
Entropic Causal Inference: Graph Identifiability

Spencer Compton^{1,2} Kristjan Greenewald^{2,3} Dmitriy Katz^{2,3} Murat Kocaoglu⁴

Abstract

Entropic causal inference is a recent framework for learning the causal graph between two variables from observational data by finding the information-theoretically simplest structural explanation of the data, i.e., the model with smallest entropy. In our work, we first extend the causal graph identifiability result in the two-variable setting under relaxed assumptions. We then show the first identifiability result using the entropic approach for learning causal graphs with more than two nodes. Our approach utilizes the property that ancestry between a source node and its descendants can be determined using the bivariate entropic tests. We provide a sound sequential peeling algorithm for general graphs that relies on this property. We also propose a heuristic algorithm for small graphs that shows strong empirical performance. We rigorously evaluate the performance of our algorithms on synthetic data generated from a variety of models, observing improvement over prior work. Finally we test our algorithms on real-world datasets.

1. Introduction

Causal reasoning is essential for high-quality decision-making, as, for instance, it improves interpretability and enables counterfactual reasoning (Athey, 2015; Morgan & Winship, 2015; Moraffah et al., 2020). By learning the relationships between causes and effects, we can predict how various interventions would affect a system. Advances in causality enable us to better answer questions such as “*Why does this phenomenon occur in the system?*” or “*What could happen if the system were perturbed in this particular way?*” Moreover, causal inference methods are being utilized to tackle key challenges for reliability of ML systems,

such as domain adaptation (Magliacane et al., 2018; Zhang et al., 2020) and generalization (e.g. via causal transportability or imputation) (Bareinboim & Pearl, 2012; Pearl & Bareinboim, 2014; Squires et al., 2020).

Structural causal models (SCMs) represent relationships in a system of random variables (Pearl, 2009). In particular, each variable is modeled with a structural equation that characterizes how the variable is realized. Causal graphs are directed acyclic graphs (DAGs) that are used to represent such systems, where nodes and edges correspond to variables and the causal relations between these variables, respectively. A variable’s structural equation is a function of the variable’s corresponding node’s parents in the graph.

Learning such causal graphs can be done through a series of interventions. However, in many settings it is not possible to perform such interventions. A large amount of literature has focused on learning the causal graph from observational data with additional “faithfulness assumptions” (Spirtes et al., 2000), though in general it is impossible to fully learn the causal graph without stronger assumptions on the generative model. A variety of stronger assumptions and corresponding methodologies exist in the literature (Shimizu et al., 2006; Hoyer et al., 2008; Loh & Bühlmann, 2014; Peters & Bühlmann, 2014). Most of these methods, however, are limited to continuous variables and thus cannot handle categorical data, especially in the multivariate setting.

A recent framework explicitly designed to handle categorical data is *entropic causal inference* (Kocaoglu et al., 2017a;b; Compton et al., 2020). At a high level, the underlying assumption of this approach is that true causal mechanisms in nature are often “simple,” taking inspiration from the Occam’s razor principle. The authors adopt an information-theoretic realization of this principle by using “entropy” to measure the complexity of a causal model. As we further explore in this work, entropic causal inference provides a means to measure the amount of randomness a generative model would require to produce an observed distribution. As Occam’s razor prefers simpler explanations, entropic causal inference prefers generative models with small randomness. We do not expect following this preference to always lead to the discovery of true causal relationships (just as one does not expect a simpler explanation to be always be the correct one), but view this as a guiding

¹Massachusetts Institute of Technology, Cambridge, USA
²MIT-IBM Watson AI Lab, Cambridge, USA ³IBM Research, Cambridge, USA ⁴Purdue University, West Lafayette, USA. Correspondence to: Spencer Compton <scompton@mit.edu>.

intuition that mirrors nature and experimental observations. Our experiments on semi-synthetic data demonstrates that the low-entropy assumption indeed holds in certain settings.

Previously, the framework was applied to discovering the causal direction between two random variables given that the amount of randomness in the true causal relationship is small. We focus on extending this framework to learn larger causal graphs instead of just cause-effect pairs. Suppose the observed variables have n states. Our contributions follow:

1. We show pairwise identifiability with strictly relaxed assumptions compared to the previously known results. We enable learning the causal graph $X \rightarrow Y$ from observational data even when (i) the cause variable X has low entropy of $o(\log(n))$ and (ii) the exogenous noise has non-constant entropy, i.e., $\mathcal{O}(1) \ll H(E) = o(\log \log(n))$.
2. We show the first identifiability result for causal graphs with more than two observed variables, with a new peeling algorithm for general graphs.
3. We propose a heuristic algorithm that searches over all DAGs and outputs the one that requires the minimum entropy to fit to the observed distribution.
4. We experimentally evaluate our algorithms and show that entropic approaches outperform the discrete additive noise models in synthetic data. We also apply our algorithms on semi-synthetic data using the *bnlearn*¹ repository and demonstrate the applicability of low-entropy assumptions and the proposed method.

2. Related Work

Learning causal graphs from observational data has been studied extensively in the case of continuous variables. Loh & Bühlmann (2014) proposes an algorithm for learning linear structural causal models when the error variance is known. Similarly, Peters & Bühlmann (2014) show that linear models with Gaussian noise become identifiable if the noise variance is the same for all variables. A more general modeling assumption is the additive noise model (ANM). In (Shimizu et al., 2006), the authors show that for almost all linear causal models, the causal graph is identifiable if the additive exogenous noise is non-Gaussian.

In the case of discrete and/or categorical variables, causal discovery literature is much more sparse. (Cai et al., 2018) introduces a method for categorical cause-effect pairs when there exists a hidden intermediate representation that is compact. In (Daniusis et al., 2010; Janzing et al., 2015), the authors propose using an information-geometric approach called IGCI that is based on independence of cause and the causal mechanism. However IGCI can provably recover

the causal direction only in the case of deterministic relations. An extension of additive noise models to discrete data is done in (Peters et al., 2011) where identifiability is shown between two variables. The authors also propose using the regression-based algorithm of (Mooij et al., 2009) (which made continuous domain ANM applicable to arbitrary graphs) for the discrete setting as well. Without specific assumptions on the graph and the generative mechanisms, this is a heuristic algorithm, i.e., identifiability in polynomial time is not guaranteed by discrete ANM on graphs with more than two nodes.

One related idea is to use Kolmogorov complexity to determine the simplest causal model (Janzing & Schölkopf, 2010). Minimum-description length has been used as a substitute for Kolmogorov complexity (which is not computable) in a series of follow-up papers (Budhathoki & Vreeken, 2017; Marx & Vreeken, 2021). Our extension of entropic causal inference to graphs can be seen as an information-theoretic realization of this promise, where the complexity of the causal model is captured by its entropy. Other information-theoretic concepts such as interaction information (Ghassami & Kiyavash, 2017) and directed information (Etesami & Kiyavash, 2014) have also been studied in the context of causality.

3. Background and Notation

Causal Graphs and Learning: Consider a causal system where each variable is generated as a function of a subset of the rest of the observed variables and some additional randomness. Such systems are modeled by structural equations and are called structural causal models (SCMs). Let $X_1, X_2, \dots, X_{|V|}$ be the set of observed variables. Accordingly, there exists functions f_i and exogenous noise terms E_i such that $X_i = f_i(\text{Pa}_i, E_i)$. This equation in a causal system should be understood as an assignment operator since changing Pa_i affects X_i whereas changing X_i does not affect Pa_i . We say the set of variables Pa_i *cause* X_i . A directed acyclic graph (DAG) can be used to summarize these causal relations, which is called the *causal graph*. We denote the causal graph by $G = (V, \mathcal{E})$ where V is the set of observed nodes and \mathcal{E} is the set of directed edges. There are $|V|$ nodes, $X_1, X_2, \dots, X_{|V|}$, where each X_i corresponds to an observable random variable. Edges are constructed by adding a directed arrow from every node in the set Pa_i to X_i for all i . Pa_i then becomes the set of parents of X_i in G . We assume causal sufficiency, i.e., that there are no unobserved confounders, and that there is no selection bias. Under these assumptions, $\text{Pa}_i \perp\!\!\!\perp E_i$. Additionally, for simplicity of presentation we denote the number of states of all variables as n (i.e. $|X_i| = n$ for all i). Note that our proofs do not require each observed variable to strictly have the same number of states; we merely need them scale together,

¹<https://www.bnlearn.com/bnrepository/>

i.e. if $X_1 \in [n_1]$ and $X_2 \in [n_2]$ then $\frac{n_1}{n_2} = \Theta(1)$. All big-o notation in the paper is relative to n . Our goal is to infer the directed causal graph from the observed joint distribution $p(X_1, X_2, \dots, X_{|V|})$ using the assumptions of the entropic causality framework as needed.

Even without making any parametric assumptions, we can learn some properties of the graph from purely observational data. Algorithms relying on conditional independence tests (such as the PC or IC algorithms (Spirtes et al., 2000; Pearl, 2009)) can identify the *Markov equivalence class* (MEC) of G , i.e. the set of graphs that produce distributions with the exact same set of conditional independence relations. Moreover, a graph’s Markov equivalence class uniquely determines its *skeleton* (the set of edges, ignoring orientation) and *unshielded colliders* (the induced subgraphs of the form $X \rightarrow Z \leftarrow Y$). A Markov equivalence class is summarized by a mixed graph called the *essential graph*, which has the same skeleton and contains a directed edge if all graphs in the equivalence class orient the edge in the same direction. All other edges are undirected. The problem of determining the true causal graph from observational data thus reduces to orienting these remaining undirected edges, given enough samples to perform conditional independence tests reliably.

Entropic Causality Framework: Without interventional data, one needs additional assumptions to refine the graph structure further than the equivalence class. The key assumption of the entropic causality framework is that, in nature, true causal models are often “simple.” In information-theoretic terms, this is formalized as the entropy of exogenous variables often being small. Previous work has shown guarantees for identifying the direction between a causal pair X, Y where $Y = f(X, E), X \perp\!\!\!\perp E$ for some exogenous variable E from observational data. The work of (Kocaoglu et al., 2017a) showed that when the support size of the exogenous variable (i.e. the Renyi-0 entropy $H_0(E) = |E|$) is small, with probability 1 it is impossible to factor the model in the reverse direction (as $X = g(Y, \tilde{E})$) with an exogenous variable with small support size (i.e. $|\tilde{E}|$ must be large). Thus, one can identify the causal pair direction by fitting the smallest cardinality exogenous variable in both directions and checking which direction enables the smaller cardinality. Kocaoglu et al. (2017a) conjectured that this approach also would work well for Shannon entropy. Compton et al. (2020) resolved this conjecture, showing identifiability for causal pairs under particular generative assumptions.

Definition 3.1 ((α, β) -support). A discrete random variable X is said to have (α, β) -support if at least α states of X have probability of at least β .

(Compton et al., 2020) assumes that the cause variable X has $(\Omega(n), \Omega(\frac{1}{n}))$ -support and that the Shannon entropy of the exogenous variable (i.e. $H(E) = H_1(E)$) is small. In

this paper, we say an event holds “with high probability” if the probability of the event not holding is bounded by $O(\frac{1}{n^\alpha})$ for any constant $\alpha > 0$. (Compton et al., 2020) showed that when $H(E) = O(1), H(\tilde{E}) = \Omega(\log(\log(n)))$ with high probability. The high probability statement is with respect to the selection of the function f , i.e., for all but a vanishing (in n) fraction of functions f , identifiability holds. Moreover, they showed that this approach was robust to only having a polynomial number of samples, whereas the result of (Kocaoglu et al., 2017a) that assumed small $|E|$ required knowing the exact joint distribution, e.g. from an oracle or infinite samples.

Algorithmically, one can provably orient causal pairs under the assumptions of (Compton et al., 2020) by comparing the minimum entropy exogenous variable needed to factor the pair in both directions (i.e. comparing the minimum $H(E)$ for which there exists a function f and $E \perp\!\!\!\perp X$ such that $Y = f(X, E)$, and the analogous quantity minimizing $H(\tilde{E})$). Finding this minimum entropy exogenous variable is an optimization problem equivalent to the *minimum-entropy coupling problem* for the conditionals, specifically, the minimum $H(E)$ in the direction $X \rightarrow Y$ is the same as the minimum-entropy coupling for $[(Y|X = i)], \forall i \in [n]$ (Kocaoglu et al., 2017a; Cicalese et al., 2017; Painsky et al., 2019). Accordingly, we denote the entropy of the minimum-entropy coupling for a variable X conditioned on a set S as $\text{MEC}(X|S)$. Compton et al. (2020) showed $\text{MEC}(Y|X) < \text{MEC}(X|Y)$ with high probability.

4. Tightening the Entropic Identifiability Result for Cause-Effect Pairs

In this work, we leverage results for the bivariate entropic causality setting to learn general graphs. Theorem 1 of (Compton et al., 2020) provides identifiability guarantees in the bivariate setting. However, the assumptions of their theorem are not general enough to imply an identifiability result on graphs with more than 2 nodes. Specifically, a fundamental challenge in applying bivariate causality to discover each edge in a larger graph is confounding due to the other variables, i.e., when one considers a pair of variables, the remaining variables act as confounders. These confounders cannot be controlled for since we do not know the causal graph and conditioning on other variables unknowingly creates additional dependencies. One natural approach to handle confounding is to recursively discover source nodes by conditioning on the common causes that are discovered so far in the graph. This idea will form the basis for our peeling algorithm to be proposed in Section 5.1. We are interested in learning graphs where the exogenous variable for every node has small entropy (in particular, $H(E_i) = o(\log(\log(n)))$). When conditioning

on the source nodes, some nodes X (e.g. the children of the source nodes) will thus have conditional entropies of order $H(X|\text{sources}) = o(\log(\log(n)))$ since for the children of source nodes, the only remaining randomness on X will be due to the low-entropy exogenous variable. This creates problems when attempting to orient edges connected to these variables conditioned on the source nodes. Specifically, Theorem 1 of Compton et al. (2020) requires the cause variable X to have $(\Omega(n), \Omega(\frac{1}{n}))$ -support which enforces $H(X) = \Omega(\log(n))$ – and this is not satisfied for the above nodes with $o(\log(\log(n)))$ entropy.

In the following bivariate result, we instead only require $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support, and simultaneously relax the exogenous variable constraint from $H(E) = O(1)$ to $H(E) = o(\log(\log(n)))$. This condition can be satisfied for X with $H(X) = O(1)$ as needed.

Theorem 4.1. *Consider the SCM $Y = f(X, E)$, $X \perp\!\!\!\perp E$, where $X, Y \in [n], E \in [m]$. Suppose E is any random variable with entropy $H(E) = o(\log(\log(n)))$. Let X have $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support. Let f be sampled uniformly randomly from all mappings $f : [n] \times [m] \rightarrow [n]$. Suppose n is sufficiently large. Then, with high probability, any \tilde{E} that satisfies $X = g(Y, \tilde{E})$, $\tilde{E} \perp\!\!\!\perp Y$ for some g , entails $H(\tilde{E}) \geq \Omega(\log(\log(n)))$.*

While interesting in its own right, we apply this tightened bivariate identifiability result to the general graph case in Section 5. Note that the assumption of a uniformly random f (also used in (Compton et al., 2020), (Kocaoglu et al., 2017a)) is not meant as a description of how nature generates causal functions, but as the least-restrictive option for putting a measure on the space of possible functions so that high-probability statements can be made rigorously. Theorem 4.1 can be immediately adapted to any alternative distribution on the space of f that does not assign any individual value of f probability mass more than $n^{c'}$ times the probability mass assigned by the uniform distribution, for some constant c' .

Proof overview for Theorem 4.1. Here we provide the intuition behind the proof strategy, the full proof is given in Appendix A. It is simple to show that the minimum entropy required to fit the function in the incorrect direction, $H(\tilde{E})$, is lower-bounded as $H(\tilde{E}) \geq \max_y H(X|Y = y)$. The overarching goal of our proof method is then to show that there exists a state y' of Y such that $H(X|Y = y') = \Omega(\log(\log(n)))$.

To accomplish this, we start by showing that the $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support of X implies existence of a subset S of $\Omega(n)$ states of X that each have probability $\Omega(\frac{1}{n \log(n)})$ and are all relatively close in probability to each other. We call this subset S , the *plateau* states. If one envisions them as adjacent in the PMF of X , these states would

have similar heights and thus look like a plateau.

Now, we conceptualize the realization of f as a balls-and-bins game, where each element of $X \times E$ (a ball) is mapped i.i.d. uniformly randomly to a state of Y (a bin). Using balls-and-bins arguments, it is our hope to show that there is a bin that receives $\Omega(\frac{\log(n)}{\log(\log(n))})$ *plateau balls* of the form $(X \in S, E = e_1)$, where e_1 is the most probable state of E , and that this will cause the corresponding conditional distribution to have large entropy. The primary intuition is that a bin receiving many plateau balls would cause the corresponding conditional distribution to have many plateau states that all have near-uniform probabilities, and this near-uniform subset of the conditional distribution would contribute a significant fraction of the probability mass to guarantee that its entropy is large. With the stronger assumptions on (α, β) -support by (Compton et al., 2020), this proof method suffices. However, as we are assuming a weaker notion of (α, β) -support, it is not clear that the plateau balls would make up a significant fraction of the conditional’s mass to guarantee large entropy.

In a sense, the plateau balls are probability masses that are “helping” us make some conditional entropy large. The proof of (Compton et al., 2020) takes the perspective that all remaining mass from non-plateau states are “hurting” our effort to make a conditional distribution with large entropy. To accommodate our relaxed assumptions, we take a more nuanced perspective on *helpful* and *hurtful* mass. Consider a non-plateau state x of X that contributes a small amount of mass towards the conditional distribution corresponding to a state y of Y . With the perspective of (Compton et al., 2020), this would be viewed as *hurtful* mass because it is from a non-plateau state of X . But intuitively, in terms of its contribution to $H(X|Y = y)$, it does not matter whether x is a plateau state or not. Through careful analysis, we can show that if $P(X = x|Y = y)$ is small then they are not “too hurtful.” We follow this intuition to make a new definition of the good mass, where we set a threshold \mathcal{T} , define the first \mathcal{T} mass we receive from a non-plateau state of X as *helpful* mass for the state of Y , and the surplus beyond \mathcal{T} from the non-plateau state of X as *hurtful* mass for the state of Y . As before, all mass from plateau states will be helpful. With this new perspective and a careful analysis, we show that there is a state y' that receives many plateau balls, and has much more helpful mass than hurtful mass. This then enables us to show that $H(X|Y = y')$ is large, proving the theorem.

5. Learning Graphs via Entropic Causality

Now, we focus on how to leverage the capability of correctly orienting causal pairs to learn causal graphs *exactly*. In comparison, traditional structure learning methods only learn the Markov equivalence class of graphs from observational data.

For example, given the line graph $X \rightarrow Y \rightarrow Z$, such methods would deduce the true graph is either $X \rightarrow Y \rightarrow Z$ or $X \leftarrow Y \leftarrow Z$, but not that it is exactly $X \rightarrow Y \rightarrow Z$.

As was discussed in Section 3, learning the entire graph can be reduced to correctly orienting each edge in the skeleton. However, we cannot naively use a pairwise algorithm, as the rest of the observed variables can act as confounders. We examine how different pairwise oracles can enable us to characterize the value of using minimum entropy couplings to learn causal graphs. One example of a natural-feeling oracle is one that can correctly orient any edges that have no active confounding. Such an oracle enables learning of directed trees and complete graphs. However, it cannot be used to learn all general graphs (see Appendix C for an example). We propose an alternative oracle, that can distinguish between a source node and any node it can reach:

Definition 5.1 (Source-pathwise oracle). A source-pathwise oracle for a DAG G always orients $A \rightarrow B$ if A is a source and there exists a directed path from A to B in G .

Let us formalize our entropic method for causal pairs as the following oracle:

Definition 5.2 (MEC oracle). A minimum entropy coupling (MEC) oracle returns $X \rightarrow Y$ if $\text{MEC}(Y|X) < \text{MEC}(X|Y)$ and $X \leftarrow Y$ otherwise, given the joint distribution $p(X, Y)$.

