the denominator of the lower and upper bounds respectively. Specifically, for $p \rightarrow \infty$ there exists a special case for which perfect denoising before running the Chow-Liu algorithm is possible, while in other cases that is not possible. Thus the minimax bound ought to be identical to noiseless case when $p \rightarrow \infty$ and $1 - (4q(1 - 2q))^p \rightarrow 1$ in the large dimensional regime. We continue by providing the marginal case of a trivial tree structure and showing that perfect denoising is possible in this case before running the Chow-Liu algorithm.

First notice that if $p \rightarrow \infty$, then the sample size $n \rightarrow \infty$, even in the noiseless regime. Consider the case of $\mathbb{E}[X_i X_j] \rightarrow 1$ for all $(i, j) \in \mathcal{E}$. Because an infinite number of samples are available, we can estimate perfectly the correlations of the observables and we find $\hat{\mathbb{E}}[Y_i Y_j] = \mathbb{E}[Y_i Y_j] = (1 - 2q)^2$ for all $(i, j) \in \mathcal{V}$. The latter as information is sufficient to find that $\mathbb{E}[X_i X_j] \rightarrow 1$ for all $(i, j) \in \mathcal{E}$. The hidden layer \mathbf{X} take two values, $(+1, +1, \dots) \triangleq +1^p$ (p values $+1$) or $(-1, -1, \dots) \triangleq -1^p$ (p values -1), because $\mathbb{E}[X_i X_j] \rightarrow 1$ for all $(i, j) \in \mathcal{V}$ and the later allows us to denoise each sample. Define as $d_H(\mathbf{X}, \mathbf{Y})$ the Hamming distance between \mathbf{X} and \mathbf{Y} . At this point we can perform perfect denoising for each sample \mathbf{y}_s of infinite length p and find the hidden sample \mathbf{x}_s with probability 1 because

$$\begin{aligned} \mathbb{P}(\mathbf{X} = \mathbf{x}_s | \mathbf{Y} = \mathbf{y}_s) &= \frac{\mathbb{P}(\mathbf{Y} = \mathbf{y}_s | \mathbf{X}_s = \mathbf{x}_s) \mathbb{P}(\mathbf{X}_s = \mathbf{x}_s)}{\sum_{\mathbf{x}} \mathbb{P}(\mathbf{Y} = \mathbf{y}_s | \mathbf{X} = \mathbf{x}) \mathbb{P}(\mathbf{X} = \mathbf{x})} \\ &= \frac{q^{d_H(\mathbf{x}_s, \mathbf{y}_s)} (1 - q)^{p - d_H(\mathbf{x}_s, \mathbf{y}_s)}}{q^{d_H(\mathbf{x}_s, \mathbf{y}_s)} (1 - q)^{p - d_H(\mathbf{x}_s, \mathbf{y}_s)} + q^{d_H(-\mathbf{x}_s, \mathbf{y}_s)} (1 - q)^{p - d_H(-\mathbf{x}_s, \mathbf{y}_s)}} \end{aligned} \quad (284)$$

and the last holds for both of the cases $\mathbf{x}_s = +1^p$ or $\mathbf{x}_s = -1^p$ because of symmetry. Further for any observation \mathbf{y}_s for any $q \in (0, 1/2)$ we have

$$\lim_{p \rightarrow \infty} \frac{d_H(-\mathbf{x}_s, \mathbf{y}_s) - d_H(\mathbf{x}_s, \mathbf{y}_s)}{p} \stackrel{a.s.}{=} 1 - q - q = 1 - 2q. \quad (285)$$

We combine (284) and (285) to find

$$\lim_{p \rightarrow \infty} \mathbb{P}(\mathbf{X} = \mathbf{x}_s | \mathbf{Y} = \mathbf{y}_s) = \lim_{p \rightarrow \infty} \frac{1}{1 + \left(\frac{q}{1-q}\right)^{\frac{d_H(-\mathbf{x}_s, \mathbf{y}_s) - d_H(\mathbf{x}_s, \mathbf{y}_s)}{p} p}} = 1 \quad (286)$$

and

$$\lim_{p \rightarrow \infty} \mathbb{P}(\mathbf{X} = -\mathbf{x}_s | \mathbf{Y} = \mathbf{y}_s) = \lim_{p \rightarrow \infty} \frac{1}{1 + \left(\frac{1-q}{q}\right)^{\frac{d_H(-\mathbf{x}_s, \mathbf{y}_s) - d_H(\mathbf{x}_s, \mathbf{y}_s)}{p}}} = 0. \quad (287)$$