We aim to show that our MEC oracle is a source-pathwise oracle for graphs with the following assumptions:

Assumption 5.3 (Low-entropy assumption). Consider an SCM where $X_i = f_i(\text{Pa}_i, E_i)$, $\text{Pa}_i \perp\!\!\!\perp E_i, \forall i$, where $X_i \in [n], E_i \in [m]$. Suppose $|V| = O(1)$, $H(E_i) = o(\log(\log(n)))$ and E_i has $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support for all i , and f_i are sampled uniformly randomly from all mappings $f_i: [m] \times [n]^{|\text{Pa}_i|} \rightarrow [n]$.

We are now ready to show the main result of our paper. We show that, under certain generative model assumptions, applying entropic causality on pairs of observed variables acts as a source-pathwise oracle for DAGs:

Theorem 5.4. *For any SCM under Assumption 5.3, the MEC oracle is a source-pathwise oracle for the causal graph with high probability for sufficiently large n .*

Characterizing entropic causality as a source-pathwise oracle enables us to identify the true causal graph for general graphs. We outline the key intuitions of our proof:

Proof overview for Theorem 5.4. Suppose X_{src} is a source and Y is a node such that there is a path from X_{src} to Y . To show the MEC oracle is a source-pathwise oracle, we show that $\text{MEC}(X_{\text{src}}|Y) > \text{MEC}(Y|X_{\text{src}})$. As in Theorem 4.1, we will accomplish this by showing there is a state y' of Y such that $H(X_{\text{src}}|Y = y') = \Omega(\log(\log(n)))$.

We begin by conceptualizing the realization of all f_i as a balls-and-bins game. Every node X_i is a uniformly random function f_i of $\text{Pa}(X_i) \cup E_i$. Let us define each ball as the concatenation of X and all E_i other than E_{src} . More formally, we denote each ball as $(X = x, E_1 = e_1, \dots, E_{\text{src}-1} = e_{\text{src}-1}, E_{\text{src}+1} = e_{\text{src}+1}, \dots, E_{|V|} = e_{|V|})$, and each ball has a corresponding probability mass of $P(X = x) \times \prod_{i \neq \text{src}} P(E_i = e_i)$. To view the realization of f_i as a balls-and-bins game, we consider the nodes in an arbitrary topological order for the graph. When we process a node X_i , we group balls according to their configuration of $(\text{Pa}(X_i) \cup E_i)$. This is because balls with the same configuration correspond to the same cell of the function f_i . For each group of balls that all share the same configuration, we uniformly randomly sample a state of X_i to assign all the balls in the group. This is essentially realizing one cell of f_i . Groups are assigned independently of other groups. In this sense, each realization of f_i is a balls-and-bins game where we group balls by their configuration, and throw them together into states of X_i (bins).

Let us define plateau balls as those who have a plateau state of X_{src} and have the most probable state of E_i for every $i \neq \text{src}$. As was done in Theorem 4.1, our goal is to show that there will be a state y' of Y such that y' receives many plateau balls and much more helpful mass than hurtful mass. However, it is not immediately clear how to show this in the graph setting. For intuition, we explore two special cases.

Consider the case of a line graph (Figure 1(a)). For simplicity of this proof overview, assume all X_i other than X_{src} are deterministic functions of their parents (i.e., $H(E_i) = 0$). Using techniques similar to Theorem 4.1, we can show there are many bins of X_2 that receive many plateau balls and much more helpful mass than hurtful mass. Moreover, we can then use similar techniques to show a non-negligible proportion of those bins will have their corresponding balls mapped together to a bin of X_3 where it does not encounter much hurtful mass. We can repeat this argument again to show some of these desirable bins “survive” from X_3 to Y . While only a scalingly small fraction of these desirable bins “survive” each level, this still ensures the survival of at least one bin if the number of vertices is constant. This will accomplish our goal of having a state y' of Y with many plateau balls and much more helpful mass than hurtful mass.

On the other hand, consider the case of a diamond graph in Figure 1(b). Again, assume for simplicity that all X_i other than X_{src} are deterministic functions of their parents. Note that when we realize f_Y , two balls are always mapped independently unless they share the same configuration of $\text{Pa}(Y) = \{X_2, X_3\}$. We observe that X_2 and X_3 are both independent deterministic functions of X_{src} . Accordingly, the probability of two particular states $x, x' \in X_{\text{src}}$ satisfying $f_2(x) = f_2(x')$ and $f_3(x) = f_3(x')$ is equal to




Figure 1. Graphs colored according to the Random Function Graph Decomposition (Definition 5.5) used in the proof of Theorem 5.4.

$\frac{1}{n^2}$. Therefore, almost all pairs of balls are mapped to Y from X_{src} independently. Since everything is almost-independently mapped to Y , we can treat it like a bivariate problem and use techniques similar to Theorem 4.1.

We are able to prove correctness for both of these graphs, but we do so in ways that are essentially opposite. For the line graph, we utilize strong dependence as bins with desired properties “survive” throughout the graph. For the diamond graph, we utilize strong independence as balls are all mapped to Y essentially independently. To combine the intuitions of these two cases into a more general proof, we introduce the Random Function Graph Decomposition:

Algorithm 1 Learning general graphs with oracle

```

1:  $\mathcal{R} \leftarrow \{1, \dots, |V|\}$  {set of remaining nodes}
2:  $\mathcal{I} \leftarrow \emptyset$  {set of pairs found to be conditionally independent}
3:  $\mathcal{T} \leftarrow []$  {list of nodes in topological order}
4: while  $|\mathcal{R}| > 0$  do
5:    $\mathcal{N} \leftarrow \emptyset$  {set of nodes discovered as non-sources}
6:    $\mathcal{C} \leftarrow \{1, \dots, |V|\} \setminus \mathcal{R}$  {condition on previous sources}
7:   for all  $(X_i, X_j) \in \{\mathcal{R} \times \mathcal{R}\}$  do
8:     if  $X_i \notin \mathcal{N}$  and  $X_j \notin \mathcal{N}$  and  $(X_i, X_j) \notin \mathcal{I}$  then
9:       if  $\text{CI}(X_i, X_j | \mathcal{C})$  then
10:         $\mathcal{I} \leftarrow \mathcal{I} \cup (X_i, X_j)$ 
11:       else if  $\text{Oracle}(X_i, X_j | \mathcal{C})$  orients  $X_i \rightarrow X_j$  then
12:         $\mathcal{N} \leftarrow \mathcal{N} \cup \{X_j\}$  { $X_j$  is not a source}
13:       else
14:         $\mathcal{N} \leftarrow \mathcal{N} \cup \{X_i\}$  { $X_i$  is not a source}
15:       end if
16:     end if
17:   end for
18:    $\mathcal{S} \leftarrow \mathcal{R} \setminus \mathcal{N}$  {the remaining nodes that are a source}
19:    $\mathcal{R} \leftarrow \mathcal{R} \setminus \mathcal{S}$  {remove sources from remaining nodes}
20:   for all  $X_i \in \mathcal{S}$  do
21:     append  $X_i$  to  $\mathcal{T}$ 
22:   end for
23: end while {Now,  $\mathcal{T}$  is a valid topological ordering}
24: for all  $(i, j) \in \{1, \dots, |V|\}^2$  where  $i < j$  do
25:   if  $\text{CI}(\mathcal{T}(i), \mathcal{T}(j) | \{\mathcal{T}(1), \dots, \mathcal{T}(j-1)\} \setminus \mathcal{T}(i))$  then
26:     no edge between  $\mathcal{T}(i)$  and  $\mathcal{T}(j)$ 
27:   else
28:     orient  $\mathcal{T}(i) \rightarrow \mathcal{T}(j)$ 
29:   end if
30: end for
    
```

Definition 5.5 (Random Function Graph Decomposition). Given a DAG and a pair of nodes (X_{src}, Y) , Random Function Graph Decomposition colors the nodes iteratively fol-

lowing any topological order of the nodes as follows:

1. Color the node with a new color if X_{src} is a parent of the node or if the node has parents of different colors.
2. Color the node with the color of its parents if all of the node’s parents have the same color.

Using the Random Function Graph Decomposition, we claim that when a node is assigned a new color as in step 1, we utilize independence as in the diamond graph (Figure 1(b)), and when a node inherits its color as in step 2 we utilize dependence as in the line graph in Figure 1(a). We illustrate Figure 1(c) as an example. With a careful analysis, we utilize these intuitions to prove the MEC oracle is a source-pathwise oracle with high probability.

5.1. Peeling Algorithm for Learning Graphs

In the previous section, we have shown how entropic causality can be used as a source-pathwise oracle. Next, we show how to learn general graphs with a source-pathwise oracle. Our algorithm will iteratively determine the graph’s sources, condition on the discovered sources, determine the graph’s sources after conditioning, and so on. Doing this will enable us to find a valid lexicographical ordering of the graph. Given a lexicographical ordering, we can learn the skeleton with $O(n^2)$ conditional independence tests.

Now, we outline how we iteratively find the sources. In each stage, we consider all the remaining nodes as candidate sources. It is our goal to remove all non-sources from our set of candidates. To do this, we iterate over all pairs of candidates and do a conditional independence test *conditioned on the sources that are found so far*. If the pair is conditionally independent, we do nothing. We note that this will never happen for a pair where one node is a true source and the other node is reachable from the source through a directed path: Conditioning on previously found sources cannot d-separate such paths. Otherwise, the pair is conditionally dependent. We then use the source-pathwise oracle to orient between the two nodes, and eliminate the sink node of the orientation as a candidate (i.e., if we orient $A \rightarrow B$, we eliminate B as a candidate source).




Figure 2. Performance of methods in the unconstrained setting in the triangle graph $X \rightarrow Y \rightarrow Z, X \rightarrow Z$: 50 datasets are sampled for each configuration from the unconstrained model $X = f(\text{Pa}_X, E_X)$. The x -axis shows entropy of the exogenous noise. The exogenous noise of the first variable is fixed to be large (≈ 3.3 bits), hence it is a high entropy source (HES). Entropic methods consistently outperform the ANM algorithm in almost all regimes.

Suppose two nodes are dependent conditioned on the past sources. Then either the pair contains a source node and a descendant of the source node, or it contains two non-source nodes. In the former case when the pair contains a source node the source-pathwise oracle will always orient correctly and the non-source node will be eliminated. In the latter case when the pair are two non-sources we can safely eliminate either as a source candidate and accordingly oracle output is irrelevant. By the end of this elimination process, we can show that only true sources will remain as candidates in each step, which enables us to obtain a valid lexicographical ordering, and thus learn the causal graph. We summarize this procedure as Algorithm 1. The following theorem shows the correctness of Algorithm 1 given a source-pathwise oracle:

Theorem 5.6. *Algorithm 1 learns any faithful causal graph $D = (V, E)$ with $\mathcal{O}(|V|^2)$ calls to a source-pathwise oracle and $\mathcal{O}(|V|^2)$ conditional independence tests.*

Note that Theorem 5.6 relies on a weaker notion of faithfulness in the causal graph. In particular, it purely relies on X_i and X_j being dependent conditioned on some set S , if there is a directed path from X_i to X_j (or vice versa) and S forms a topological prefix of the graph. Note that under Assumption 5.3 this faithfulness property will hold with very high probability (at least $1 - O(1/2^{\Omega(n)})^2$), meaning it is not positing any assumption beyond Assumption 5.3.

Finally, we show that we can use entropic causality together with Algorithm 1 for learning general causal graphs:

²Given the $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support for E_i , one can show that X_i will have a support size of at least $\Omega(n)$, with probability at least $1 - O(1/2^{\Omega(n)})$. The probability of any two particular states of X_i producing the same conditional distribution of $X_j | X_i = x_i$ is at most $1/n$ (by considering the realization of the last probability mass from X_i). Accordingly, the probability of this happening for $\Omega(n)$ states of X_i is bounded by $1/n^{\Omega(n)}$.

Corollary 5.7. *For any SCM under Assumption 5.3, using entropic causality for pairwise comparisons in Algorithm 1 learns, with high probability, the causal graph that is implied by the SCM.*

6. Experiments

We first introduce a heuristic that we call the *entropic enumeration* algorithm. In this algorithm, we enumerate over all possible causal graphs consistent with the skeleton and calculate the minimum entropy needed to generate the observed distribution from the graph with independent noise at each node. The minimum entropy needed to generate the joint distribution with some graph D is $\sum_{X_i} \text{MEC}(X_i | \text{Pa}_D(X_i))$ where $\text{Pa}_D(X_i)$ denotes the parents of X_i in D . The graph requiring the least randomness is then selected.

We are not aware of any provably correct method for causal discovery between categorical variables that are non-deterministically related. For discrete variables, the only such method other than entropic causality is the discrete additive noise model (Peters et al., 2011). We compare entropic causality to discrete ANM for learning causal graphs, using the graph extension of ANM proposed by Mooij et al. (2009). To isolate the role of our algorithms in identifying causal graphs beyond the equivalence class, we support every algorithm in our comparisons with the skeleton of the true graph (obtainable from conditional independence tests given enough data). We evaluate performance via the structural Hamming distance (SHD) from the estimated graph to the true causal graph. See the Appendix for implementation details.

Performance on Synthetic Data. Figure 2 compares the performance of entropic peeling, entropic enumeration and discrete ANM algorithms for the triangle graph, i.e., the graph with edges $X \rightarrow Y, Y \rightarrow Z$, and $X \rightarrow Z$. Every

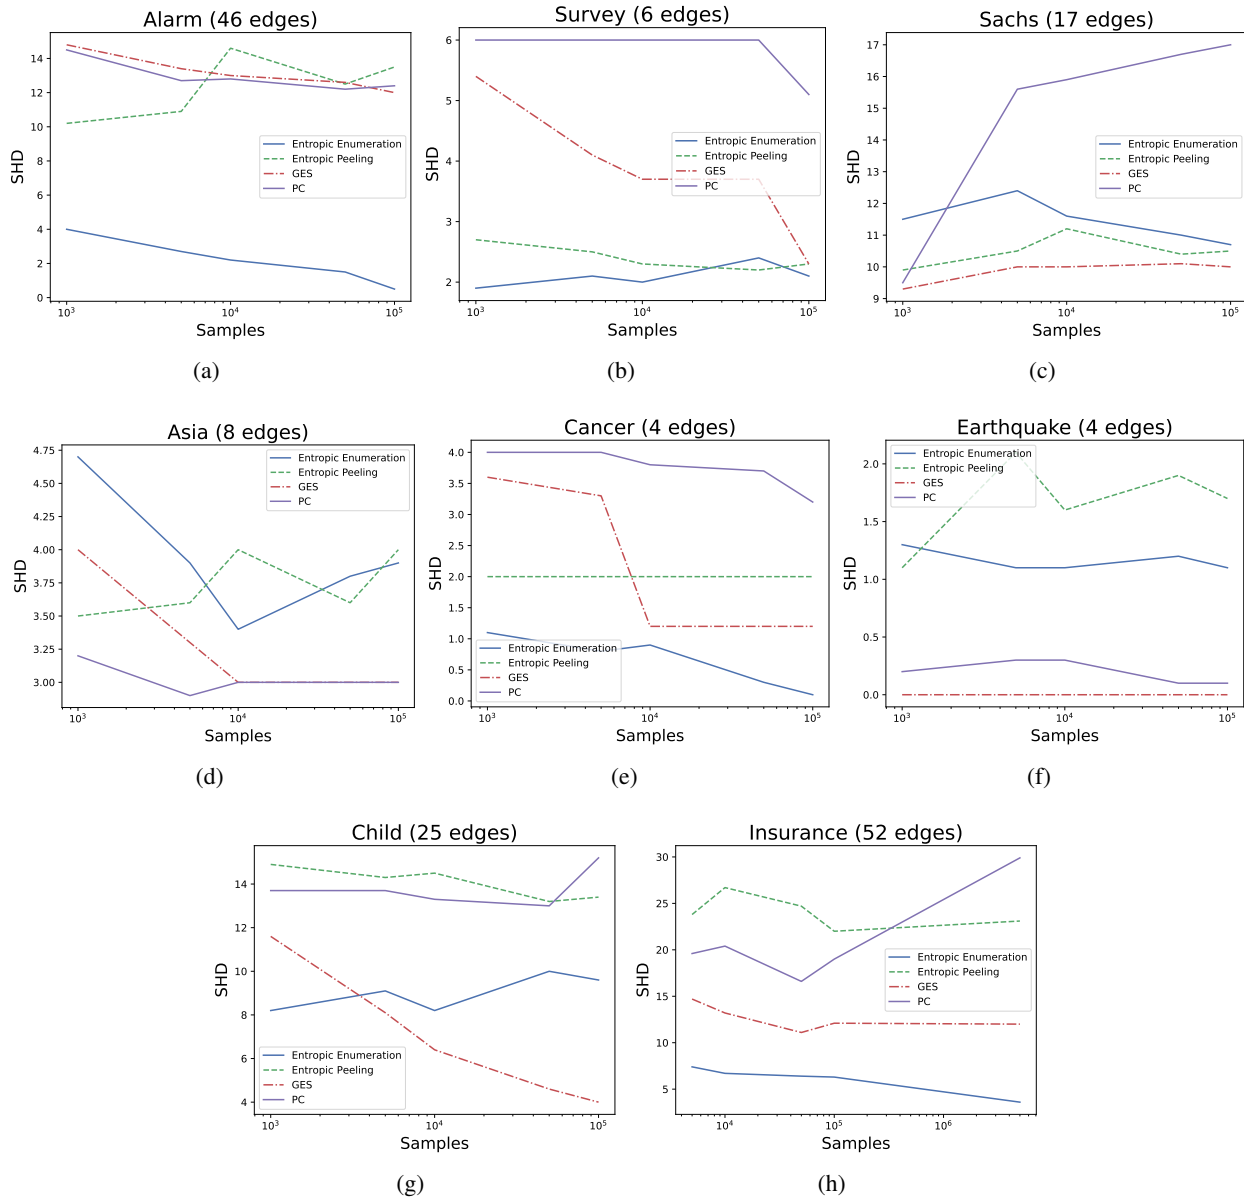


Figure 3. Performance of methods on networks from the *bnlearn* repository with varying samples: 10 datasets are sampled for each configuration from the *bnlearn* network. E.g., entropic enumeration exactly recovers Alarm, no algorithm correctly learns half of Sachs.

datapoint is obtained by averaging the SHD to the graph for 50 instances of structural models. To ensure that the entropy of the exogenous nodes are close to the value on the x-axis, their distributions are sampled from a Dirichlet distribution with a parameter that is obtained through a binary search. We observe that the entropic methods consistently outperform the ANM approach. Importantly, we observe how entropic methods are able to near-perfectly learn the *exact* triangle graph in almost all regimes, even though all triangle graphs are in the same Markov equivalence class and thus traditional structure learning algorithms like PC

or GES cannot learn anything. With enough samples, entropic enumeration learns the graph near-perfectly until the exogenous noise nears $\log(n)$, exceeding our theoretical guarantee of $o(\log(\log(n)))$. In Figure 2, we fix the source node to have high entropy. Our motivation is that if all nodes have essentially zero randomness, then we expect the performance of any method to degrade as there is no randomness in samples to observe causality or faithfulness. In Figure 9 in Appendix, we do not fix a high-entropy source and still observe that entropic methods outperform ANM in almost all regimes. Experiments with different and larger graphs




Figure 4. Performance of methods in the ANM setting on the line graph $X \rightarrow Y \rightarrow Z$: 25 datasets are sampled for each configuration from the ANM model $X = f(\text{Pa}_X) + N$. The x -axis shows entropy of the additive noise.

can be seen in the Appendix.

Performance in Discrete Additive Noise Regime. In Figure 4, we compare the performance of the entropic algorithms and discrete ANM *when the true SCM is a discrete additive noise model*. Using the discrete ANM generative model, we observe that entropic enumeration out-performs the discrete ANM method with few samples and matches its performance with many samples. This demonstrates that even though entropic methods are designed for the general unconstrained SCM class, they perform similarly to ANM which was designed specifically for this setting.

Effect of Finite Samples. We observe that entropic methods, particularly enumeration, work well even in regimes with low samples. Experiments focusing on the impact of finite samples can be found in the Appendix.

Performance on Real-World Data. Due to the computational cost of discrete ANM, we compared entropic causality against GES and PC algorithms to evaluate how well it learned real-world causal graphs from the *bnlearn* repository beyond their equivalence class. Figure 3 shows performance on the eight networks we evaluated. Of particular interest is Figure 3(a), where entropic enumeration almost perfectly identifies a graph with 46 edges from its skeleton and finite samples. Again, we do not claim that the assumptions of entropic causality are universally true in nature, but instead that there are real settings such as Figure 3(a) where the framework enables us to learn causal graphs. Our experiments, exceeding our best theoretical guarantees, show that even when the number of nodes is the same as the number of states, entropic causality can be used for learning the causal graph with a moderate number of samples.