As a consequence there exists one case for which perfect denoising is possible before running the Chow-Liu algorithm. Because we want Theorems 1.2 and 1.4 to reduce to the noiseless case for $p \rightarrow \infty$, the above best case scenario must be covered. However, perfect denoising is not possible in general (for instance $\mathbb{E}[X_i X_j] < 1$ and finite p).

H.2 A Structure-Preserving Case

Lemma 30 considers a special case of tree structures for the hidden variables, the set of edges is a set with disconnected edges, no edge is connected to any other. Then we show that the same structure is preserved for the observable variables.

Lemma 30 *Let $F = (\mathcal{V}, \mathcal{E})$ be a forest with $|\mathcal{V}| = p$ and $|\mathcal{E}| = p/2 \in \mathbb{N}$ such that no edge is connected to any other edge. Assume that $X_i \in \{-1, +1\}$ and $\mathbb{E}[X_i] = 0$ for all $i \in [1, \dots, p]$. If \mathbf{Y} is the output of the BSC channel (in the hidden model) with distribution $p_{\dagger}(\mathbf{y})$, then $p_{\dagger}(\mathbf{y})$ also factorizes with respect to F .*

Proof The pair variables (Y_i, Y_j) for $(i, j) \in \mathcal{E}$ are independent because of the disconnected edges of the hidden layer. The latter directly gives the factorization as

$$p_{\dagger}(\mathbf{y}) = \prod_{(i,j) \in \mathcal{E}} p_{\dagger}(y_i, y_j) = \prod_{i \in \mathcal{V}} p_{\dagger}(y_i) \prod_{(i,j) \in \mathcal{E}} \frac{p_{\dagger}(y_i, y_j)}{p_{\dagger}(y_i)p_{\dagger}(y_j)}, \quad (288)$$

because $|\mathcal{V}| = p$, $|\mathcal{E}| = p/2$ and the marginal distributions are uniform. ■

Acknowledgments

This work was supported in part by DARPA and SSC Pacific under contract N66001-15-C-4070 and the United States National Science Foundation under award CCF-1453432, and the United States National Institutes of Health under award 1R01DA040487.

References

- Animashree Anandkumar and Ragupathyraj Valluvan. Learning loopy graphical models with latent variables: Efficient methods and guarantees. *The Annals of Statistics*, pages 401–435, 2013. URL <https://www.jstor.org/stable/23566568>.
- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *The Journal of Machine Learning Research*, 15(1):2773–2832, 2014. URL <http://jmlr.org/papers/v15/anandkumar14b.html>.

- M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002. URL <https://www.irisa.fr/aspi/legland/ensta/ref/arulampalam02a.pdf>.
- Onureena Banerjee, Laurent El Ghaoui, and Alexandre d Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *Journal of Machine learning research*, 9(3):485–516, 2008. URL <http://www.jmlr.org/papers/volume9/banerjee08a/banerjee08a.pdf>.
- Michael Banf and Seung Y Rhee. Enhancing gene regulatory network inference through data integration with Markov random fields. *Scientific Reports*, 7, 2017. URL <https://www.nature.com/articles/srep41174>.
- Rina Foygel Barber, Mathias Drton, et al. High-dimensional Ising model selection with bayesian information criteria. *Electronic Journal of Statistics*, 9(1):567–607, 2015. URL <https://projecteuclid.org/euclid.ejs/1427203129>.
- Lionel Barnett, Joseph T Lizier, Michael Harré, Anil K Seth, and Terry Bossomaier. Information flow in a kinetic Ising model peaks in the disordered phase. *Physical Review Letters*, 111(17):177203, 2013. URL <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.111.177203>.
- George Bennett. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962. URL https://www.jstor.org/stable/2282438?seq=1#page_scan_tab_contents.
- Martin Bilodeau. Graphical lassos for meta-elliptical distributions. *Canadian Journal of Statistics*, 42(2):185–203, 2014. URL https://www.jstor.org/stable/43185177?seq=1#page_scan_tab_contents.
- Guy Bresler. Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing*, pages 771–782. ACM, 2015. URL <https://arxiv.org/abs/1411.6156>.
- Guy Bresler and Mina Karzand. Learning a tree-structured ising model in order to make predictions. *Annals of Statistics*, 48(2):713–737, 2020. URL <https://projecteuclid.org/euclid.aos/1590480031>.
- Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591, 2009. URL <https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.81.591>.
- Arun Tejasvi Chaganty and Percy Liang. Estimating latent-variable graphical models using moments and likelihoods. In *International Conference on Machine Learning*, pages 1872–1880, 2014. URL <http://proceedings.mlr.press/v32/chaganty14.html>.
- Venkat Chandrasekaran, Pablo A Parrilo, and Alan S Willsky. Latent variable graphical model selection via convex optimization. In *Communication, Control, and Computing*