7. Conclusion

In this work, we have extended the entropic causality framework to graphs. An identifiability result was proven, and

two algorithms were presented and experimentally evaluated — a theoretically-motivated sequential peeling algorithm and a heuristic entropic enumeration algorithm that performs better on small graphs. Overall, we observed strong experimental results in settings much more general than the assumptions used in our theory, indicating that a much stronger theoretical analysis might be possible. We note however that the quantity $H(E_i) = \Theta(\log(\log(n)))$ appears to be approximately a phase transition for the balls-and-bins setting, and posit that the development of novel tools may be required for such an extension of the theory.

We suggest such an advancement may involve an increased focus on a total entropy criterion (i.e., an extension of comparing $H(X) + H(E)$ to $H(Y) + H(\tilde{E})$ in the bivariate case), as in our proposed algorithm of entropic enumeration. Experiments indicate that this performs well, and one might argue that it appears to be more conceptually justified. For one, it mirrors Occam’s razor in that it prefers the causal graph with minimal total randomness required. While we do not claim that this methodology will always discover the true generative model (as Occam’s razor does not require the simplest explanation to *always* be true), we believe these intuitions mirror nature more often than not, as confirmed by our experimental results. Moreover, such an approach appears to fare better with counter-examples for exogenous-based criterion such as the traveling ball scenario of (Janzing, 2019) discussed in (Compton et al., 2020). Showing theoretical guarantees for this approach’s performance is of interest in future work, and can be framed more generally as, “*Under what conditions is the true generative model the most information-theoretically efficient way to produce a distribution?*”

References

- Athey, S. Machine learning and causal inference for policy evaluation. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 5–6, 2015.
- Bareinboim, E. and Pearl, J. Transportability of causal effects: Completeness results. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 26, 2012.
- Budhathoki, K. and Vreeken, J. Mdl for causal inference on discrete data. In *2017 IEEE International Conference on Data Mining (ICDM)*, pp. 751–756. IEEE, 2017.
- Cai, R., Qiao, J., Zhang, K., Zhang, Z., and Hao, Z. Causal discovery from discrete data using hidden compact representation. *Advances in neural information processing systems*, 2018:2666, 2018.
- Cicalese, F., Gargano, L., and Vaccaro, U. How to find a joint probability distribution of minimum entropy (almost) given the marginals. In *IEEE International Symposium on Information Theory (ISIT)*, pp. 2173–2177, 2017.
- Compton, S., Kocaoglu, M., Greenewald, K., and Katz, D. Entropic causal inference: Identifiability and finite sample results. In *Advances in Neural Information Processing Systems*, volume 33, pp. 14772–14782, 2020.
- Cunha, J. Lecture notes of 15-859m: Randomized algorithms; lec 8: Balls and bins/two choices, February 2011.
- Daniusis, P., Janzing, D., Mooij, J., Zscheischler, J., Steudel, B., Zhang, K., and Schölkopf, B. Inferring deterministic causal relations. In *26th Conference on Uncertainty in Artificial Intelligence (UAI 2010)*, pp. 143–150. AUAI Press, 2010.
- Etesami, J. and Kiyavash, N. Directed information graphs: A generalization of linear dynamical graphs. In *2014 American Control Conference*, pp. 2563–2568. IEEE, 2014.
- Ghassami, A. and Kiyavash, N. Interaction information for causal inference: The case of directed triangle. In *IEEE International Symposium on Information Theory (ISIT)*, pp. 1326–1330, 2017.
- Hoyer, P., Janzing, D., Mooij, J. M., Peters, J., and Schölkopf, B. Nonlinear causal discovery with additive noise models. *Advances in Neural Information Processing Systems*, 21:689–696, 2008.
- Janzing, D. The cause-effect problem: Motivation, ideas, and popular misconceptions. In *Cause Effect Pairs in Machine Learning*, pp. 3–26. Springer, 2019.
- Janzing, D. and Schölkopf, B. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory*, 56(10):5168–5194, 2010.
- Janzing, D., Steudel, B., Shajarisales, N., and Schölkopf, B. Justifying information-geometric causal inference. In *Measures of Complexity*, pp. 253–265. Springer, 2015.
- Kocaoglu, M., Dimakis, A., Vishwanath, S., and Hassibi, B. Entropic causal inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017a.
- Kocaoglu, M., Dimakis, A. G., Vishwanath, S., and Hassibi, B. Entropic causality and greedy minimum entropy coupling. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 1465–1469. IEEE, 2017b.
- Loh, P.-L. and Bühlmann, P. High-dimensional learning of linear causal networks via inverse covariance estimation. *The Journal of Machine Learning Research*, 15(1):3065–3105, 2014.
- Magliacane, S., van Ommen, T., Claassen, T., Bongers, S., Versteeg, P., and Mooij, J. M. Domain adaptation by using causal inference to predict invariant conditional distributions. In *Advances in Neural Information Processing Systems*, pp. 10869–10879, 2018.
- Marx, A. and Vreeken, J. Formally justifying mdl-based inference of cause and effect. *arXiv preprint arXiv:2105.01902*, 2021.
- Mooij, J., Janzing, D., Peters, J., and Schölkopf, B. Regression by dependence minimization and its application to causal inference in additive noise models. In *Proceedings of the 26th International Conference on Machine Learning*, pp. 745–752, 2009.
- Moraffah, R., Karami, M., Guo, R., Raglin, A., and Liu, H. Causal interpretability for machine learning-problems, methods and evaluation. *ACM SIGKDD Explorations Newsletter*, 22(1):18–33, 2020.
- Morgan, S. L. and Winship, C. *Counterfactuals and causal inference*. Cambridge University Press, 2015.
- Painsky, A., Rosset, S., and Feder, M. Innovation representation of stochastic processes with application to causal inference. *IEEE Transactions on Information Theory*, 66(2):1136–1154, 2019.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Pearl, J. and Bareinboim, E. External validity: From do-calculus to transportability across populations. *Statistical Science*, 29(4):579–595, 2014.

- Peters, J. and Bühlmann, P. Identifiability of gaussian structural equation models with equal error variances. *Biometrika*, 101(1):219–228, 2014.
- Peters, J., Janzing, D., and Schölkopf, B. Causal inference on discrete data using additive noise models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2436–2450, 2011.
- Scutari, M. Learning bayesian networks with the bnlearn R package. *Journal of Statistical Software*, 35(3):1–22, 2010. doi: 10.18637/jss.v035.i03.
- Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A., and Jordan, M. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10), 2006.
- Spirites, P., Glymour, C. N., Scheines, R., and Heckerman, D. *Causation, Prediction, and Search*. MIT Press, 2000.
- Squires, C., Shen, D., Agarwal, A., Shah, D., and Uhler, C. Causal imputation via synthetic interventions. *arXiv preprint arXiv:2011.03127*, 2020.
- Wajc, D. Negative association: definition, properties, and applications. *Manuscript, available from <https://googl/j2ekqM>*, 2017.
- Zhang, K., Gong, M., Stojanov, P., Huang, B., LIU, Q., and Glymour, C. Domain adaptation as a problem of inference on graphical models. *Advances in Neural Information Processing Systems*, 33, 2020.

A. Proof of Theorem 4.1

A.1. Proof Outline

Following the approach described in the proof overview in the main text (with the descriptions of helpful and hurtful mass), we first introduce the *surplus* of a state of Y to characterize the amount of hurtful mass it receives:

Recall that S is the set of ‘‘plateau states’’ of X , i.e., those whose probabilities are close to one another.

Definition A.1 (Surplus). We define the surplus of a state y of Y as $z_y = \sum_{j \notin S} \max(0, P(X = j, Y = i) - \mathcal{T})$.

Intuitively, only values from states of X outside the plateau states which exceed the threshold will significantly ‘‘hurt’’ conditional entropy $H(X|Y = y)$. We will show there is a state y' of Y where $z_{y'}$ is small and y' receives $\Omega(\frac{\log(n)}{\log(\log(n))})$ plateau balls. To bound $z_{y'}$, we will characterize it as the sum of contributions from three types of balls from $(X \setminus S) \times E$.³

Definition A.2 (Ball characterizations). We characterize three types of balls:

1. **Dense balls.** Consider a set L of states of X , where a state of X is in L if $P(X = x) \geq \frac{1}{\log^3(n)}$. Dense balls are all balls of the form $(x \in L, e \in E)$. We call these dense balls, because the low-entropy of E will prevent the collective mass of these balls from ‘‘expanding’’ well.
2. **Large balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $\geq \frac{\mathcal{T}}{2}$.
3. **Small balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $< \frac{\mathcal{T}}{2}$.

We will show there are a non-negligible fraction of bins such that z_y is small. To do so, we will bound the contribution from dense balls by showing that the small entropy of E prevents ‘‘spread’’ in a sense, as there cannot be many states of Y that receive much contribution towards z_y from these dense balls. We will bound contribution from large balls by bounding the number of large balls, and showing that a non-negligible number of bins receive no large balls. Finally, we will bound contribution from small balls by showing how they often are mapped to states of Y that have yet to receive $\frac{\mathcal{T}}{2}$ mass from the corresponding state of X , meaning they often don’t immediately increase z_y .

Finally, we will show (with high probability) the existence of a bin y' with small $z_{y'}$ that will receive many plateau balls, and how this will imply $H(X|Y = y') = \Omega(\log(\log(n)))$.

³In the proofs, with a slight abuse of notation, we use X, E both for the observed and exogenous variables, respectively and their supports.

A.2. Complete Proof

Bounding $H(\tilde{E})$ via $H(X|Y = y)$. Because $\tilde{E} \perp\!\!\!\perp Y$, it must be true that $H(\tilde{E}) \geq \max_y H(X|Y = y)$. This is simple to prove by data processing inequality and is shown in Step 1 of the proof of Theorem 1 by (Compton et al., 2020). We aim to show there exists a y' such that $H(X|Y = y') = \Omega(\log(\log(n)))$.

Showing existence of a near-uniform plateau. First, we aim to find a subset of the support of X whose probabilities are multiplicatively close to one another. Here, we have a looser requirement for closeness than (Compton et al., 2020). Instead of requiring these probabilities to be within a constant factor of each other, we allow them to be up to a factor of $\log^{c_{\text{close}}}(n)$ apart where c_{close} is a constant such that $0 < c_{\text{close}} < 1$. While there are multiple values of c_{close} that would be suitable for our analysis, for simplicity of presentation we choose $c_{\text{close}} = \frac{1}{4}$ throughout. This set of states of X that are multiplicatively close to one another will be called the *plateau* of X . We begin by showing how the $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support assumption implies a plateau of states of X :

Lemma A.3 (Plateau existence). *Suppose X has $(c_{\text{support}}n, \frac{1}{c_{\text{lb}}n \log(n)})$ -support for constants $0 < c_{\text{support}} \leq 1$ and $c_{\text{lb}} \geq 1$. Additionally, assume n is sufficiently large such that $\frac{\log(2c_{\text{lb}}/c_{\text{support}})}{\log(\log(n))} \leq 1$. Then, there exists a subset $S \subseteq [n]$ of the support of X , such that the following three statements hold:*

1. $\frac{\max_{i \in S} P(X=i)}{\min_{i \in S} P(X=i)} \leq \log^{c_{\text{close}}}(n)$
2. $\min_{i \in S} P(X = i) \geq \frac{1}{c_{\text{lb}}n \log(n)}$
3. $|S| \geq \frac{c_{\text{close}}c_{\text{support}}n}{6}$, for any $0 < c_{\text{close}} < 1$.

Proof. By definition of $(c_{\text{support}}n, \frac{1}{c_{\text{lb}}n \log(n)})$ support, there are at least $c_{\text{support}}n$ states of X with probability in range $[\frac{1}{c_{\text{lb}}n \log(n)}, 1]$. Moreover, at most $\frac{c_{\text{support}}n}{2}$ states will have probabilities in range $[\frac{2}{c_{\text{support}}n}, 1]$. Otherwise, they would have total probability mass > 1 which is impossible. Therefore, there are at least $\frac{c_{\text{support}}n}{2}$ states with probabilities in range $[\frac{1}{c_{\text{lb}}n \log(n)}, \frac{2}{c_{\text{support}}n}]$.

Now, we aim to divide the range $[\frac{1}{c_{\text{lb}}n \log(n)}, \frac{2}{c_{\text{support}}n}]$ into a number of contiguous segments such that all values in any segment are multiplicatively within $\log^{c_{\text{close}}}(n)$ of each other. To do so, we can create segments $[\frac{1}{c_{\text{lb}}n \log(n)} \times (\log^{c_{\text{close}}}(n))^i, \frac{1}{c_{\text{lb}}n \log(n)} \times (\log^{c_{\text{close}}}(n))^{i+1}]$ from $i = 0$ until the smallest i that satisfies $\frac{1}{c_{\text{lb}}n \log(n)} \times (\log^{c_{\text{close}}}(n))^{i+1} \geq \frac{2}{c_{\text{support}}n}$. Accordingly, we need $\lceil \frac{\log((2/(c_{\text{support}}n))/(1/(c_{\text{lb}}n \log(n))))}{\log(\log^{c_{\text{close}}}(n))} \rceil \leq 1 + \frac{1}{c_{\text{close}}}$

$\frac{\log(2c_{\text{lb}}/c_{\text{support}})}{c_{\text{close}} \log(\log(n))} \leq \frac{3}{c_{\text{close}}}$ groups. Hence one group must have at least $\frac{c_{\text{support}}n/2}{3/c_{\text{close}}} = \frac{c_{\text{close}}c_{\text{support}}n}{6}$ states of X that are multiplicatively within $\log^{c_{\text{close}}}(n)$ and have probability at least $\frac{1}{c_{\text{lb}}n \log(n)}$. \square

Lower-bounding the most probable state of E . Our proof method focuses on a balls-and-bins game where states of $X \times E$ are balls and states of Y are bins. We focus first on *plateau balls*, which are balls corresponding to states of S (the set of plateau states of X) and the highest probability state of E . In particular, they are balls of the form $(X \in S, E = e_1)$ where e_1 is the most probable state of E . To show that these plateau balls have enough probability mass to be helpful, we first show that $P(E = e_1)$ is relatively large:

Lemma A.4. *If $H(E) \leq c_{\text{close}} \log(\log(n))$ then $P(E = e_1) \geq \frac{1}{\log^{c_{\text{close}}}(n)}$*

Proof. For any distribution with entropy H , its state with the highest probability has at least probability 2^{-H} (see Lemma 5 of (Compton et al., 2020)). Thus if $H(E) \leq c_{\text{close}} \log(\log(n))$ then $P(E = e_1) \geq 2^{-c_{\text{close}} \log(\log(n))} = \frac{1}{\log^{c_{\text{close}}}(n)}$. \square

Introducing surplus. We now begin proving how there exists a bin that receives a large amount of mass that helps the bin have large conditional entropy (such helpful mass includes the plateau balls), and not much mass that hurts the conditional entropy making it small. To formalize this hurtful mass, recall the *surplus* quantity described in Definition A.1. This surplus is a way of quantifying the probability mass received by a state of Y that is hurtful towards making the conditional entropy large. We define surplus, with the threshold of \mathcal{T} specified as $\frac{12}{n \log(n)}$ as follows:

Definition A.5 (Surplus, $\mathcal{T} = \frac{12}{n \log(n)}$). We define the surplus of a state i of Y as $z_i = \sum_{j \notin S} \max(0, P(X = j, Y = i) - \frac{12}{n \log(n)})$, where S is the set of plateau states of X .

Characterizing balls-and-bins. Now we will show that there are a non-negligible number of states of Y where the surplus is small. Recall from our proof outline that we view the process of realizing the random function f as a balls-and-bins game. In particular, each element of $X \times E$ (a ball) is i.i.d. uniformly randomly assigned to a state of Y (a bin). Only balls of the form $(x \in X \setminus S, e \in E)$ affect a bin's surplus. To bound surplus for bins, we characterize it as the sum of contributions from three types of balls from $(X \setminus S) \times E$, and restate this characterization from the proof outline:

Definition A.6 (Ball characterizations). We characterize three types of balls:

1. **Dense balls.** Consider a set L of states of X , where a state of X is in L if $P(X = x) \geq \frac{1}{\log^3(n)}$. Dense balls are all balls of the form $(x \in L, e \in E)$. We call these dense balls, because the low-entropy of E will prevent the collective mass of these balls from “expanding” well.
2. **Large balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $\geq \frac{\mathcal{T}}{2}$.
3. **Small balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $< \frac{\mathcal{T}}{2}$.

Bounding the harmful effects of dense balls. Recall that $\mathcal{T} = \frac{12}{n \log(n)}$. Now, we show how to bound the contribution of dense balls towards surplus. By our assumptions, $Y = f(X, E)$, and $H(E)$ is small, meaning there is not much randomness in our function. We defined L as states of X with probability at least $\frac{1}{\log^3(n)}$, so $|L| \leq \log^3(n)$. We would like to show that there are not too many bins where the dense balls contribute a significant amount to surplus. If $H(E) = 0$, this would be easy to show as then there would only be $|L| \leq \log^3(n)$ dense balls and thus they could only affect the surplus of $\log^3(n)$ bins. However, we aim to show this claim in the more general setting where $H(E) = o(\log(\log(n)))$. To accomplish this, we follow the same intuition to show that the limited entropy of E prevents this small number of states of X from greatly “spreading” to significantly affect a large number of states of Y . In particular, we show:

Lemma A.7 (Limited expansion). *Suppose Y can be written as a function $f(X, E)$ and $X \perp\!\!\!\perp E$. Consider any subset R of the support of X . For any subset T of the support of Y that satisfies $\forall t \in T : P(X \in R, Y = t) > \delta$, the cardinality of T is upper bounded as $|T| \leq \frac{H(E) + \log(|R|) + 2}{\delta \log(\frac{1}{\delta})}$.*

Proof. Consider a variable X' , whose distribution is obtained from the distribution of X by keeping only the states in R , and then normalized. More formally, for any $i \in R$, $P(X' = i) = \frac{P(X=i)}{P(X \in R)}$, and for any $i \notin R$, $P(X' = i) = 0$.

Recall $Y = f(X, E)$. Using the same f, E , we define $Y' = f(X', E)$. Note that $P(X \in R, Y = i) \leq P(Y' = i)$. If $P(X \in R, Y = i) \geq \delta$, then it must be true that $P(Y' = i) \geq \delta$. Moreover, this implies that if there exists such a subset T then $H(Y') \geq |T| \delta \log(\frac{1}{\delta}) - 2$ (note the negative two is from the fact that modifying a distribution by adding non-negative numbers to probabilities can decrease entropy by at most 2). Moreover, by data-processing inequality note that $H(Y') \leq H(X') + H(E|X') \leq H(X') + H(E) \leq \log(|R|) + H(E)$, where previous inequality is due to the fact that conditioning reduces entropy. This implies the desired inequality for the cardinality of set T . \square

To more directly use this for our goal, we present:

Corollary A.8. *There exist no subset $|T| = n/4$ such that $\forall t \in T : P(X \in L, Y = t) \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$*

Proof. Note that $|L| \leq \log^3(n)$. By Lemma A.7, any such T must satisfy:

$$|T| \tag{1}$$

$$\leq \frac{H(E) + \log(|L|) + 2}{1/(n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}) \cdot \log(n)} \tag{2}$$

$$\leq \frac{5 \log^2(\log(n)) \cdot n \cdot \log^{2c_{\text{close}}(n)}}{\log(n)} \tag{3}$$

$$\leq \frac{5 \log^2(\log(n)) \cdot n \cdot \log^{1/2}(n)}{\log(n)} \tag{4}$$

$$\leq \frac{n}{4} \tag{5}$$

We obtain Step 4 by previously setting $c_{\text{close}} = \frac{1}{4}$. We obtain Step 5 when n is sufficiently large such that $\frac{5 \log^2(\log(n))}{\log^{1/2}(n)} \leq \frac{1}{4}$. It can be shown that $n \geq 5$ is sufficient. \square

As a result, dense balls cannot significantly affect the surplus of many bins.

Bounding the harmful effects of large balls. We now show how large balls cannot significantly affect the surplus of too many bins, by showing there is a non-negligible number of bins that receive no large balls.

Lemma A.9 (Avoided big). *Given a balls-and-bins game with $c \cdot n \ln(n)$ balls mapped uniformly randomly to n bins, at least $\frac{n^{1-c}}{2}$ bins will receive no balls with high probability if c is a constant such that $0 < c \leq \frac{1}{3}$.*

Proof. This follows directly from (Wajc, 2017). By (Wajc, 2017), with high probability the number of empty bins will be $n^{1-c} \pm O(\sqrt{n \log(n)})$. For sufficiently large n , $O(\sqrt{n \log(n)}) \leq \frac{n^{2/3}}{2} \leq \frac{n^{1-c}}{2}$ and thus the number of empty bins is at least $\frac{n^{1-c}}{2}$ with high probability. \square

Note how this relates to the coupon collector's problem, where it is well-known that $\Theta(n \log(n))$ trials are necessary and sufficient to receive at least one copy of all coupons with high probability. This is analogous to the number of balls needed such that every bin has at least one ball. The result of Lemma A.9 is intuitive from the coupon collector's problem, because the number of trials needed concentrates very well. Meaning, with a constant-factor less number of trials than the expectation required, there are many coupons that have not yet been collected with high probability.

Corollary A.10. *As there are at most $\frac{1}{\tau/2} \leq \frac{n \log(n)}{6} \leq \frac{1}{4} \cdot n \ln(n)$ large balls, with high probability there are at least $\frac{n^{3/4}}{2}$ bins that receive no large balls.*

Bounding the harmful effects of small balls. For the small balls, we will also show that they cannot contribute too much surplus to too many states of Y . We will notably use that all small balls correspond to a state of X where $P(X = x) \leq \frac{1}{\log^3(n)}$. We will utilize this to show that most small balls are assigned to a state of Y that has not yet received $> \frac{\tau}{2}$ mass from its corresponding state of X , and accordingly would not increase the surplus. To accomplish this, we define a surplus quantity that only takes into account small balls:

Definition A.11 (Small ball surplus). We define the small ball surplus of a state y of Y as

$$z_y^{\text{small}} = \sum_{x \notin (S \cup L)} \max \left(\left(\sum_{\substack{e: \\ P(X=x, E=e) < \frac{\tau}{2}}} P(X = x, E = e, Y = y) \right) - \tau \right)$$

With this notion of surplus constrained to small balls, we show the following:

Lemma A.12 (Small ball limited surplus). *With high probability, there are at most $\frac{n}{4}$ values of i , i.e., number of bins, where $z_i^{\text{small}} \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$.*

Proof. We will consider all small balls in an arbitrary order. Let $x(t)$ be the corresponding state of X for the t -th small ball, $e(t)$ the corresponding state of E , and $w_{\text{ball}}(t)$ be the ball's probability mass (i.e., $P(X = x(t), E = e(t))$). Recall that for all small balls it must hold that $x(t) \notin L$ and thus $P(X = x(t)) < \frac{1}{\log^3(n)}$. We define the total small ball surplus as $Z^{\text{small}} = \sum_{y \in Y} z_y^{\text{small}}$. Now, we will consider all small balls in an arbitrary order and realize their corresponding entry of f to map them to a state of Y . Initially, we have not realized the entry of f for any balls and thus all $z_y^{\text{small}} = 0$ and $Z^{\text{small}} = 0$. As we map small balls to states of Y , we define $\Delta(t)$ as the increase of Z^{small} after mapping the t -th ball to a state of Y . By definition, $\sum_t \Delta(t)$ is equal to Z^{small} after all values of f have been completely realized.

Our primary intuition is that we will show for many small balls it holds that $\Delta(t) = 0$. As a result, we expect Z^{small} to not be very large.

As a result, we expect Z^{small} to not be very large. Let $y(t)$ be equal to $f(x(t), e(t))$, the state of Y that the t -th ball is mapped to. As f is realized for each config-

uration, let $w_Y^t(y, x)$ denote the total mass of balls assigned to state y of Y so far from state x of X , i.e., $w_Y^t(y', x') := \sum_{x', e: f(x', e)=y'} w_{\text{ball}}(t)$.

We upper-bound the expectation of $\Delta(t)$:

Claim 1. Regardless of the realizations of all $\Delta(t')$ for $t' < t$, it holds that $\Delta(t)$ is a random variable with values in range $[0, w_{\text{ball}}(t)]$ and $E[\Delta(t)] \leq \frac{w_{\text{ball}}(t)}{\log^2(n)}$.

Proof. The only conditions under which $\Delta(t)$ takes a positive value (which is upper-bounded by $w_{\text{ball}}(t)$), is when $w_Y^t(y(t), x(t)) > \frac{\mathcal{T}}{2}$ before the t -th ball is realized. Recall that $P(x(t)) \leq \frac{1}{\log^3(n)}$. Accordingly, the number of states y' of Y where $w_Y^t(y', x(t)) > \frac{\mathcal{T}}{2}$ is upper-bounded by $\frac{P(X=x(t))}{\mathcal{T}/2} \leq \frac{1/\log^3(n)}{6/(n \log(n))} = \frac{n \log(n)}{6 \log^3(n)} \leq \frac{n}{\log^2(n)}$. This is due to the fact that balls partition the total mass of $P(X = x(t))$ since we have $P(X = x(t)) = \sum_e P(X = x(t), E = e)$. This implies that the probability that the t -th ball will be mapped to a state y' of Y such that $w_Y^t(y', x(t))$ already exceeds the threshold of $\mathcal{T}/2$ (in other words where we might have $\Delta(t) > 0$) is upper-bounded by $\frac{n/\log^2(n)}{n} = \frac{1}{\log^2(n)}$ due to the fact that the function f is realized independently and uniformly randomly for each pair of (x, e) , i.e., for every distinct ball. Accordingly, $E[\Delta(t)] \leq \frac{w(t)}{\log^2(n)}$. \square

This enables us to upper-bound the sum of $\Delta(t)$:

Claim 2. $\sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$ with high probability.

Proof. We will transform $\Delta(t)$ into a martingale. In particular, we define $\Delta'(t) = \Delta'(t-1) + \Delta(t) - E[\Delta(t)|\Delta(1), \dots, \Delta(t-1)]$. We define $\Delta'(0) = 0$, and note that $\Delta'(c)$ is a martingale. By Azuma's inequality, we show $|\sum_t \Delta'(t)| \leq \frac{1}{8 \log(n)}$ with high probability:

$$\begin{aligned} P[|\Delta(t)| > \varepsilon] &< 2e^{-\frac{\varepsilon^2}{2 \sum c_i^2}} \\ &\leq 2e^{-\frac{\left(\frac{1}{8 \log(n)}\right)^2}{2(\max_i c_i) \cdot \sum c_i}} \\ &\leq 2e^{-\frac{\left(\frac{1}{8 \log(n)}\right)^2}{2 \times \mathcal{T}/2 \cdot 1}} \\ &= 2e^{-\frac{-n \log(n)}{12 \times 8 \times |\mathcal{V}| \times \log(n)}} \end{aligned}$$

Accordingly, by definition of $\Delta'(t)$ this implies $|\left(\sum_t \Delta(t)\right) - \sum_c E[\Delta(t)|\Delta(1), \dots, \Delta(c-1)]| \leq \frac{1}{8 \log(n)}$. By Claim 16 we know all $E[\Delta(t)|\Delta(1), \dots, \Delta(c-1)] \leq \frac{w_{\text{confing}(c)}}{\log^2(n)}$ and accordingly, $\sum_t E[\Delta(t)|\Delta(1), \dots, \Delta(c-1)] \leq \frac{1}{\log^2(n)}$. Together, these imply $\sum_t \Delta(t) \leq \frac{1}{8 \log(n)} + \frac{1}{\log^2(n)}$ with high probability, and for sufficiently large n it holds that

$\frac{1}{\log^2(n)} \leq \frac{1}{8 \log(n)}$. Thus, our high-probability on $|\Delta'(t)|$ implies that $\sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$ with high probability. \square

Finally, we conclude that our upper-bound on $\sum_t \Delta(t)$ implies an upper-bound on the number of states of Y with non-negligible small ball support:

Claim 3. If $\sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$, then there are at most $\frac{n}{4}$ bins where $z_i^{\text{small}} \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$.

Proof. $Z^{\text{small}} = \sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$. Given this upper-bound for total small ball surplus, we can immediately upper-bound the number of states of Y with small ball surplus greater than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ by the quantity $\frac{1/(4 \log(n))}{1/(n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)})} \leq \frac{n \cdot \log(\log(n)) \cdot \log^{1/2}(n)}{4 \log(n)} \leq \frac{n}{4}$. We obtain this by using $c_{\text{close}} = \frac{1}{4}$ and for sufficiently large n such that $\log(\log(n)) \leq \log^{1/2}(n)$. \square

Combining the three ball types: many bins with small surplus. Now, we combine all these intuitions to show there are many bins that have a small amount of surplus. We have shown that, with high probability, there are at most $n/4$ bins with non-negligible mass from dense balls by Corollary A.8, and at most $n/4$ bins with non-negligible mass from small balls Lemma B.18. Combining these sets, there are at most $n/2$ bins with non-negligible mass from dense balls or small balls. By Corollary B.16, with high probability at least $\frac{n^{3/4}}{2}$ bins will receive no large balls. Our goal is to show the intersection of the sets is large, so there are many bins that have small surplus.

Lemma A.13. *Let there be two sets $A, B \subseteq [n]$, where $|A| \geq \frac{n}{2}$ and A and B are both independently uniformly random subsets of size $|A|$ and $|B|$, respectively. It holds that $P(|A \cap B| \geq \frac{|B|}{4}) \geq 1 - 2e^{-\frac{|B|}{8}}$.*

Proof. To accomplish this, we will heavily utilize properties of negative association (NA). Lemma 8 of (Wajc, 2017) shows that permutation distributions are NA. Lemma 9 of (Wajc, 2017) shows closure properties of NA random variables. In particular, they show that concordant monotone functions defined on disjoint subsets of a set of NA variables are also NA. Accordingly, consider concordant monotone functions where each bin i has a random variable \mathcal{A}_i that takes value 1 if it is the first $|A|$ values of a permutation distribution and value 0 otherwise. These random variables are thus NA. Suppose we first realize the set B , independently of the realization of A . Then, a bin $y \in B$ would be in $A \cap B$

if $A_y = 1$. It is clear this formulation of the random process has a bijective mapping with the true random process, so $P(|A \cap B| \geq \frac{|B|}{4}) = P(\sum_{y \in B} \mathcal{A}_y \geq \frac{|B|}{4})$. By Theorem 5 of (Wajc, 2017), we can use Hoeffding's upper tail bound to show $P(\sum_{y \in B} \mathcal{A}_y < \frac{|B|}{4}) \leq P(|\sum_{y \in B} \mathcal{A}_y - E[\sum_{y \in B} \mathcal{A}_y]| > \frac{|B|}{4}) \leq 2e^{-\frac{|B|}{8}}$. \square

Corollary A.14. *With high probability, there are at least $\frac{n^{3/4}}{8}$ bins with surplus $z_y \leq \frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}$.*

Proof. We have defined three types of balls, and have proven results that show how there are many bins with negligible bad contribution for each type of ball. Now, we combine these with Lemma A.13 to show there are many bins where there is not much bad contribution in total. By Corollary A.8 there are at most $n/4$ bins with more than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}(n)}$ mass from dense balls. By Lemma B.18, there are at most $n/4$ bins with small ball surplus more than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}(n)}$. Let A be the set of bins with at most $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}(n)}$ mass from dense balls and at most $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}(n)}$ small ball surplus. By combining Corollary A.8 and Lemma B.18 we know $|A| \geq \frac{n}{2}$ with high probability. Let B be the set of bins that receive no big balls. By Corollary B.16, it holds that $|B| \geq \frac{n^{3/4}}{2}$ with high probability. By Lemma A.13, it holds that $|A \cap B| \geq \frac{n^{3/4}}{8}$ with failure probability at most $2e^{-\frac{2n^{3/4}}{16}}$. Moreover, all such bins will have total surplus at most $\frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$, because they receive no large balls and total surplus is then upper-bounded by the sum of small ball surplus and total mass from dense balls. \square

Existence of a small surplus bin with many plateau balls. Recall plateau balls, which are balls of $X \times E$ that take the form $(x \in S, E = e_1)$, where e_1 is the most probable state of E . We show that at least one of the bins with small surplus will receive many plateau balls with high probability:

Lemma A.15. *There exists a bin with surplus at most $\frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ and at least $\frac{\log(n)}{2 \log(\log(n))}$ plateau balls.*

Proof. Note that total surplus is independent of how plateau balls are mapped. Accordingly, we have determined a set of $\frac{n^{3/4}}{8}$ bins with small enough surplus. We aim to show that one of these bins receives a large number of plateau balls with high probability. We will rely on negative association (NA) in the balls-and-bins process to prove our result.

Claim 4. Indicator variables for if a bin receives some threshold of balls in a i.i.d. uniformly random balls-and-bins game are NA.

Proof. This follows immediately by using results of (Wajc, 2017). By Theorem 10 of (Wajc, 2017), the random variables of the number of balls assigned to each bin are NA. By Lemma 9 of (Wajc, 2017), concordant monotone functions define on disjoint subsets of a set of NA random variables are NA. Accordingly, if we have an indicator variable for whether a bin receives at least some number of balls, these indicator variables are NA. \square

Now, we lower-bound the expectation of these indicator variables:

Claim 5. Suppose cn balls ($c \leq 1$) are thrown i.i.d. uniformly randomly into n bins. The probability that a particular bin receives at least $k = \frac{d \log(n)}{\log(\log(n))}$ balls is at least $\frac{1}{en^d}$ given that $\frac{d}{c} \leq \log(\log(n))$.

Proof. We use the method outlined by (Cunha, 2011). We lower-bound the probability of a bin receiving at least k balls as follows:

$$\begin{aligned} \binom{cn}{k} \cdot \left(\frac{1}{n}\right)^k \cdot \left(1 - \frac{1}{n}\right)^{cn-k} &\geq \left(\frac{cn}{k}\right)^k \cdot \frac{1}{n^k} \cdot \frac{1}{e} \\ &\geq \frac{1}{e} \cdot \left(\frac{c}{k}\right)^k \\ &= \frac{1}{e} \cdot \left(\frac{c \log(\log(n))}{d \log(n)}\right)^{\log_{\log(n)}(n^d)} \\ &\geq \frac{1}{e} \cdot \left(\frac{1}{\log(n)}\right)^{\log_{\log(n)}(n^d)} \\ &= \frac{1}{en^d} \end{aligned} \quad (6)$$

We obtain Step 6 by using $\frac{d}{c} \leq \log(\log(n))$. \square

By Lemma A.3 there are at least $\frac{c_{\text{close}} c_{\text{support}}}{6} \cdot n = \frac{c_{\text{support}}}{24} \cdot n$ plateau balls. Now, consider NA indicator variables \mathcal{B}_i for whether or not a particular bin receives at least $\frac{\log(n)}{2 \log(\log(n))}$ plateau balls. By Claim 4, these indicator variables are NA. By Claim 5, it holds that $E[\mathcal{B}_i] \geq \frac{1}{en^{0.5}}$ for sufficiently large n where $\frac{1/2}{c_{\text{support}}/24} = \frac{12}{c_{\text{support}}} \leq \log(\log(n))$. Finally, we can upper-bound the probability that $\mathcal{B}_i = 0$ for all bins with small enough surplus, of which there are at least $\frac{n^{3/4}}{8}$. Using marginal probability bounds for NA variables shown in Corollary 3 of (Wajc, 2017), all such $\mathcal{B}_i = 0$ with probability at most $\left(\frac{1}{en^{0.5}}\right)^{\frac{n^{3/4}}{8}}$. \square

Proving large conditional entropy. Finally, we show how the existence of a bin with small surplus and many plateau balls implies that the bin has large conditional entropy:

Lemma A.16 (High-entropy conditional). *Given a bin y' that has $z_{y'} \leq \frac{2}{n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}}$, and receives $\frac{\log(n)}{2 \log(\log(n))}$ plateau balls, then $H(X|Y = y') = \Omega(\log(\log(n)))$.*

Proof. To show $H(X|Y = y')$ is large, we first define the vector v such that $v(x) = P(X = x, Y = y')$. Similarly, we define $\bar{v}(x) = \frac{v}{P(Y=y')}$, meaning $\bar{v}(x) = P(X = x|Y = y')$ and $|\bar{v}|_1 = 1$. Our underlying goal is to show $H(\bar{v})$ is large. To accomplish this, we will split the probability mass of v into three different vectors $v_{\text{initial}}, v_{\text{plateau}}, v_{\text{surplus}}$ such that $v = v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}$. The entries of v_{plateau} will correspond to mass from plateau states of X , v_{initial} will correspond to the first \mathcal{T} mass from non-plateau states of X , and v_{surplus} will correspond to mass that contributes to the surplus $z_{y'}$. We more formally define the three vectors as follows:

- v_{plateau} . The vector of probability mass from plateau states of X . $v_{\text{plateau}}(x)$ is 0 if $x \notin S$ and $v_{\text{plateau}}(x) = P(X = x, Y = y')$ if $x \in S$.
- v_{initial} . For non-plateau states of X , their first \mathcal{T} probability mass belongs to v_{initial} . $v_{\text{initial}}(x) = \min(P(X = x, Y = y'), \mathcal{T})$ if $x \notin S$ and $v_{\text{initial}}(x) = 0$ otherwise.
- v_{surplus} . For non-plateau states of X , their probability mass beyond the first \mathcal{T} mass belongs to v_{surplus} . This corresponds to the surplus quantity. $v_{\text{surplus}}(x) = \max(0, P(X = x, Y = y') - \mathcal{T})$ if $x \notin S$ and $v_{\text{surplus}}(x) = 0$ otherwise. By this definition, $z_{y'} = |v_{\text{surplus}}|_1$.

To show $H(X|Y = y') = H(\bar{v})$ is large, we divide our approach into two steps:

1. Show there is substantial helpful mass: $|v_{\text{initial}} + v_{\text{plateau}}|_1 = \Omega\left(\frac{1}{n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}}(n)}\right)$
2. Show the distribution of helpful mass has high entropy: $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) = \Omega(\log(\log(n)))$.
3. Show that, even after adding the hurtful mass, the conditional entropy is large: $H(X|Y = y') = H(\bar{v}) \geq H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - O(1) = \Omega(\log(\log(n)))$

In the first step, we are showing that the distribution when focusing on just the helpful mass of $v_{\text{initial}}, v_{\text{plateau}}$ has high a substantial amount of probability mass. In the second step, we prove how this distribution of helpful mass has high entropy. In the third step, we show that the hurtful mass of v_{surplus} does not decrease entropy more than a constant.

First, we show that there is a substantial amount of helpful mass:

Claim 6. $|v_{\text{initial}} + v_{\text{plateau}}|_1 = \frac{1}{2c_{\text{lb}}n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}}(n)}$

Proof. Recall that the bin y' received $\frac{\log(n)}{2 \log(\log(n))}$ plateau balls. As defined in Lemma A.3, the set S of plateau

states is defined such that $\frac{\max_{x \in S} P(X=x)}{\min_{x \in S} P(X=x)} \leq \log^{c_{\text{close}}}(n)$ and $\min_{x \in S} P(X = x) \geq \frac{1}{c_{\text{lb}}n \log(n)}$. Also recall that by Lemma A.4 the most probably state of E has large probability. In particular, $P(E = e_1) \geq \frac{1}{\log^{c_{\text{close}}}(n)}$. Let the subset $S' \subseteq S$ be the subset of plateau states of X such that their plateau ball is mapped to y' . In particular, for every $x \in S'$ it holds that $f(x, e_1) = y'$. Accordingly, $P(X = x, Y = y') \geq P(X = x) \cdot P(E = e_1)$ for $x \in S'$. Thus, the total weight from plateau states of X is at least $|S'| \cdot \min_{x \in S'} P(X = x) \cdot P(E = e_1) \geq |S'| \cdot \frac{\max_{x \in S'} P(X=x)}{\log^{c_{\text{close}}}(n)}$. $P(E = e_1) \geq \frac{1}{2c_{\text{lb}}n \log(\log(n)) \log^{2c_{\text{close}}}(n)}$. \square

Next, we show the distribution of helpful mass has high entropy:

Claim 7. $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) \geq \frac{\log(\log(n))}{4}$

Proof. Let us define $\bar{v}_{\text{helpful}} = \frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}$ to be the vector of helpful mass, and we will show $H(\bar{v}_{\text{helpful}})$ is large by upper-bounding $\max_x \bar{v}_{\text{helpful}}(x)$.