- (Allerton), *2010 48th Annual Allerton Conference on*, pages 1610–1613. IEEE, 2010. URL <https://ieeexplore.ieee.org/document/5707106/>.
- C Chow and Cong Liu. Approximating discrete probability distributions with dependence trees. *IEEE transactions on Information Theory*, 14(3):462–467, 1968. URL <https://ieeexplore.ieee.org/abstract/document/1054142>.
- Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Optimal phylogenetic reconstruction. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, pages 159–168. ACM, 2006. URL <https://dl.acm.org/doi/abs/10.1145/1132516.1132540>.
- Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Phylogenies without branch bounds: Contracting the short, pruning the deep. In *Annual International Conference on Research in Computational Molecular Biology*, pages 451–465. Springer, 2009. URL https://link.springer.com/chapter/10.1007/978-3-642-02008-7_32.
- Constantinos Daskalakis, Sebastien Roch, et al. Alignment-free phylogenetic reconstruction: Sample complexity via a branching process analysis. *The Annals of Applied Probability*, 23(2):693–721, 2013. URL <https://projecteuclid.org/euclid.aoap/1360682027>.
- Constantinos Daskalakis, Nishanth Dikkala, and Gautam Kamath. Testing Ising models. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1989–2007. SIAM, 2018. URL <https://dl.acm.org/doi/10.5555/3174304.3175435>.
- Randal Douc, Eric Moulines, Jimmy Olsson, Ramon Van Handel, et al. Consistency of the maximum likelihood estimator for general hidden Markov models. *the Annals of Statistics*, 39(1):474–513, 2011. URL <https://projecteuclid.org/euclid.aos/1297779854>.
- Mathias Drton and Marloes H Maathuis. Structure learning in graphical modeling. *Annual Review of Statistics and Its Application*, 4:365–393, 2017. URL <https://www.annualreviews.org/doi/abs/10.1146/annurev-statistics-060116-053803?journalCode=statistics>.
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006a. URL <https://www.iacr.org/archive/eurocrypt2006/40040493/40040493.pdf>.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006b. URL https://link.springer.com/chapter/10.1007/11681878_14.