For non-plateau states of X , it follows from Claim 20 that $\max_{x \notin S} \bar{v}_{\text{helpful}}(x) \leq \frac{\mathcal{T}}{|v_{\text{initial}} + v_{\text{plateau}}|_1} \leq \frac{\mathcal{T}}{2c_{\text{lb}}n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}}(n)} = \frac{\mathcal{T}}{24c_{\text{lb}} \log(\log(n)) \cdot \log^{2c_{\text{close}}}(n)}$.

For plateau states of X , in Claim 20 we also developed the lower-bound of $|v_{\text{initial}} + v_{\text{plateau}}|_1 \geq |S'| \cdot \frac{\max_{x \in S'} P(X=x)}{\log^{c_{\text{close}}}(n)}$. $P(E = e_1) \geq \frac{\log(n) \cdot \max_{x \in S'} P(X=x)}{2 \log^{2c_{\text{close}}}(n) \log(\log(n))}$. Accordingly, we can upper-bound $\max_{x \in S'} \bar{v}_{\text{helpful}}(x) \leq \frac{\max_{x \in S'} P(X=x)}{|v_{\text{initial}} + v_{\text{plateau}}|_1} \leq \frac{2 \log(\log(n)) \log^{2c_{\text{close}}}(n)}{\log(n)}$.

Accordingly, we can lower-bound the entropy of $H(\bar{v}_{\text{helpful}}) = \sum_x \bar{v}_{\text{helpful}}(x) \cdot \log\left(\frac{1}{\bar{v}_{\text{helpful}}(x)}\right) \geq \sum_x \bar{v}_{\text{helpful}}(x) \cdot \log\left(\frac{1}{\max_{x'} \bar{v}_{\text{helpful}}(x')}\right) = \log\left(\frac{1}{\max_{x'} \bar{v}_{\text{helpful}}(x')}\right) \geq \log\left(\frac{24c_{\text{lb}} \log(n)}{\log^{2c_{\text{close}}}(n) \log(\log(n))}\right) = (1 - 2c_{\text{close}}) \log(\log(n)) - \log(\log(\log(n))) - \log(24c_{\text{lb}}) = \frac{\log(\log(n))}{2} - \log(\log(\log(n))) - \log(24c_{\text{lb}}) \geq \frac{\log(\log(n))}{4}$ for sufficiently large n where $\frac{\log(\log(n))}{2} \geq \log(\log(\log(n))) + \log(24c_{\text{lb}})$. \square

Finally, we show the hurtful mass does not decrease entropy much, and thus our conditional distribution has high entropy:

Claim 8. $H(X|Y = y') = H(\bar{v}) \geq \Omega(1) \cdot H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - O(1) = \Omega(\log(\log(n)))$

Proof. We lower-bound $H(\bar{v})$ with the main intuitions that $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) = \Omega(\log(\log(n)))$ and

$\frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} = \Omega(1)$. We more precisely obtain this lower-bound for $H(\bar{v})$ as follows:

$$\begin{aligned}
 H(\bar{v}) &= H\left(\frac{v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}\right) \\
 &= \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x) + v_{\text{surplus}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\
 &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x) + v_{\text{surplus}}(x)} \\
 &\geq \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\
 &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x)} - 2 \tag{7} \\
 &\geq \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\
 &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x)} - 2 \\
 &= \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) \\
 &\quad - 2 \\
 &= \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}}|_1 + z_{y'}} H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - 2 \\
 &\geq \frac{1}{1 + 2c_{\text{lb}}} \cdot H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - 2 \tag{8} \\
 &= \Omega(\log(\log(n))) \tag{9}
 \end{aligned}$$

To obtain Step 10, we note that all summands are manipulated from the form $\sum_x p_x \log(\frac{1}{p_x})$ to $\sum_x p'_x \log(\frac{1}{p'_x})$ where $p'_x \leq p_x$ for all x . As the derivative of $p \log(\frac{1}{p})$ is non-negative for $0 \leq p \leq \frac{1}{e}$, the value of at most two summands can decrease, and they can each decrease by at most one. To obtain Step 11, we use Claim 20. To obtain Step 12, we use Claim 22. \square

Thus, we have shown $H(X|Y = y') = \Omega(\log(\log(n)))$. \square

Corollary A.17. *Under our assumptions, $H(X|Y = y') = \Omega(\log(\log(n)))$ and thus $H(\tilde{E}) = \Omega(\log(\log(n)))$.*

B. Proof of Theorem 5.4

B.1. Proof Outline

For much of this proof, we follow intuitions and use terminology from the proof of Theorem 4.1. Consider a pair of variables X and Y such that X is a source and there is a path from X to Y . We aim to show that $\text{MEC}(Y|X) < \text{MEC}(X|Y)$. It is simple for us to show that $\text{MEC}(Y|X) = o(\log(\log(n)))$. To show $\text{MEC}(X|Y) =$

$\Omega(\log(\log(n)))$, we will use an approach similar to Theorem 4.1 in that we will show existence of a state y' of Y such that $H(X|Y = y')$ is large. For showing there is a large $H(X|Y = y')$, we will show that there is a y' where its surplus is small and it receives many plateau balls. While we can factor Y as a function of X and small-entropy E (i.e., $Y = f(X, E)$), this is *not* a uniformly random function so we cannot simply apply the result of Theorem 4.1. In fact, a key difficulty is that this graph setting with more than two variables results in correlations between mappings. For example, in a graph such as the line graph (Figure 1(a)) with each node being a uniformly random deterministic function of its parents, one can show that conditioning on $f_Y(f_{X_2}(X = x)) = y$ almost doubles the probability that $f_Y(f_{X_2}(X = x')) = y$. Our new proof method must be able to withstand the dependencies that are introduced by this setting.

To provide some intuition, we give a very high-level overview for how to show existence of a large $H(X_{\text{src}}|Y = y')$ for two particular graphs, and we then expand to generalize these intuitions.

First, we consider the line graph. For simplicity, suppose that all nodes are deterministic functions of their parents (i.e., all $H(E_i) = 0$). Using the method from Theorem 4.1, we can see that there exists a large $H(X_{\text{src}}|X_2 = x'_2)$. This is because we can show there is a bin of X_2 that has small surplus and receives $\Omega(\frac{\log(n)}{\log(\log(n))})$ balls. However, this analysis is actually loose in a sense. For a c where $0 < c < 1$, we can actually show there are n^c such bins that have small surplus and receive $\Omega(\frac{\log(n)}{\log(\log(n))})$ plateau balls. Now, when we look at how X_2 is mapped to Y , each of the bins of X_2 will “stick together.” More formally, each bin of X_2 will have all of its mass mapped together to a uniformly random state of Y . This is because it is a deterministic function, but our proof will utilize a similar idea for when the function is not deterministic but the entropy is still small. It is then our hope that a good fraction of the bins with our desired properties (small surplus and many plateau balls) at X_2 , will be mapped to a state of Y that does not have much surplus. In this sense, we have “heavy bins” and a non-negligible proportion of them are “surviving” from one node to the next because they aren’t mapped to a bin with too much surplus. Through careful analysis, we are able to show that at least one such bin survives to the node of Y , and thus $H(X|Y = y')$ is large. This proof method would hold if we extend this line graph to any constant length.

Second, we consider the diamond graph (Figure 1(b)). Again, we assume all functions are deterministic for simplicity. Recall that for the line graph, our proof method was to show that there were many heavy bins at X_2 , and then some heavy bins kept “sticking together” and “surviving” until we reached Y . This was because if two states

of X were mapped to the same state of X_2 , then they would “stick together” and would always be mapped to the same state for later nodes (e.g. if $f_{X_2}(x) = f_{X_2}(x')$ then $f_{X_3}(f_{X_2}(x)) = f_{X_3}(f_{X_2}(x'))$). However, this is very far from what is happening in diamond graph. In diamond graph, observe that $Y = f_Y(X_2, X_3)$. By definition of our graph, X_2 and X_3 are independent deterministic functions of X . Two states x and x' of X will be mapped to Y independently unless both $f_{X_2}(x) = f_{X_2}(x')$ and $f_{X_3}(x) = f_{X_3}(x')$. As these are independent, the probability of this happening is $\frac{1}{n^2}$. Thus, the expected number of pairs that are not mapped to Y independently of each other is $\binom{n}{2} \times \frac{1}{n^2} < \frac{1}{2}$. Accordingly, essentially all states of X will be mapped to a state of Y i.i.d. uniformly randomly. This enables us to more directly use the result and techniques of Theorem 4.1 and treat X and Y as a bivariate problem.

While we are able to show how both of these graphs will result in a large $H(X_{\text{src}}|Y = y')$, we do so very differently. For the line graph we show that there are bins with the properties we desire (small surplus and many plateau balls), that they will “stick together” as we move down through the graph, and at least one will “survive” to Y and thus $H(X_{\text{src}}|Y = y')$. For the diamond graph we show that when we get to Y , almost everything will be mapped independently randomly again, and that we can more directly use our bivariate techniques. There is a strong sense in which these two proof methods are opposites of each other (utilizing probability mass staying together throughout the graph as opposed to being independent at the end), yet we would like one unified approach for handling general graphs. To accomplish this, we introduce the *Random Function Graph Decomposition* to combine intuitions of these two settings into a characterization for all graphs.

Definition B.1 (Random Function Graph Decomposition). For the Random Function Graph Decomposition we specify a source X and a node Y such that there is a path from X to Y . We ignore all nodes not along a path from X to Y . We define the remaining nodes as the set V_{decomp} . Then, we consider the nodes of V_{decomp} an arbitrary valid topological ordering and color each node as follows:

- If X is a parent of the node, or if the node has multiple parents and they are not all the same color, we create a new color for this node.
- Otherwise, all of the node’s parent(s) have the same color, and this node will inherit said color.

At a high-level, when a new color is created for a node, then everything is being mapped to the node almost-independently (similar to the intuition of the diamond graph). When a node inherits its color, there is a sense in which

things “stick together” (similar to the intuition of the line graph). Let $\text{color-root}(Y)$ be the earliest node in any topological ordering that has the same color as Y in the Random Function Graph Decomposition (it can be shown that $\text{color-root}(Y)$ is unique). We aim to use the Random Function Graph Decomposition to show that everything will be mapped to $\text{color-root}(Y)$ mostly independently. This will result in there being some bins with our desired properties (small surplus, many plateau balls) at $\text{color-root}(Y)$. Then, we will show that at least one of these bins survives throughout all bins with the same color from $\text{color-root}(Y)$ to Y , implying existence of a large $H(X_{\text{src}}|Y = y')$.

In particular, to show that balls are mapped to $\text{color-root}(Y)$ mostly independently, we introduce the notion of *related* mass. More concretely, we define $\text{related}_1(x)$ mass as the amount of mass of balls that are ever mapped to the same state as x among any variable. We define $\text{related}_2(x)$ mass as the amount of mass of balls that are mapped to the same state as x for variables of at least two distinct colors in the Random Function Graph Decomposition. Inductively, we will show there are $\Omega(n)$ plateau balls such that $\text{related}_1(x) = O(\frac{1}{n})$ and $\text{related}_2(x) = O(\frac{1}{n^2})$. Moreover, we show that the quantity $\text{related}_2(x)$ upper-bounds mass that can have some dependence with x in how it is mapped to $\text{color-root}(Y)$. With this upper-bound on dependence, we are able to use techniques of Theorem 4.1 to show there are many bins of $\text{color-root}(Y)$ with many plateau balls and not much surplus. Finally, we show that, within the color of $\text{color-root}(Y)$ and Y , at least one of these bins “survives” to Y and accordingly $\text{MEC}(X_{\text{src}}|Y)$ is large.

B.2. Complete Proof

We must show that for a source X_{src} and a node Y such that there is a path from X_{src} to Y , $\text{MEC}(X_{\text{src}}|Y) > \text{MEC}(Y|X_{\text{src}})$.

Upper bounding $\text{MEC}(Y|X_{\text{src}})$. It is simple to show that $\text{MEC}(Y|X_{\text{src}})$ is small:

Claim 9. $\text{MEC}(Y|X_{\text{src}}) \leq o(|V| \log(\log(n)))$

Proof. Y can be written as a function of X_{src} and the set of all E_i excluding $E_{X_{\text{src}}}$. As X_{src} is independent of these E_i , and their total entropy is $\sum_i H(E_i) = o(|V| \log(\log(n)))$, the claim holds since $|V| = \mathcal{O}(1)$. \square

Bounding $\text{MEC}(X_{\text{src}}|Y)$ via $H(X_{\text{src}}|Y = y)$. Our method for lower-bounding $\text{MEC}(X_{\text{src}}|Y)$ is substantially more involved. As in Theorem 4.1, we will lower-bound it by $\text{MEC}(X_{\text{src}}|Y) \geq \max_y H(X_{\text{src}}|Y = y)$ (see Theorem 4.1 for proof). Our proof aims to show there is a conditional entropy such that $\max_y H(X_{\text{src}}|Y = y) = \Omega(\log(\log(n)))$.

Showing existence of a near-uniform plateau. A key step in our approach, as in the proof of Theorem 4.1, is that we

will find a subset of the support of X whose probabilities are multiplicative close to one another. In particular, we will find a subset of X_{src} where their probabilities are within a factor of $\log^{c_{\text{close}}}(n)$ of each other, where $0 < c_{\text{close}} < 1$. For our analysis, we require a value of c_{close} that is $\Omega(1)$ yet below some threshold. While there are multiple values of c_{close} that satisfy this condition, we will use $c_{\text{close}} = 1/4$. This set of states of X_{src} that are multiplicatively close to one another will be called the *plateau* of X_{src} . We use Lemma A.3 proven in Theorem 4.1 to show how the $(\Omega(n), \Omega(\frac{1}{n \log(n)}))$ -support assumptions implies a plateau of states of X :

Lemma B.2 (Plateau existence). *Suppose X has $(c_{\text{support}}n, \frac{1}{c_{\text{lb}}n \log(n)})$ -support for constants $0 < c_{\text{support}} \leq 1$ and $c_{\text{lb}} \geq 1$. Additionally, assume n is sufficiently large such that $\frac{\log(2c_{\text{lb}}/c_{\text{support}})}{\log(\log(n))} \leq 1$. Then, there exists a subset $S \subseteq [n]$ of the support of X , such that the following three statements hold:*

1. $\frac{\max_{i \in S} P(X=i)}{\min_{i \in S} P(X=i)} \leq \log^{c_{\text{close}}}(n)$
2. $\min_{i \in S} P(X = i) \geq \frac{1}{c_{\text{lb}}n \log(n)}$
3. $|S| \geq \frac{c_{\text{close}}c_{\text{support}}n}{6}$, for any $0 < c_{\text{close}} < 1$.

Characterization as a balls-and-bins game. Our proof method of Theorem 4.1 characterizes a balls-and-bins game where states of $X \times E$ are balls and states of Y are bins. As we realized an entry $f(x, e)$ as a uniformly random state of Y , we characterized this as a ball (a state of $X \times E$) being assigned to a uniformly random bin (a state of Y). In the graph setting of this theorem, such a characterization is more complicated. Any node X_i is a uniformly random function of $\text{Pa}(X_i)$ and E_i . We define E^* to be the Cartesian product of all E_i other than E_X . Using this, we characterize balls as being states of $X \times E^*$. Note how any random variable in our SCM is a deterministic function of $X \times E^*$. In particular, it is the composition of (potentially many) f_i terms. For simplicity of notation, we let $f_T^*(x \times e^*)$ denote the value of a set of variables T for a particular state of $x \times e^*$. In the characterization of our balls-and-bins game, all balls with the same configuration of $\text{Pa}(X_i)$ and E_i are mapped uniformly randomly together to a state of X_i . In other words, configurations are realized i.i.d. uniformly randomly. Using our notation, this means two balls (x_a, e_a^*) and (x_b, e_b^*) are mapped independently to variable X_i if and only if $f_{\text{Pa}(X_i) \cup E_i}^*(x_a, e_a^*) \neq f_{\text{Pa}(X_i) \cup E_i}^*(x_b, e_b^*)$.

Lower-bounding the most probable state of E_i and E^* . We focus first on *plateau balls*, which are balls corresponding to states of S (the set of plateau states of X) and the highest probability state of E^* . In particular, they are balls of the form $(X \in S, E^* = e_1^*)$ where e_1^* is the most probable state

of E^* . To show that these plateau balls have enough probability mass to be helpful, we first use Lemma A.4 proven in Theorem 4.1 that implies all $\max_e H(E_i = e) \geq \frac{1}{\log^{c_{\text{close}}}(n)}$:

Lemma B.3. *If $H(E) \leq c_{\text{close}} \log(\log(n))$ then $P(E = e_1) \geq \frac{1}{\log^{c_{\text{close}}}(n)}$*

This implies a lower-bound on the probability $P(E^* = e_1^*)$:

Lemma B.4. *If all $\max_e P(E_i = e) \geq \frac{1}{\log^{c_{\text{close}}}(n)}$, then $P(E^* = e_1^*) \geq \frac{1}{\log^{c_{\text{close}}|V|-1}(n)}$.*

Proof. As E^* is the Cartesian product of $|V| - 1$ variables E_i , it holds that $\max P(E^* = e^*) \geq (\min_i \max_e P(E_i = e))^{|V|-1} \geq (\frac{1}{\log^{c_{\text{close}}}(n)})^{|V|-1} \geq \frac{1}{\log^{c_{\text{close}}|V|-1}(n)}$. \square

Introducing surplus. In Theorem 4.1, we prove how there exists a bin that receives a large amount of mass that helps the bin have large conditional entropy (such helpful mass includes the plateau balls), and not much mass that hurts the conditional entropy making it small. To formalize this hurtful mass, we introduced the *surplus* quantity described in Definition A.1. This surplus is a way of quantifying the probability mass received by a state of Y that is hurtful towards making the conditional entropy large. The proof of Theorem 4.1 achieves a lower-bound for $\max_y H(X|Y = y)$ by proving existence of a state y' of Y where y' receives many plateau balls and the surplus is small. Likewise, we will also prove existence of such a state of Y with many plateau balls and small surplus, in the graph setting. We formalize the notion of surplus as follows:

Definition B.5 (Surplus, $\mathcal{T} = \frac{120}{n \log(n)}$). We define the surplus of a state i of Y as $z_i = \sum_{j \notin S} \max(0, P(X = j, Y = i) - \frac{120}{n \log(n)})$.

Introducing the Random Function Graph Decomposition. In Appendix B.1, we introduced intuitions from considering the diamond graph in Figure 1(b) and the line graph in Figure 1(a). In the proof outline for a diamond graph, we utilized the intuition that almost all balls were independently assigned to Y . This enables us to use techniques from Theorem 4.1, as almost all balls were independently assigned to a uniformly random state of Y , closely mirroring the setting of Theorem 4.1. In the proof outline for line graph, we used techniques of Theorem 4.1 to show that there would be many bins that received many plateau balls and small surplus. Then, we showed that at least one of these bins would mostly “survive” and remain in-tact to Y . While our intuitions for both of these graphs enabled us to show existence of a large $H(X_{\text{src}}|Y = y)$, but they did so with near-opposite methods. Our intuition for the diamond graph exploits independence (everything is assigned almost independently to Y), while our intuition for the line graph exploits dependence (some bins with our desired properties

“survive” from the second node onwards). We introduce the *Random Function Graph Decomposition* to combine intuitions of these two graphs into a characterization for all graphs:

Definition B.6 (Random Function Graph Decomposition). For the Random Function Graph Decomposition we specify a source X and a node Y such that there is a path from X to Y . We ignore all nodes not along a path from X to Y . We define the remaining nodes as the set V_{decomp} . Then, we consider the nodes of V_{decomp} an arbitrary valid topological ordering and color each node as follows:

- If X is a parent of the node, or if the node has multiple parents and they are not all the same color, we create a new color for this node.
- Otherwise, all of the node’s parent(s) have the same color, and this node will inherit said color.

At a high-level, when a new color is created for a node, then we will see that plateau balls are being mapped to the node almost-independently (similar to the intuition of the diamond graph). When a node inherits its color, there is a sense in which things “stick together” (similar to the intuition of the line graph). For some node $X_i \in V_{\text{decomp}}$, we define $\text{color}(X_i)$ to be the node’s color in the Random Function Graph Decomposition. Under a fixed topological ordering, let $\text{color-root}(Y)$ be the earliest node that has the same color as Y in the Random Function Graph Decomposition (it can be shown that $\text{color-root}(Y)$ is unique). We aim to use the Random Function Graph Decomposition to show that everything will be mapped to $\text{color-root}(Y)$ mostly independently. This will result in there being some bins with our desired properties (small surplus, many plateau balls) at $\text{color-root}(Y)$. Then, we will show that at least one of these bins survives throughout all bins with the same color from $\text{color-root}(Y)$ to Y , implying existence of a large $H(X_{\text{src}}|Y = y')$.

Introducing related mass. To show how plateau balls are mapped to $\text{color-root}(Y)$ mostly independently, we introduce the concept of *related* mass. Related mass introduces a measure of how much mass has come into contact with a particular plateau ball:

Definition B.7 (Related mass). We define related mass of two types as follows.

- For a plateau state x of X , we define $\text{related}_1(x)$ mass as the amount of mass of balls from non-plateau states of X that are ever mapped to the same state as the plateau ball of x among any variable in the Random Function Graph Decomposition. In other words, x', e^* contributes to $\text{related}_1(x)$ if it satisfies the following for some X_i : x' together with some realization e^* contributes to the same bin of X_i that x is mapped to together with e_1^* . More formally, we define $\mathcal{B}_1(x)$ as the

set of balls whose mass counts towards $\text{related}_1(x)$, where $\mathcal{B}_1(x) = \{x' \in X \setminus S, e^* \in E^* | \exists X_i \in V_{\text{decomp}} \text{ s.t. } f_{X_i}^*(x', e^*) = f_{X_i}^*(x, e_1^*)\}$. Accordingly, $\text{related}_1(x) = \sum_{x', e^* \in \mathcal{B}_1(x)} P(X = x') \cdot P(E^* = e^*)$.

- For a plateau state x of X , we define $\text{related}_2(x)$ mass as the amount of mass of balls from non-plateau states of X that are ever mapped to the same state as the plateau ball of x among variables of at least two distinct colors in the Random Function Graph Decomposition. In other words, x', e^* contributes to $\text{related}_2(x)$ if it satisfies the following for some X_i, X_j with distinct colors: x' together with some realization e^* contributes to the same bin of X_i that x is mapped to together with e_1^* ; same holds for X_j . More formally, we define $\mathcal{B}_2(x)$ as the set of balls whose mass counts towards $\text{related}_2(x)$, where $\mathcal{B}_2(x) = \{x' \in X \setminus S, e^* \in E^* | \exists X_i, X_j \in V_{\text{decomp}} \text{ s.t. } f_{X_i}^*(x', e^*) = f_{X_i}^*(x, e_1^*), f_{X_j}^*(x', e^*) = f_{X_j}^*(x, e_1^*), \text{color}(X_i) \neq \text{color}(X_j)\}$. Accordingly, $\text{related}_2(x) = \sum_{x', e^* \in \mathcal{B}_2(x)} P(X = x') \cdot P(E^* = e^*)$.

Now, we will consider an arbitrary topological ordering of V_{decomp} . In this ordering, we define $\text{order}(X_i)$ for $X_i \in V_{\text{decomp}}$ as the index of X_i in the topological ordering. We introduce a modification of $\text{related}_1(x)$ where $\text{related}_1^{\text{order}(X_i)}(x)$ only considers nodes of V_{decomp} that are strictly earlier in the topological ordering than X_i . We define $\text{related}_2^{\text{order}(X_i)}(x)$ analogously. It is our goal to show that there are many plateau states $x \in S$ such that $\text{related}_2^{\text{order}(\text{color-root}(Y))}(x)$ is small. This will enable us to show how there are many plateau balls that are mapped to $\text{color-root}(Y)$ independently of almost all other mass.

Upper-bounding related mass. To show independence in how some plateau balls are mapped to $\text{color-root}(Y)$, we bound $\text{related}_2^{\text{order}(\text{color-root}(Y))}(x)$ for some plateau states $x \in S$.

To show this, we will process nodes in the topological ordering. After processing the first i nodes, we will argue that there is a large set S_{indep}^i with upper-bounds on all $\text{related}_1^i(x)$ and $\text{related}_2^i(x)$.

Lemma B.8. *With high probability, after processing the first i nodes in the topological ordering, there exists a set S_{indep}^i such that $|S_{\text{indep}}^i| = \frac{|S|}{6^i}$, all $x \in S_{\text{indep}}^i$ satisfy $\text{related}_1^i(x) \leq \frac{6i}{n}$ and $\text{related}_2^i(x) \leq \frac{18 \times i \times (i-1)}{n^2}$, and all $x, x' \in S_{\text{indep}}^i$ satisfy $f_{X_j}^*(x, e_1^*) \neq f_{X_j}^*(x', e_1^*)$ for all $1 \leq j \leq i$.*

Proof. We begin with the following claims.

Claim 10. Lemma B.8 holds for $i = 1$.

Proof. To find a subset of S to be S_{indep}^i , we will choose any arbitrary subset of size $\frac{S}{6}$. By definition of related_1 and related_2 , all plateau balls are different states of X_{src} so $\text{related}_1^1(x) = \text{related}_2^1(x) = 0$ for every $x \in S$, and $f_{X_{\text{src}}}^*(x, e_1^*) \neq f_{X_{\text{src}}}^*(x', e_1^*)$ for all $x, x' \in S$. \square

Claim 11. Lemma B.8 holds for i if it holds for all $j < i$.

Proof. First, we realize f_{X_i} for all cells other than those corresponding to configurations of $\text{Pa}(X_i) \cup E_{X_i}$ that contain an element of S_{indep}^{i-1} . Now, we consider the process of realizing the entries of f_{X_i} corresponding to elements of S_{indep}^{i-1} in an arbitrary order. We define a random variable for every element of S_{indep}^{i-1} . For the j -th element, we define \mathcal{S}_j as follows:

- If the element is mapped to a bin that another element of S_{indep}^{i-1} has been mapped to, then $\mathcal{S}_j = -1$.
- Otherwise, if the element $x \in S_{\text{indep}}^{i-1}$ is mapped to a bin that contains total mass at least $\frac{6}{n}$, or total mass from $\mathcal{B}_1^{i-1}(x)$ of at least $\frac{6\text{related}_1^{i-1}(x)}{n}$, then $\mathcal{S}_j = 0$.
- Else, then $\mathcal{S}_j = 1$.

The intuition behind \mathcal{S}_j is that we will count an element of $x \in S_{\text{indep}}^i$ as being eligible for S_{indep}^i if it lands in a bin with no other value of S_{indep}^{i-1} , and if it lands in a bin that will not increase $\text{related}_1^i(x)$ or $\text{related}_2^i(x)$ by too much.

Claim 12. Consider the set comprised of each element $x \in S_{\text{indep}}^{i-1}$ that satisfies the following. Suppose x is assigned to a bin such that before x is mapped to the bin, the bin has total mass at most $\frac{6}{n}$ and total mass intersecting from $\mathcal{B}_1^{i-1}(x)$ of at most $\frac{6\text{related}_1^{i-1}(x)}{n}$. Moreover, suppose x is the only element of S_{indep}^{i-1} that is ever assigned to this bin. Then, this set of all such x would meet the desired properties required of S_{indep}^i .

Proof. The increase of the quantity $\text{related}_1^i(x)$ is bounded by the amount of other mass in the bin that x is assigned to. Accordingly, $\text{related}_1^i(x) \leq \text{related}_1^{i-1}(x) + \frac{6}{n} \leq \frac{6 \times (i-1)}{n} + \frac{6}{n} = \frac{6i}{n}$. The increase of the quantity $\text{related}_2^i(x)$ is bounded by the amount of mass from $\mathcal{B}_1^{i-1}(x)$ in the bin x is assigned to. Accordingly, $\text{related}_2^i(x) \leq \text{related}_2^{i-1}(x) + \frac{6\text{related}_1^{i-1}(x)}{n} \leq \frac{18 \times (i-1) \times (i-2)}{n^2} + \frac{36(i-1)}{n^2} = \frac{18 \times i \times (i-1)}{n^2}$. \square

Moreover, we claim that $\sum \mathcal{S}_i$ serves as a lower bound for the set of elements eligible for S_{indep}^i referenced in Claim 12.

Claim 13. The number of elements of S_{indep}^{i-1} that are eligible for S_{indep}^i by satisfying Claim 12 is at least $\sum_j \mathcal{S}_j$.

Proof. For each bin, consider the sum of \mathcal{S}_j for variables corresponding to elements of S_{indep}^{i-1} that were assigned to the bin (if any). If the sum is nonpositive, then we trivially claim the set of elements meeting the criteria in this bin is at least the sum, as there will be at least 0 such elements. Otherwise, the sum must be 1. This implies there is exactly one element of S_{indep}^{i-1} assigned to the bin, and that it met the criteria when it was assigned, because its corresponding $\mathcal{S}_j = 1$. Moreover, as no other elements could have been assigned to the bin later, it still meets the criteria. Combining both cases, we see that the sum of \mathcal{S}_j for each bin is a lower-bound for the number of elements satisfying the criteria in said bin, and thus globally the sum of all \mathcal{S}_j is a lower-bound for how many elements meet the criteria in total. \square

We aim to now use the sum of \mathcal{S}_j as a lower-bound for the size of the set of elements meeting the criteria. To do so, we will first lower-bound $E[\mathcal{S}_j]$.

Claim 14. Regardless of the realization of any previous randomness, $E[\mathcal{S}_j] \geq \frac{1}{3}$.

Proof. \mathcal{S}_j is equal to -1 only if it is assigned to a bin with another element of S_{indep}^{i-1} . The number of such bins is upper-bounded by $|S_{\text{indep}}^{i-1}| \leq |S_{\text{indep}}^1| \leq \frac{n}{6}$. Otherwise, \mathcal{S}_j is equal to 0 only if the bin had mass at least $\frac{n}{6}$ or it has mass from the corresponding $\mathcal{B}_1^{i-1}(x)$ of at least $\frac{6\text{related}_1^{i-1}(x)}{n}$. There can only be at most $\frac{n}{6}$ bins satisfying the former, and at most $\frac{n}{6}$ bins satisfying the latter. Accordingly, there are at least $n - 3 \times \frac{n}{6} = \frac{n}{2}$ where if the corresponding element is assigned to it, then $\mathcal{S}_j = 1$. Hence $E[\mathcal{S}_j] \geq \frac{1}{2} - \frac{1}{6} = \frac{1}{3}$. \square

As we need a set S_{indep}^i with cardinality $|S_{\text{indep}}^i| = \frac{|S_{\text{indep}}^{i-1}|}{6}$, we show the following:

Claim 15. $\sum_j \mathcal{S}_j \geq \frac{|S_{\text{indep}}^{i-1}|}{6}$ with high probability.

Proof. We will modify the variables to make a martingale and then utilize Azuma's inequality. We define $\mathcal{S}'_j = \mathcal{S}'_{j-1} + \mathcal{S}_j - E[(\mathcal{S}_j | \mathcal{S}_1, \dots, \mathcal{S}_{j-1})]$. Accordingly, the sequence of \mathcal{S}' is a martingale of length $|S_{\text{indep}}^{i-1}|$ where $|\mathcal{S}'_{j-1} - \mathcal{S}'_j| \leq 1$. Thus, we can use Azuma's inequality to show $P(|\mathcal{S}'_{|S_{\text{indep}}^{i-1}|} -$

$\mathcal{S}'_1| \geq \frac{|S_{\text{indep}}^{i-1}|}{6}) \leq 2e^{-\frac{|S_{\text{indep}}^{i-1}|}{72}} = 2e^{-\frac{|S|}{726^{i-1}}} = 2e^{-\Omega(n)}$. By definition, $\sum_j \mathcal{S}_j = \sum_j \mathcal{S}'_j + \sum_j E[\mathcal{S}_j]$. By our result with Azuma's inequality, we then claim that with high probability it holds that $\sum_j \mathcal{S}_j \geq -\frac{|S_{\text{indep}}^{i-1}|}{6} + \frac{|S_{\text{indep}}^{i-1}|}{3} = \frac{|S_{\text{indep}}^{i-1}|}{6}$. \square

Combining Claim 12 and Claim 15, we have now shown that there exists a valid set S_{indep}^i of size $|S_{\text{indep}}^i| = \frac{|S_{\text{indep}}^{i-1}|}{6}$, completing the proof of Claim 11. \square

By induction Lemma B.8 holds. \square

Corollary B.9. *There exists a subset of plateau states $S_{\text{indep}} \subseteq S$ such that $|S_{\text{indep}}| \geq \frac{|S|}{6^{|V|}} = \Omega(n)$ and every $x \in S_{\text{indep}}$ satisfies related $_2^{\text{order}(\text{color-root}(Y))}(x) \leq \frac{18 \times |V| \times (|V|-1)}{n^2}$. Moreover, for all pairs $x, x' \in S_{\text{indep}}$ it holds that they never share a state, meaning $f_{X_i}^*(x) \neq f_{X_i}^*(x')$ for all $X_i \in V_{\text{decomp}}$.*

Proof. One such set is simply $|S_{\text{indep}}^{|V|}|$ as shown in Lemma B.8. \square

Characterizing balls. Recall that each variable assigns balls with the same configuration of its parents and exogenous variable together. We aim to show a similar result at $\text{color-root}(Y)$. To do so, we will characterize the balls within configurations into types:

Definition B.10 (Ball characterizations). We characterize three types of balls:

1. **Dense balls.** Consider a set L of states of X , where a state of X is in L if $P(X = x) \geq \frac{1}{\log^3(n)}$. Dense balls are all balls of the form $(x \in L, e \in E)$. We call these dense balls, because the low-entropy of E will prevent the collective mass of these balls from “expanding” well.
2. **Large balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $\geq \frac{T}{2}$.
3. **Small balls.** For all balls of the form $(x \in X \setminus (S \cup L), e \in E)$ where the ball has mass $< \frac{T}{2}$.

We use $\mathcal{T} = \frac{120|V|}{n \log(n)}$.

Now, we will show that for every variable X_i there are many bins without too much surplus, such that the plateau configurations have many bins that they may be assigned to that will help us obtain a bin with small surplus and many plateau configurations.

Definition B.11 (Configuration and ball characterizations). We characterize three types of configurations/balls:

1. **Large configurations.** For all configurations of the form $(\text{Pa}(X_i) \cup E_i)$ where the configuration has balls of total mass $\geq \frac{T}{2}$.

2. **Dense ball.** Consider a set L of states of X_{src} , where a state of X_{src} is in L if $P(X_{\text{src}} = x) \geq \frac{1}{\log^3(n)}$. Dense balls are all balls of the form $(x \in L, e \in E^*)$. We call these dense balls, because the low-entropy of E will prevent the collective mass of these balls from being distributed well throughout.

3. **Small ball.** For all balls of the form $(x \in X_{\text{src}} \setminus (S \cup L), e \in E^*)$ where the ball has mass $< \frac{T}{2}$.

Bounding dense ball surplus. Recall the following used in Theorem 4.1 to bound contributions from dense balls:

Lemma B.12 (Limited expansion). *Suppose Y can be written as a function $f(X, E)$ and $X \perp\!\!\!\perp E$. Consider any subset R of the support of X . For any subset T of the support of Y that satisfies $\forall t \in T : P(X \in R, Y = t) > \delta$, the cardinality of T is upper bounded as $|T| \leq \frac{H(E) + \log(|R|) + 2}{\delta \log(\frac{1}{\delta})}$.*

We use the following corollary:

Corollary B.13. *There exist no subset $|T| = n/4$ such that $\forall t \in T : P(X \in L, Y = t) \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$*

These imply the following for our graph setting. While this may seem strictly weaker than Corollary A.8, we will utilize that the event of a bin having too much mass from dense balls is now independent from how large configurations are mapped.

Corollary B.14. *Let $\mathcal{C}_{\text{large}}$ denote the set of large configurations of $\text{Pa}(X_i) \cup E_i$ as defined in Definition B.11. Let C be a random variable denoting the configuration of the corresponding ball of $\text{Pa}(X_i) \cup E_i^*$. We claim that dense balls in configurations other than $\mathcal{C}_{\text{large}}$ are not distributed well throughout X_i . In particular, there exists no subset $|T| = n/4$ such that $\forall t \in T : P(X \in L, Y = t, C \notin \mathcal{C}_{\text{large}}) \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$.*

Proof. Note that the results of Lemma A.7 and Corollary A.8 still hold in this setting as any $X_i \in V_{\text{decomp}}$ can be written as a function of X_{src} and $\cup_j E_j$. This corollary trivially follows from Corollary A.8, as it is strictly weaker in that we add a restriction that $C \notin \mathcal{C}_{\text{large}}$. Any set T that contradicts Corollary B.14 would immediately contradict Corollary A.8. \square

Bounding large configuration surplus. To bound contribution to surplus by large configurations, we bound the number of bins that receive any mass from large configurations. Recall Lemma A.9 from the proof of Theorem 4.1:

Lemma B.15 (Avoided big). *Given a balls-and-bins game with $c \cdot n \ln(n)$ balls mapped uniformly randomly to n bins, at least $\frac{n^{1-c}}{2}$ bins will receive no balls with high probability if c is a constant such that $0 < c \leq \frac{1}{3}$.*

Accordingly, we can use the following:

Corollary B.16. *As there are at most $\frac{1}{\mathcal{T}/2} \leq \frac{n \log(n)}{60|V|} \leq \frac{1}{40|V|} \cdot n \ln(n)$ large balls, with high probability there are at least $\frac{n^{1-\frac{1}{40|V|}}}{2}$ bins that receive no large balls.*

Bounding small ball surplus. Here we will bound the surplus from small balls. Note that, while this proof is not short, it is using the same ideas as the corresponding section in the proof of Theorem 4.1. However, there are some subtle differences that necessitate a separate proof for the graph setting. We use identical text from the proof of Theorem 4.1 when applicable.

For the small balls, we will also show that they cannot contribute too much surplus to too many states of any X_i . We will notably use that all small balls correspond to a state of X_{src} where $P(X_{\text{src}} = x) \leq \frac{1}{\log^3(n)}$. We will utilize this to show that most small balls are assigned to a state of X_i that has not yet received $> \frac{\mathcal{T}}{2}$ mass from its corresponding state of X_{src} , and accordingly would not increase the surplus. To accomplish this, we define a surplus quantity that only takes into account small balls:

Definition B.17 (Small ball surplus). We define the small ball surplus of a state y of Y as

$$z_j^{\text{small}} = \sum_{x_{\text{src}} \notin (S \cup L)} \max \left(0, -\mathcal{T} + \sum_{\substack{(X=x_{\text{src}}, E^*=e, C \notin \mathcal{C}_{\text{large}}) \\ (X=x_{\text{src}}, E^*=e, X_i=j)}} P(X_{\text{src}} = x, E^* = e, X_i = j) \right).$$

With this notion of surplus constrained to small balls, we show the following:

Lemma B.18 (Small ball limited surplus). *With high probability, there are at most $\frac{n}{4}$ values of i , i.e., number of bins, where $z_i^{\text{small}} \geq \frac{n}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}$.*

Proof. We will consider configurations $C \notin \mathcal{C}_{\text{large}}$ in an arbitrary order, and within each configuration consider balls in an arbitrary order. Let $x_i(c)$ be the corresponding state of X_i for the c -th configuration, let $x_{\text{src}_c}(t)$ be the corresponding state of X_{src} for the t -th ball in the c -th configuration. $e_{\text{src}_c}(t)$ be the corresponding state of E^* for the t -th ball in the c -th configuration, and $w_{c,\text{ball}}(t)$ be the t -th ball's probability mass in the c -th configuration (i.e., $P(X_{\text{src}} = x_{\text{src}_c}(t), E = e_{\text{src}_c}(t))$). Moreover, we define $w_{\text{config}}(c)$ as the weight of all such balls with configuration c . Recall that for all small balls it must hold that $x_{\text{src}_c}(t) \notin L$ and thus $P(X_{\text{src}} = x_{\text{src}_c}(t)) < \frac{1}{\log^3(n)}$. We define the total small ball surplus as $Z^{\text{small}} = \sum_{j \in X_i} z_j^{\text{small}}$. Now, we will

consider all non-large configurations in an arbitrary order and realize their corresponding entry of f to map them to a state of X_i . Initially, we have not realized the entry of f for any balls and thus all $z_j^{\text{small}} = 0$ and $Z^{\text{small}} = 0$. As we map configurations to states of X_i , we define $\Delta(c)$ as the increase of Z^{small} after mapping the c -th configuration to a state of X_i . By definition, $\sum_c \Delta(c)$ is equal to Z^{small} after all values of f have been completely realized.