- David Edwards, Gabriel CG De Abreu, and Rodrigo Labouriau. Selecting high-dimensional mixed graphical models using minimal AIC or BIC forests. *BMC Bioinformatics*, 11(1):18, 2010. URL <https://bmcbioinformatics.biomedcentral.com/articles/10.1186/1471-2105-11-18>.
- Péter L Erdős, Michael A Steel, László A Székely, and Tandy J Warnow. A few logs suffice to build (almost) all trees: Part ii. *Theoretical Computer Science*, 221(1-2):77–118, 1999.
- E Yagmur Erten, Joseph T Lizier, Mahendra Piraveenan, and Mikhail Prokopenko. Criticality and information dynamics in epidemiological models. *Entropy*, 19(5):194, 2017. URL <https://www.mdpi.com/1099-4300/19/5/194>.
- Jianqing Fan, Han Liu, Yang Ning, and Hui Zou. High dimensional semiparametric latent graphical model for mixed data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(2):405–421, 2017. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssb.12168>.
- Xiequan Fan, Ion Grama, and Quansheng Liu. Hoeffding’s inequality for supermartingales. *Stochastic Processes and their Applications*, 122(10):3545–3559, 2012. URL <https://arxiv.org/abs/1109.4359>.
- Michael Finegold and Mathias Drton. Robust graphical modeling of gene networks using classical and alternative t-distributions. *The Annals of Applied Statistics*, pages 1057–1080, 2011. URL <https://projecteuclid.org/euclid.aoas/1310562216>.
- Rina Foygel and Mathias Drton. Extended Bayesian information criteria for Gaussian graphical models. In *Advances in Neural Information Processing Systems*, pages 604–612, 2010. URL <https://arxiv.org/abs/1011.6640>.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, 2008. URL <https://www.ncbi.nlm.nih.gov/pubmed/18079126>.
- Xin Gao, Daniel Q Pu, Yuehua Wu, and Hong Xu. Tuning parameter selection for penalized likelihood estimation of Gaussian graphical model. *Statistica Sinica*, pages 1123–1146, 2012. URL <https://www.jstor.org/stable/24309977>.
- Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*, volume 9. Walter de Gruyter, 2011.
- Surbhi Goel, Daniel M Kane, and Adam R Klivans. Learning ising models with independent failures. In *Conference on Learning Theory*, pages 1449–1469, 2019. URL <http://proceedings.mlr.press/v99/goel19a.html>.
- Søren Højsgaard, David Edwards, and Steffen Lauritzen. *Graphical Models with R*. Springer Science & Business Media, 2012.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(1):695–709, 2005. URL jmlr.org/papers/volume6/hyvarinen05a/old.pdf.

- Aapo Hyvärinen. Some extensions of score matching. *Computational statistics & data analysis*, 51(5):2499–2512, 2007. URL <https://www.cs.helsinki.fi/u/ahyvarin/papers/CSDA07.pdf>.
- Leon Isserlis. On a formula for the product-moment coefficient of any order of a normal frequency distribution in any number of variables. *Biometrika*, 12(1/2):134–139, 1918. URL https://www.jstor.org/stable/2331932?seq=1#page_scan_tab_contents.
- Ali Jalali, Christopher C Johnson, and Pradeep K Ravikumar. On learning discrete graphical models using greedy methods. In *Advances in Neural Information Processing Systems*, pages 1935–1943, 2011. URL <https://papers.nips.cc/paper/4290-on-learning-discrete-graphical-models-using-greedy-methods>.
- Andrew H Jazwinski. *Stochastic Processes and Filtering Theory*. Courier Corporation, 2007.
- Dionysios S Kalogerias and Athina P Petropulu. Grid based nonlinear filtering revisited: Recursive estimation & asymptotic optimality. *IEEE Transactions on Signal Processing*, 64(16):4244–4259, 2016. URL <https://ieeexplore.ieee.org/document/7457725>.
- David Karger and Nathan Srebro. Learning Markov networks: Maximum bounded tree-width graphs. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 392–401. Society for Industrial and Applied Mathematics, 2001. URL <https://www.biostat.wisc.edu/~page/markovnet.pdf>.
- S. A. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. What can we learn privately? In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 531–540, 2008. doi: 10.1109/FOCS.2008.27. URL <http://dx.doi.org/10.1109/FOCS.2008.27>.
- Minje Kim and Paris Smaragdis. Single channel source separation using smooth nonnegative matrix factorization with Markov random fields. In *Machine Learning for Signal Processing (MLSP), 2013 IEEE International Workshop on*, pages 1–6. IEEE, 2013. URL <https://ieeexplore.ieee.org/document/6661909/>.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.
- Steffen L Lauritzen. *Graphical Models*, volume 17. The Clarendon Press Oxford University Press, New York, 1996.
- Chuan Li and Michael Wand. Combining Markov random fields and convolutional neural networks for image synthesis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. URL <https://ieeexplore.ieee.org/document/7780641>.
- Guosheng Lin, Chunhua Shen, Anton van den Hengel, and Ian Reid. Efficient piecewise training of deep structured models for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3194–3203, 2016a. URL http://openaccess.thecvf.com/content_cvpr_2016/papers/Lin_Efficient_Piecewise_Training_CVPR_2016_paper.pdf.

- Lina Lin, Mathias Drton, Ali Shojaie, et al. Estimation of high-dimensional graphical models using regularized score matching. *Electronic Journal of Statistics*, 10(1):806–854, 2016b. URL <https://projecteuclid.org/euclid.ejs/1459967424>.
- Han Liu, Min Xu, Haijie Gu, Anupam Gupta, John Lafferty, and Larry Wasserman. Forest density estimation. *Journal of Machine Learning Research*, 12(Mar):907–951, 2011. URL <http://jmlr.csail.mit.edu/papers/volume12/liu11a/liu11a.pdf>.
- Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Deep learning Markov random field for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. URL <https://ieeexplore.ieee.org/document/8006236>.
- Shiqian Ma, Lingzhou Xue, and Hui Zou. Alternating direction methods for latent variable Gaussian graphical model selection. *Neural computation*, 25(8):2172–2198, 2013. URL https://www.mitpressjournals.org/doi/abs/10.1162/NECO_a_00379?journalCode=neco.
- Hirotsugu Matsuda, Naofumi Ogita, Akira Sasaki, and Kazunori Satō. Statistical mechanics of population: The lattice Lotka-Volterra model. *Progress of Theoretical Physics*, 88(6): 1035–1049, 1992. URL <http://dx.doi.org/10.1143/ptp/88.6.1035>.
- Alan Morningstar and Roger G. Melko. Deep learning the ising model near criticality. *Journal of Machine Learning Research*, 18(163):1–17, 2018. URL <http://jmlr.org/papers/v18/17-527.html>.
- Preetam Nandy, Alain Hauser, and Marloes H Maathuis. High-dimensional consistency in score-based and hybrid structure learning. *arXiv preprint arXiv:1507.02608*, 2015. URL <https://projecteuclid.org/euclid.aos/1536307246>.
- Konstantinos E Nikolakakis, Dionysios S Kalogierias, and Anand D Sarwate. Information thresholds for non-parametric structure learning on tree graphical models. *arXiv preprint arXiv:1909.09596*, 2020. URL <https://arxiv.org/abs/1909.09596>.
- Judea Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference (Morgan Kaufmann Series in Representation and Reasoning), 1988.
- Yury Polyanskiy and Yihong Wu. Strong data-processing inequalities for channels and Bayesian networks. In *Convexity and Concentration*, pages 211–249. Springer, 2017. URL <http://people.lids.mit.edu/yp/homepage/data/simple-IMA.pdf>.
- Maxim Raginsky. Strong data processing inequalities and ϕ -Sobolev inequalities for discrete channels. *IEEE Transactions on Information Theory*, 62(6):3355–3389, 2016. URL <https://arxiv.org/pdf/1411.3575.pdf>.
- Pradeep Ravikumar, Martin J Wainwright, John D Lafferty, et al. High-dimensional Ising model selection using L1-regularized logistic regression. *The Annals of Statistics*, 38(3): 1287–1319, 2010. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.67.6940&rep=rep1&type=pdf>.