Our primary intuition is that we will show for many small balls it holds that they have zero contribution towards their configuration's quantity $\Delta(c)$. As f is realized for each configuration, let $w_{X_i}(x_i, x_{\text{src}})$ denote the total mass of balls assigned to state x_i of X_i so far from state x_{src} of X_{src} , i.e., $w_{X_i}(x_i, x_{\text{src}}) := \sum_{c' < c, t: x_i(c) = x_i'} w_{c,\text{ball}}(t')$. Note that this quantity is shared among all configurations.

Claim 16. Regardless of the realizations of all $\Delta(c')$ for $c' < c$, it holds that $\Delta(c)$ is a random variable with values in range $[0, w_{\text{config}}(c)]$ and $E[\Delta(c)] \leq \frac{w_{\text{config}}(c)}{\log^2(n)}$.

Proof. Let us define $\Delta_t(c)$ as the contribution of the t -th ball to $\Delta(c)$. By definition, $\sum_t \Delta_t(c) = \Delta(c)$. We aim to show $E[\Delta_t(c)] \leq \frac{w_{c,\text{ball}}(t)}{\log^2(n)}$. This would immediately imply the desired bound on $E[\Delta(t)]$ by linearity of expectation.

The only conditions under which $\Delta_t(c)$ takes a non-negative value (which is upper-bounded by $w_{c,\text{ball}}(t)$), is when $w_{X_i}(x_i(c), x_{\text{src}_c}(t)) > \frac{\mathcal{T}}{2}$ before the entry of f for the c -th configuration is realized (other. Recall that $P(x_{\text{src}_c}(t)) \leq \frac{1}{\log^3(n)}$. Accordingly, the number of states x_i' of X_i where $w_{X_i}^c(x_i', x_{\text{src}_c}(t)) > \frac{\mathcal{T}}{2}$ is upper-bounded by $\frac{P(X_{\text{src}} = x_{\text{src}_c}(t))}{\mathcal{T}/2} \leq \frac{1/\log^3(n)}{60|V|/(n \log(n))} = \frac{n \log(n)}{60|V| \log^3(n)} \leq \frac{n}{\log^2(n)}$. This is due to the fact that balls partition the total mass of $P(X_{\text{src}} = x_{\text{src}_c}(t))$ since we have $P(X_{\text{src}} = x_{\text{src}_c}(t)) = \sum_e P(X = x_{\text{src}_c}(t), E = e)$. This implies that the probability that the t -th ball of configuration c will be mapped to a state x_i' of X_i such that $w_{X_i}^c(x_i', x_{\text{src}_c}(t))$ already exceeds the threshold of $\mathcal{T}/2$ (in other words where we will have $\Delta_t(c) > 0$) is upper-bounded by $\frac{n/\log^2(n)}{n} = \frac{1}{\log^2(n)}$ due to the fact that the function f is realized uniformly randomly. Accordingly, $E[\Delta_t(c)] \leq \frac{w_{c,\text{ball}}(t)}{\log^2(n)}$ and thus $E[\Delta(c)] \leq \frac{w_{\text{config}}(c)}{\log^2(n)}$. \square

This enables us to upper-bound the sum of $\Delta(t)$:

Claim 17. $\sum_c \Delta(c) \leq \frac{1}{4 \log(n)}$ with high probability.

Proof. We will transform $\Delta(c)$ into a martingale. In particular, we define $\Delta'(c) = \Delta'(c-1) + \Delta(c) - E[\Delta(c) | \Delta(1), \dots, \Delta(c-1)]$. We define $\Delta'(0) = 0$, and note that $\Delta'(c)$ is a martingale. By Azuma's inequality, we

show $|\sum_c \Delta'(c)| \leq \frac{1}{8 \log(n)}$ with high probability:

$$\begin{aligned} P[|\Delta(c)| > \varepsilon] &< 2e^{-\frac{\varepsilon^2}{2 \sum c_i^2}} \\ &\leq 2e^{-\frac{\left(\frac{1}{8 \log(n)}\right)^2}{2(\max_i c_i) \cdot \sum c_i}} \\ &\leq 2e^{-\frac{\left(\frac{1}{8 \log(n)}\right)^2}{2 \times 7/2 \cdot 1}} \\ &= 2e^{-\frac{-n \log(n)}{120 \times 8 \times |V| \times \log(n)}} \end{aligned}$$

Accordingly, by definition of $\Delta'(c)$ this implies $|\left(\sum_c \Delta(c) - \sum_c E[\Delta(c)|\Delta(1), \dots, \Delta(c-1)]\right)| \leq \frac{1}{8 \log(n)}$. By Claim 16 we know all $E[\Delta(c)|\Delta(1), \dots, \Delta(c-1)] \leq \frac{w_{\text{config}(c)}}{\log^2(n)}$ and accordingly, $\sum_c E[\Delta(c)|\Delta(1), \dots, \Delta(c-1)] \leq \frac{1}{\log^2(n)}$. Together, these imply $\sum_c \Delta(c) \leq \frac{1}{8 \log(n)} + \frac{1}{\log^2(n)}$ with high probability, and for sufficiently large n it holds that $\frac{1}{\log^2(n)} \leq \frac{1}{8 \log(n)}$. Thus, our high-probability on $|\Delta'(c)|$ implies that $\sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$ with high probability. \square

Finally, we conclude that our upper-bound on $\sum_t \Delta(t)$ implies an upper-bound on the number of states of Y with non-negligible small ball support:

Claim 18. If $\sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$, then there are at most $\frac{n}{4}$ bins where $z_i^{\text{small}} \geq \frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$.

Proof. By definition, $Z^{\text{small}} = \sum_t \Delta(t) \leq \frac{1}{4 \log(n)}$. Given this upper-bound for total small ball surplus, we can immediately upper-bound the number of states of X_i with small ball surplus greater than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ by the quantity $\frac{1/(4 \log(n))}{1/(n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)})} \leq \frac{n \cdot \log(\log(n)) \cdot \log^{1/2}(n)}{4 \log(n)} \leq \frac{n}{4}$. We obtain this by using $c_{\text{close}} = \frac{1}{4}$ and for sufficiently large n such that $\log(\log(n)) \leq \log^{1/2}(n)$. \square

This concludes the proof of the lemma. \square

Concluding many bins with small surplus. Now, we combine all these intuitions to show there are many bins that have a small amount of surplus. We have shown that, with high probability, there are at most $n/4$ bins with non-negligible mass from dense balls by Corollary A.8, and at most $n/4$ bins with non-negligible surplus from small balls from non-large configurations Lemma B.18. Combining these sets, there are at most $n/2$ bins with non-negligible mass from dense balls or surplus from small balls. By Corollary B.16, with high probability at least $\frac{n^{1-\frac{1}{40|V|}}}{2}$ bins will receive no large configurations mapped to it. Our goal

is to show the intersection of the sets is large, so there are many bins that have small surplus. We use Lemma A.13 proven in Theorem 4.1:

Lemma B.19. *Let there be two sets $A, B \subseteq [n]$, where $|A| \geq \frac{n}{2}$ and A and B are both independently uniformly random subsets of size $|A|$ and $|B|$, respectively. It holds that $P(|A \cap B| \geq \frac{|B|}{4}) \geq 1 - 2e^{-\frac{|B|}{8}}$.*

Corollary B.20. *With high probability, there are at least $\frac{n^{1-\frac{1}{40|V|}}}{8}$ bins with surplus $z_y \leq \frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$.*

Proof. We have defined three types of balls, and have proven results that show how there are many bins with negligible bad contribution for each type of ball. Now, we combine these with Lemma A.13 to show there are many bins where there is not much bad contribution in total. By Corollary A.8 there are at most $n/4$ bins with more than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ mass from dense balls. By Lemma B.18, there are at most $n/4$ bins with small ball surplus more than $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$. Let A be the set of bins with at most $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ mass from dense balls and at most $\frac{1}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ small ball surplus. By combining Corollary A.8 and Lemma B.18 we know $|A| \geq \frac{n}{2}$ with high probability. Let B be the set of bins that receive no large configurations. By Corollary B.16, it holds that $|B| \geq \frac{n^{1-\frac{1}{40|V|}}}{2}$ with high probability. A and B are independent, as A is undetermined by the mapping of large configurations. By Lemma A.13, it holds that $|A \cap B| \geq \frac{n^{1-\frac{1}{40|V|}}}{8}$ with failure probability at most $2e^{-\frac{n^{1-\frac{1}{40|V|}}}{16}}$. Moreover, all such bins will have total surplus at most $\frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$, because they receive no large configurations and total surplus is then upper-bounded by the sum of small ball surplus and total mass from dense balls. \square

Existence of many desirable bins at color-root(Y).

We have shown that at each node X_i there are many bins without much surplus. If we restrict this calculation of surplus to not include mass from plateau configurations in S_{indep} for color-root(Y), then the set of bins that do not have much surplus is independent of the assignment of such configurations with plateau balls not having much related mass. Consider the set $S_{\text{indep}}^{\text{order}(\text{color-root}(Y))}$. By Corollary B.9, we know $|S_{\text{indep}}^{\text{order}(\text{color-root}(Y))}| \geq \frac{|S|}{6|V|}$, no two corresponding plateau balls ever share a state before color-root(Y), and all related $S_{\text{indep}}^{\text{order}(\text{color-root}(Y))}(x) \leq \frac{18 \times |V| \times (|V|-1)}{n^2}$. Recall that color-root(Y) by its definition must create a new color, and thus either have X_{src} as a parent, or have at least two distinct colors in its parent set. Given these properties, we know that each plateau ball in this set has at

most related $\mathbb{2}^{\text{order}(\text{color-root}(Y))}(x) \leq \frac{18 \times |V| \times (|V|-1)}{n^2}$ mass in its configuration for $\text{color-root}(Y)$, and all elements in the set will be in different configurations. We aim to show that there are many bins with small surplus that receive many of these configurations corresponding to the plateau balls in the set. To do so, we use the following results shown in Theorem 4.1. First, we use negative association:

Claim 4. Indicator variables for if a bin receives some threshold of balls in a i.i.d. uniformly random balls-and-bins game are NA.

Second, we lower bound the probability of a bin researching a certain threshold:

Claim 5. Suppose cn balls ($c \leq 1$) are thrown i.i.d. uniformly randomly into n bins. The probability that a particular bin receives at least $k = \frac{d \log(n)}{\log(\log(n))}$ balls is at least $\frac{1}{en^d}$ given that $\frac{d}{c} \leq \log(\log(n))$.

Corollary B.21. *With high probability, there are $\frac{n^{1-\frac{1}{20|V|}}}{16e}$ bins with $z_j \leq \frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ and at least $\frac{\log(n)}{40|V| \log(\log(n))}$ configurations mapped from states of $\mathcal{S}_{\text{indep}}^{\text{order}(\text{color-root}(Y))}$.*

Proof. Consider the indicator variable \mathcal{B}_i if a bin met the threshold of plateau configurations. We show that with high probability $\sum_i \mathcal{B}_i$ is large enough, considering just the bins with small surplus. By Corollary B.20 we know there are at least $\frac{n^{1-\frac{1}{40|V|}}}{8}$ such bins with high probability. Now let us focus on just the configurations corresponding to $\mathcal{S}_{\text{indep}}^{\text{order}(\text{color-root}(Y))}$ that we know has cardinality $\Omega(n)$. By Claim 5, the probability of a bin receiving at least $\frac{\log(n)}{40|V| \log(\log(n))}$ such configurations is at least $\frac{1}{en \frac{1}{40|V|}}$.

Accordingly, it holds that $\sum_i E[\mathcal{B}_i] \geq \frac{n^{1-\frac{1}{20|V|}}}{8e}$.

By Hoeffding's inequality, we show $|\sum_i \mathcal{B}_i - \sum_i E[\mathcal{B}_i]| \leq \frac{n^{1-\frac{1}{20|V|}}}{16e}$ with high probability:

$$\begin{aligned} P[|S_n - E_n| > t] &< 2e^{-\frac{2t^2}{\sum c_i^2}} \\ &\leq 2e^{-\frac{2n^{2-\frac{1}{10|V|}}}{16^2 e^2 n^{1-\frac{1}{40|V|}}}} \\ &\leq 2e^{-\frac{n^{1-\frac{3}{40|V|}}}{16e^2}}. \end{aligned}$$

Therefore, $\sum_i \mathcal{B}_i \geq \frac{n^{1-\frac{1}{20|V|}}}{16e}$ with high probability. \square

Survival of desirable bins to Y . Now we aim to show that of the bins that received many plateau configurations and had small surplus, that enough will ‘‘survive’’ and keep these properties as we process nodes within the same color as Y ,

and that eventually at least one such bin will survive to Y with high probability.

Lemma B.22. *After processing i nodes of the same color as Y , with high probability there are at least $\frac{n^{1-\frac{1}{20|V|}}}{16e}$ sets of plateau balls, such that each set has cardinality at least $\frac{\log(n)}{100 \log(\log(n))}$, have been assigned together to a bin with surplus $z_j \leq \frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$, and no two sets were ever mapped to the same state within this color.*

Proof. Trivially, this holds for $i = 1$ from Corollary B.21.

First, we note that all sets of plateau balls will again be mapped together. This is because they are in the same configuration, as for any node X_i , as it inherits its color, it must be true that they all have the same values for $\text{Pa}(X_i)$ and because they are plateau balls they must have $E^* = e_1^*$.

Now, we will make a random variable \mathcal{S}_i for whether the i -th configuration survived together. Roughly, we desire \mathcal{S}_i to be 1 if it is assigned to a bin with small surplus with none of the other bins that has survived to this stage, we desire \mathcal{S}_i to be 0 if it is assigned to a bin with non-small surplus, and \mathcal{S}_i to be -1 if it lands in a small surplus bin with another bin that had survived (the intuition is that said bin would likely have a positive \mathcal{S}_j and now we must cancel them out). Now, we slightly modify \mathcal{S}_i so all \mathcal{S}_i are independent.

By Corollary B.20 we know there will be at least $\frac{n^{1-\frac{1}{40|V|}}}{8}$ small surplus bins with high probability. Let us create a subset of bad bins B_{bad} for which \mathcal{S}_i will take value -1 . Before realizing the assignment for the i -th configuration, add all small surplus bins that have already received a bin that survived to this round. Arbitrarily fill the remainder of

B_{bad} so that $|B_{\text{bad}}| = \frac{n^{1-\frac{1}{20|V|}}}{16e}$. Let us define $|B_{\text{good}}| = \frac{n^{1-\frac{1}{40|V|}}}{8} - \frac{n^{1-\frac{1}{20|V|}}}{16e}$. So, if the configuration is assigned to B_{bad} then $\mathcal{S}_i = -1$, if assigned to B_{good} then $\mathcal{S}_i = 1$, and otherwise $\mathcal{S}_i = 0$. By Hoeffding's inequality, it holds that $\sum_i \mathcal{S}_i \geq \frac{n^{1-\frac{1}{20|V|}}}{16e}$ with high probability and thus the lemma holds. \square

Concluding large conditional entropy from desirable bin. By Lemma B.22, it is clear that with high probability there is at least $\frac{n^{19/20}}{16e} > \sqrt{n}$ bin y' of Y satisfying the desired properties. Now, we seek to prove that this implies $H(X_{\text{src}}|Y = y')$. Consider a looser definition of surplus:

Definition B.23 (Relaxed Surplus, $\mathcal{T}^{\text{relax}} = \frac{120|V|}{n \log(n)}$). We define the surplus of a state i of Y as $z_i^{\text{relax}} = \sum_{j \notin \mathcal{S}} \max(0, P(X = j, Y = i) - \frac{120|V|}{n \log(n)})$.

Claim 19. There exists a bin with at least $\frac{\log(n)}{100 \log(\log(n))}$ plateau balls and relaxed surplus at most $\frac{2|V|}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$

Proof. There are two contributors towards relaxed surplus. First, when configurations were assigned to $\text{color-root}(Y)$, each plateau ball brought related₂(x) mass with it that could contribute to the surplus. By Lemma B.8, each of the plateau balls we considered satisfied $\text{related}_2(x) \leq \frac{18|V|(|V|-1)}{n^2}$. Accordingly, there is at most $n \times \frac{18|V|(|V|-1)}{n^2} \leq \frac{18|V|^2}{n}$ such mass in total. Among the at least \sqrt{n} bins that survived to Y , let us choose the one with the least initial mass from related₂. Accordingly, it must have at most $\frac{18|V|^2}{n^{1.5}}$ such mass.

Now, consider the at most $|V| - 1$ times that mass may have been acquired by landing in a bin with at most $\frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$ surplus. Combining all these masses and calculating the worst-case relaxed mass results in an upper-bound of $\frac{2}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}} \times (|V| - 1) + \frac{18|V|^2}{n^{1.5}} \leq \frac{2|V|}{n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$. This is because the definition of relaxed surplus gives enough threshold to fit within it all the mass that was within the regular surplus threshold for each of the groups we are aggregating. \square

Now, we show that this implies $H(X|Y = y')$ is large, with almost exactly the same proof as Lemma A.16:

Lemma B.24 (High-entropy conditional). *Given a bin y' that has $z_{y'}^{\text{relax}} \leq \frac{2|V|}{n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}}$, and receives $\frac{\log(n)}{100 \log(\log(n))}$ plateau balls, then $H(X_{\text{src}}|Y = y') = \Omega(\log(\log(n)))$.*

Proof. To show $H(X_{\text{src}}|Y = y')$ is large, we first define the vector v such that $v(x) = P(X_{\text{src}} = x, Y = y')$. Similarly, we define $\bar{v}(x) = \frac{v}{P(Y=y')}$, meaning $\bar{v}(x) = P(X_{\text{src}} = x|Y = y')$ and $|\bar{v}|_1 = 1$. Our underlying goal is to show $H(\bar{v})$ is large. To accomplish this, we will split the probability mass of v into three different vectors $v_{\text{initial}}, v_{\text{plateau}}, v_{\text{surplus}}$ such that $v = v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}$. The entries of v_{plateau} will correspond to mass from plateau states of X , v_{initial} will correspond to the first $\mathcal{T}^{\text{relaxed}}$ mass from non-plateau states of X_{src} , and v_{surplus} will correspond to mass that contributes to the surplus $z_{y'}$. We more formally define the three vectors as follows:

- v_{plateau} . The vector of probability mass from plateau states of X_{src} . $v_{\text{plateau}}(x)$ is 0 if $x \notin S$ and $v_{\text{plateau}}(x) = P(X_{\text{src}} = x, Y = y')$ if $x \in S$.
- v_{initial} . For non-plateau states of X_{src} , their first \mathcal{T} probability mass belongs to v_{initial} . $v_{\text{initial}}(x) = \min(P(X = x, Y = y'), \mathcal{T}^{\text{relaxed}})$ if $x \notin S$ and $v_{\text{initial}}(x) = 0$ otherwise.
- v_{surplus} . For non-plateau states of X_{src} , their probability mass beyond the first $\mathcal{T}^{\text{relaxed}}$ mass belongs to v_{surplus} . This corresponds to the surplus quantity. $v_{\text{surplus}}(x) =$

$\max(0, P(X_{\text{src}} = x, Y = y') - \mathcal{T}^{\text{relaxed}})$ if $x \notin S$ and $v_{\text{surplus}}(x) = 0$ otherwise. By this definition, $z_{y'} = |v_{\text{surplus}}|_1$.

To show $H(X_{\text{src}}|Y = y') = H(\bar{v})$ is large, we divide our approach into two steps:

1. Show there is substantial helpful mass: $|v_{\text{initial}} + v_{\text{plateau}}|_1 = \Omega\left(\frac{1}{n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}}\right)$
2. Show the distribution of helpful mass has high entropy: $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) = \Omega(\log(\log(n)))$
3. Show that, even after adding the hurtful mass, the conditional entropy is large: $H(X_{\text{src}}|Y = y') = H(\bar{v}) \geq H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - O(1) = \Omega(\log(\log(n)))$

In the first step, we are showing that the distribution when focusing on just the helpful mass of $v_{\text{initial}}, v_{\text{plateau}}$ has high a substantial amount of probability mass. In the second step, we prove how this distribution of helpful mass has high entropy. In the third step, we show that the hurtful mass of v_{surplus} does not decrease entropy more than a constant.