- Avik Ray, Sujay Sanghavi, and Sanjay Shakkottai. Improved greedy algorithms for learning graphical models. *IEEE Transactions on Information Theory*, 61(6):3457–3468, 2015. URL <https://ieeexplore.ieee.org/document/7097023/>.
- Patrick Rebeschini, Ramon Van Handel, et al. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, 25(5):2809–2866, 2015. URL <https://projecteuclid.org/euclid.aoap/1438261054#info/10.1214/14-AAP1061>.
- Alexander G Schwing and Raquel Urtasun. Fully connected deep structured networks. *arXiv preprint arXiv:1503.02351*, 2015. URL <https://arxiv.org/abs/1503.02351>.
- Tetsuya Takaishi. Multiple time series Ising model for financial market simulations. In *Journal of Physics: Conference Series*, volume 574, page 012149. IOP Publishing, 2015. URL <https://iopscience.iop.org/article/10.1088/1742-6596/574/1/012149/pdf>.
- Vincent YF Tan, Animashree Anandkumar, and Alan S Willsky. Learning Gaussian tree models: Analysis of error exponents and extremal structures. *IEEE Transactions on Signal Processing*, 58(5):2701–2714, 2010. URL <https://ieeexplore.ieee.org/document/5406101>.
- Vincent YF Tan, Animashree Anandkumar, and Alan S Willsky. Learning high-dimensional Markov forest distributions: Analysis of error rates. *Journal of Machine Learning Research*, 12(May):1617–1653, 2011. URL <https://dl.acm.org/citation.cfm?id=2021052>.
- Salvatore Torquato. Toward an Ising model of cancer and beyond. *Physical Biology*, 8(1):015017, 2011. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3151151/>.
- Alexandre B Tsybakov. Introduction to Non-parametric Estimation. Revised and Extended from the 2004 French Original. Translated by Vladimir Zaiats, 2009.
- Ramon Van Handel. Observability and nonlinear filtering. *Probability theory and related fields*, 145(1-2):35–74, 2009. URL <https://link.springer.com/article/10.1007/s00440-008-0161-y>.
- Daniel Vogel and Roland Fried. Elliptical graphical modelling. *Biometrika*, 98(4):935–951, 2011. URL https://www.jstor.org/stable/pdf/23076182.pdf?seq=1#page_scan_tab_contents.
- Daniel Vogel and David E Tyler. Robust estimators for non-decomposable elliptical graphical models. *Biometrika*, 101(4):865–882, 2014. URL <https://arxiv.org/abs/1302.5251>.
- Martin J Wainwright, Tommi S Jaakkola, and Alan S Willsky. Tree-reweighted belief propagation algorithms and approximate ML estimation by pseudo-moment matching. In *AISTATS*, 2003. URL https://www.jstor.org/stable/23566568?seq=1#page_scan_tab_contents.
- Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008. URL <https://www.nowpublishers.com/article/DownloadSummary/MAL-001>.

- Bin Wang, Zhijian Ou, and Zhiqiang Tan. Learning trans-dimensional random fields with applications to language modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. URL <https://ieeexplore.ieee.org/document/7906616>.
- Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, March 1965. doi: 10.2307/2283137. URL <http://dx.doi.org/10.2307/2283137>.
- Scott Wisdom, John Hershey, Jonathan Le Roux, and Shinji Watanabe. Deep unfolding for multichannel source separation. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 121–125. IEEE, 2016. URL <https://ieeexplore.ieee.org/document/7471649>.
- Baoyuan Wu, Bao-Gang Hu, and Qiang Ji. A coupled hidden Markov random field model for simultaneous face clustering and tracking in videos. *Pattern Recognition*, 64:361–373, 2017. URL <https://www.sciencedirect.com/science/article/abs/pii/S0031320316303387>.
- Ming Yuan and Yi Lin. Model selection and estimation in the Gaussian graphical model. *Biometrika*, 94(1):19–35, 2007. URL https://www.jstor.org/stable/20441351?seq=1#page_scan_tab_contents.
- W-X Zhou and Didier Sornette. Self-organizing Ising model of financial markets. *The European Physical Journal B*, 55(2):175–181, 2007. URL <https://epjb.epj.org/articles/epjb/abs/2007/02/b06183/b06183.html>.
- Abdelhak M Zoubir, Visa Koivunen, Yacine Chakhchoukh, and Michael Muma. Robust estimation in signal processing: A tutorial-style treatment of fundamental concepts. *IEEE Signal Processing Magazine*, 29(4):61–80, 2012. URL https://ieeexplore.ieee.org/abstract/document/6217389?casa_token=a-jZiHR6I9gAAAAA:EwBF9R3XSyWhyLB8Xab0FUDihrBT_1saGuspv5zAI8JZ_OWpvCrrwo_ByYH9WFtv0Sz7LC4Ww.
- Yiming Zuo, Yi Cui, Guoqiang Yu, Ruijiang Li, and Habtom W Ressim. Incorporating prior biological knowledge for network-based differential gene expression analysis using differentially weighted graphical LASSO. *BMC Bioinformatics*, 18(1):99, 2017. URL <https://bmcbioinformatics.biomedcentral.com/articles/10.1186/s12859-017-1515-1>.