First, we show that there is a substantial amount of helpful mass:

Claim 20. $|v_{\text{initial}} + v_{\text{plateau}}|_1 = \frac{1}{100c_{\text{pb}}n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}}$

Proof. Recall that the bin y' received $\frac{\log(n)}{100 \log(\log(n))}$ plateau balls. As defined in Lemma A.3, the set S of plateau states is defined such that $\frac{\max_{x \in S} P(X=x)}{\min_{x \in S} P(X=x)} \leq \log^{c_{\text{close}}(n)}$ and $\min_{x \in S} P(X = x) \geq \frac{1}{c_{\text{pb}}n \log(n)}$. Also recall that by Lemma A.4 the most probably state of E has large probability. In particular, $P(E = e_1) \geq \frac{1}{\log^{c_{\text{close}}(n)}}$. Let the subset $S' \subseteq S$ be the subset of plateau states of X such that their plateau ball is mapped to y' . In particular, for every $x \in S'$ it holds that $f(x, e_1) = y'$. Accordingly, $P(X_{\text{src}} = x, Y = y') \geq P(X_{\text{src}} = x) \cdot P(E = e_1)$ for $x \in S'$. Thus, the total weight from plateau states of X_{src} is at least $|S'| \cdot \min_{x \in S'} P(X_{\text{src}} = x) \cdot P(E = e_1) \geq |S'| \cdot \frac{\max_{x \in S'} P(X_{\text{src}}=x)}{\log^{c_{\text{close}}(n)}} \cdot P(E = e_1) \geq \frac{1}{100c_{\text{pb}}n \log(\log(n)) \log^{2c_{\text{close}}(n)}}$. \square

Next, we show the distribution of helpful mass has high entropy:

Claim 21. $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) \geq \frac{\log(\log(n))}{4}$

Proof. Let us define $\bar{v}_{\text{helpful}} = \frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}$ to be the vector of helpful mass, and we will show $H(\bar{v}_{\text{helpful}})$ is large by upper-bounding $\max_x \bar{v}_{\text{helpful}}(x)$.

For non-plateau states of X_{src} , it follows from Claim 20 that $\max_{x \notin S} \bar{v}_{\text{helpful}}(x) \leq \frac{\mathcal{T}^{\text{relaxed}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1} \leq \frac{\mathcal{T}^{\text{relaxed}}}{100c_{\text{lb}}n \cdot \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}} = \frac{100|V|c_{\text{lb}} \log(\log(n)) \cdot \log^{2c_{\text{close}}(n)}}{\log(n)}$.

For plateau states of X , in Claim 20 we also developed the lower-bound of $|v_{\text{initial}} + v_{\text{plateau}}|_1 \geq |S'| \cdot \frac{\max_{x \in S'} P(X=x)}{\log^{c_{\text{close}}(n)}} \cdot P(E = e_1) \geq \frac{\log(n) \cdot \max_{x \in S'} P(X=x)}{2 \log^{2c_{\text{close}}(n)} \log(\log(n))}$. Accordingly, we can upper-bound $\max_{x \in S'} \bar{v}_{\text{helpful}}(x) \leq \frac{\max_{x \in S'} P(X=x)}{|v_{\text{initial}} + v_{\text{plateau}}|_1} \leq \frac{100 \log(\log(n)) \log^{2c_{\text{close}}(n)}}{\log(n)}$.

Accordingly, we can lower-bound the entropy of $H(\bar{v}_{\text{helpful}}) = \sum_x \bar{v}_{\text{helpful}}(x) \cdot \log\left(\frac{1}{\bar{v}_{\text{helpful}}(x)}\right) \geq \sum_x \bar{v}_{\text{helpful}}(x) \cdot \log\left(\frac{1}{\max_{x'} \bar{v}_{\text{helpful}}(x')}\right) = \log\left(\frac{1}{\max_{x'} \bar{v}_{\text{helpful}}(x')}\right) \geq \log\left(\frac{1200c_{\text{lb}} \log(n)}{\log^{2c_{\text{close}}(n)} \log(\log(n))}\right) = (1 - 2c_{\text{close}}) \log(\log(n)) - \log(\log(\log(n))) - \log(1200c_{\text{lb}}) = \frac{\log(\log(n))}{2} - \log(\log(\log(n))) - \log(1200c_{\text{lb}}) \geq \frac{\log(\log(n))}{4}$ for sufficiently large n where $\frac{\log(\log(n))}{2} \geq \log(\log(\log(n))) + \log(1200c_{\text{lb}})$. \square

Finally, we show the hurtful mass does not decrease entropy much, and thus our conditional distribution has high entropy:

Claim 22. $H(X|Y = y') = H(\bar{v}) \geq \Omega(1) \cdot H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - O(1) = \Omega(\log(\log(n)))$

Proof. We lower-bound $H(\bar{v})$ with the main intuitions that $H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) = \Omega(\log(\log(n)))$ and $\frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} = \Omega(1)$. We more precisely obtain this lower-bound for $H(\bar{v})$ as follows:

$$\begin{aligned} H(\bar{v}) &= H\left(\frac{v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}\right) \\ &= \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x) + v_{\text{surplus}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\ &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x) + v_{\text{surplus}}(x)} \\ &\geq \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\ &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x)} - 2 \quad (10) \\ &\geq \sum_x \frac{v_{\text{initial}}(x) + v_{\text{plateau}}(x)}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} \times \\ &\quad \log \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{v_{\text{initial}}(x) + v_{\text{plateau}}(x)} - 2 \\ &= \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}} + v_{\text{surplus}}|_1} H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) \\ &- 2 \end{aligned}$$

$$\begin{aligned} &= \frac{|v_{\text{initial}} + v_{\text{plateau}}|_1}{|v_{\text{initial}} + v_{\text{plateau}}|_1 + z_{y'}} H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - 2 \\ &\geq \frac{1}{1 + 50c_{\text{lb}}|V|} \cdot H\left(\frac{v_{\text{initial}} + v_{\text{plateau}}}{|v_{\text{initial}} + v_{\text{plateau}}|_1}\right) - 2 \quad (11) \end{aligned}$$

$$= \Omega(\log(\log(n))) \quad (12)$$

To obtain Step 10, we note that all summands are manipulated from the form $\sum_x p_x \log\left(\frac{1}{p_x}\right)$ to $\sum_x p'_x \log\left(\frac{1}{p'_x}\right)$ where $p'_x \leq p_x$ for all x . As the derivative of $p \log\left(\frac{1}{p}\right)$ is non-negative for $0 \leq p \leq \frac{1}{e}$, the value of at most two summands can decrease, and they can each decrease by at most one. To obtain Step 11, we use Claim 20. To obtain Step 12, we use Claim 22. \square

Thus, we have shown $H(X|Y = y') = \Omega(\log(\log(n)))$. \square

Corollary B.25. *Under our assumptions, $H(X_{\text{src}}|Y = y') = \Omega(\log(\log(n)))$ and thus $H(\bar{E}) = \Omega(\log(\log(n)))$.*

C. Counterexample for General Identifiability with Unconfounded-Pairwise Oracles

We formalize the oracle first discussed in Section 5:

Definition C.1 (Unconfounded-pairwise oracle). An unconfounded-pairwise oracle is an oracle that returns the correct orientation of an edge if the edge exists in the true graph and if there is no confounding for the edge.

Consider a causal graph G_1 with four nodes and the edge set $\{X_1 \rightarrow X_2, X_1 \rightarrow X_3, X_2 \rightarrow X_4, X_2 \rightarrow X_3, X_3 \rightarrow X_4\}$. Likewise, consider G_2 with four nodes and the edge set $\{X_2 \rightarrow X_1, X_3 \rightarrow X_1, X_4 \rightarrow X_2, X_2 \rightarrow X_3, X_4 \rightarrow X_3\}$. Note that these graphs are in the same Markov equivalence class. Finally, consider a causal graph G_3 with four nodes and the edge set $\{X_1 \rightarrow X_2, X_1 \rightarrow X_3, X_4 \rightarrow X_2, X_2 \rightarrow X_3, X_4 \rightarrow X_3\}$. If we orient all edges in the skeleton for G_1 and G_2 without conditioning using an unconfounded-pairwise oracle, G_3 is a consistent output of the oracle for both G_1 and G_2 . As a result, the peeling approach would see the same set of edge orientations for G_1 and G_2 and not be able to identify the true source.

D. Proof of Theorem 5.6

We need to show that in each iteration of the **while** loop in *line 4*, the algorithm correctly identifies all non-source nodes and only the true non-sources. In the first iteration of the loop, there is no node to condition on and therefore tests are independence (not conditional independence) tests. Consider a non-source node N in the initial graph. Then there must be a directed path from some source node X_{src} to N . Due to the faithfulness assumption, two nodes with a

directed path cannot be unconditionally independent. Then either $\text{Oracle}(X_{\text{src}}, N|\emptyset)$ will return $X_{\text{src}} \rightarrow N$ in *line 12* or $\text{Oracle}(N, X_{\text{src}}|\emptyset)$ will return $X_{\text{src}} \rightarrow N$. In either case, N is added to the non-source list. Now suppose S is a source node. It is never identified as a non-source since it is either conditionally independent with the node it is compared with, found in *line 9*, or it is conditionally dependent with a non-source (for which it has a path to) and the oracle orients correctly in *line 11* or *line 13*. Since all non-sources are correctly identified and only the true non-sources are identified as non-sources, all sources are correctly identified as well in *line 15*. Since sources are incomparable in the partial order, the order in which they are added to the topological order does not impact the validity of topological order in *line 18*.

Suppose the **while** loop identified all sources correctly for all iterations $j, \forall j < i$. Let \mathcal{S}_i be the sources in \mathcal{R} , i.e., the sources that are to be discovered in iteration i . Let $\text{Pa}_{\mathcal{S}_i}$ be the set of parents of \mathcal{S}_i in the initial graph. Then $\text{Pa}_{\mathcal{S}_i} \subseteq \mathcal{C}$, i.e., the set of conditioned nodes include all the parents of the current source nodes. Therefore, \mathcal{C} blocks all backdoor paths from \mathcal{S}_i , effectively disconnecting the previous found sources from the graph: This is because $G \setminus \mathcal{C}$ is a valid Bayesian network for the conditional distribution $p(\cdot | \mathcal{C} = c)$. Therefore, in $G \setminus \mathcal{C}$, \mathcal{S}_i are source nodes and the source pathwise oracle correctly identifies all non-source nodes, similar to the base case.

This implies after the **while** loop terminates, we have a valid topological order for the nodes in the graph. Finally, the for loop in *line 19* converts the obtained total order into a partial order since either it removes an edge, or adds an edge that is consistent with the topological order. Therefore we only need to show that non-edges are correct. Suppose X_i, X_j are non-adjacent in the true graph. Then conditioned on all ancestors of X_i, X_j they are independent due to d-separation in the graph. Therefore all non-edges are identified at some iteration of the **for loop**. Furthermore, under the faithfulness assumption, no edge can be mistakenly identified as a non-edge: No two adjacent nodes at any stage of the algorithm are independent conditioned on the ancestors of the cause variable.

E. Additional Experiments

E.1. Further Synthetic Experiments and Comparisons with ANM

Figure 5 compares the performance under the additive noise model assumption, i.e., the data is sampled from $X = f(\text{Pa}_X) + N_X$ for all variables. The noise term is chosen as a uniform, zero-mean cyclic noise in the supports $\{-1, 0, 1\}, \{-2, -1, 0, 1, 2\}, \{-3, -2, -1, 0, 1, 2, 3\}$ which corresponds to different entropies $H(N_X)$. This entropy is shown on the x -axis in Figure 5. The corre-

sponding plots where the x -axis represents the number of samples are given in the Appendix. While this is the generative model that discrete ANM is designed for, its performance is either approximately matched or exceeded by entropic enumeration.

For further experiments with different number of nodes, please see Figures 5, 6, 10.

E.2. Entropy Measure for Peeling Algorithm

In this section, we compare different versions of peeling algorithm, one that uses only the exogenous entropy and one that uses the total entropy in pairwise comparisons. We randomly sample exogenous distributions according to symmetric Dirichlet distribution, which is characterized by a single parameter α . By varying α and with rejection sampling, we are able to generate distributions for the exogenous nodes E such that $H(E) \leq \theta$ for some θ . For each distribution, we then compare the structural Hamming distance of the output of our peeling algorithm with the true graph. The structural Hamming distance (SHD) is the number of edge modifications (insertions, deletions, flips) required to change one graph to another. Results are given in Figure 12. Let $H(E)$ and $H(\tilde{E})$ be the minimum exogenous entropy needed to generate Y from X and the minimum exogenous entropy needed to generate X from Y , respectively. At each step of the algorithm for every pair X, Y , the red curve compares $H(X)$ with $H(Y)$, the blue curve compares $H(E)$ with $H(\tilde{E})$, and the green curve compares $H(X) + H(E)$ with $H(Y) + H(\tilde{E})$ and orients the edge based on the minimum. As expected, comparing exogenous entropies perform better than comparing observed variables' entropies in the low-entropy regime. Interestingly, we observe that comparing total entropies consistently performs much better than either.

E.3. Entropy Percentile of True Graph

In this section, We test the hypothesis that *when true exogenous entropies are small, the true causal graph is the DAG that minimizes the total entropy*. To test this, we find the minimum entropy needed to generate the joint distribution for every directed acyclic graph that is consistent with the true graph skeleton. We then look at the percentile of the entropy of the true graph. For example, if there are 5 DAGs with less entropy than the true graph out of 100 distinct DAGs, then the percentile is $1 - 5/100 = 0.95$.

E.3.1. SYNTHETIC DATA

Figure 13 shows the results for various graphs. It can be seen that, especially for dense graphs, the true graph is the unique minimizer of total entropy needed to generate the joint distribution for a very wide range of entropy values. This presents exhaustive search as a practical algorithm for

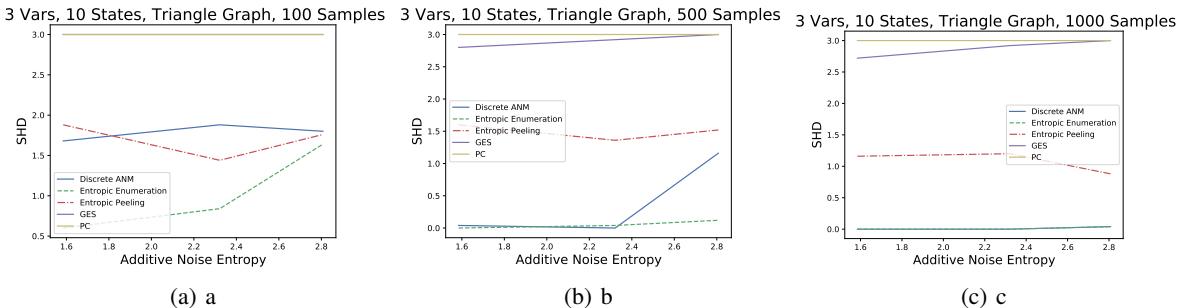


Figure 5. Performance of methods in the ANM setting in the triangle graph $X \rightarrow Y \rightarrow Z, X \rightarrow Z$: 25 datasets are sampled for each configuration from the ANM model $X = f(\text{Pa}_X) + N$. The x -axis shows entropy of the additive noise.

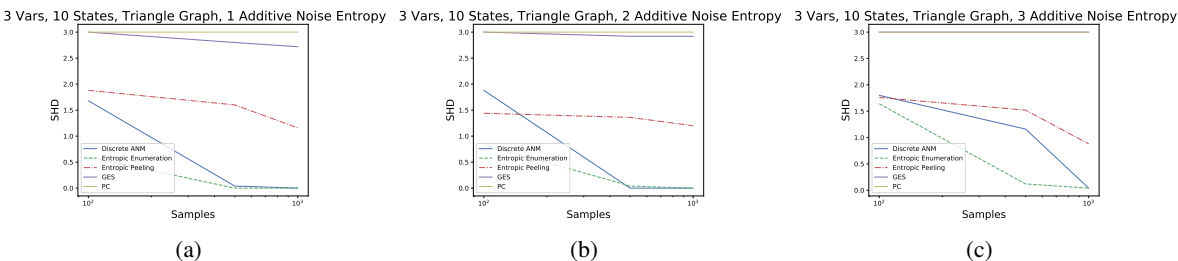


Figure 6. Performance of methods in the ANM setting in the triangle graph $X \rightarrow Y \rightarrow Z, X \rightarrow Z$: 25 datasets are sampled for each configuration from the ANM model $X = f(\text{Pa}_X) + N$. The x -axis shows the number of samples in each dataset. Entropic enumeration outperforms ANM algorithm in the low-noise low-sample regime.

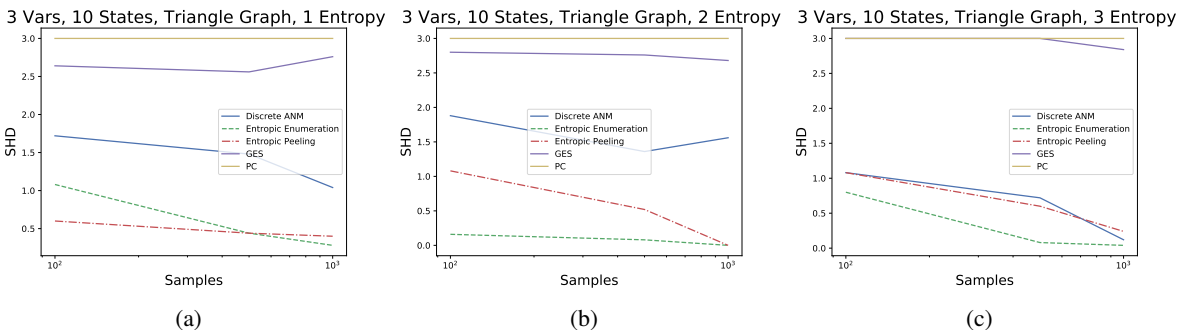


Figure 7. Performance of methods in the unconstrained setting in the triangle graph $X \rightarrow Y \rightarrow Z, X \rightarrow Z$: 25 datasets are sampled for each configuration from the unconstrained model $X = f(\text{Pa}_X, E_X)$. The x -axis shows the number of samples in each dataset. Entropic enumeration and peeling algorithms consistently outperform the ANM algorithm in almost all regimes.

graphs with a small number of nodes (or generally, those for which the MEC is small). Our experiments, contrary to our theory, show that even when the number of nodes is the same as the number of states, entropic causality can be used for learning the graph.

E.3.2. SEMI-SYNTHETIC DATA

We use the *bnlearn* repository⁴ which contains a selection of Bayesian network models curated from real data (Scutari,

2010). Each model in *bnlearn* contains the real causal DAG along with a .bif file that details the conditional distributions that define the distribution. Using this, we generate some number of samples (specified per experiment) and evaluate a collection of causal graph discovery algorithms

⁴<https://www.bnlearn.com/bnrepository/>

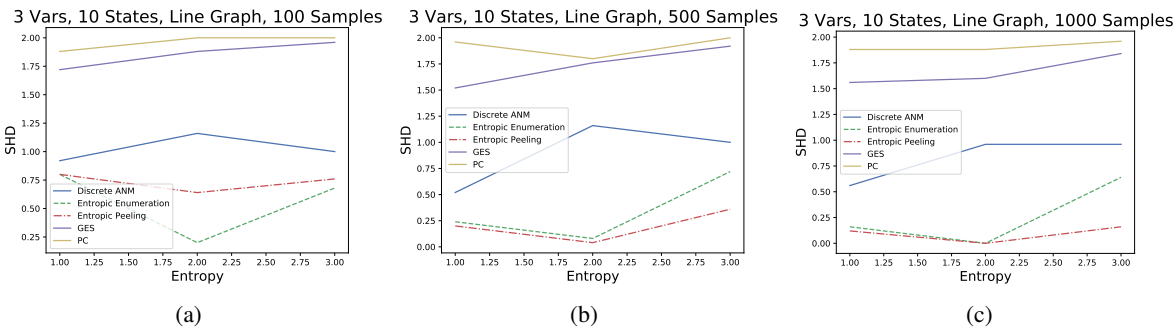


Figure 8. Performance of methods in the unconstrained setting on the line graph $X \rightarrow Y \rightarrow Z$: 25 datasets are sampled for each configuration from the unconstrained model $X = f(\text{Pa}_X, E_X)$. The x -axis shows entropy of the exogenous noise. Entropic enumeration and peeling algorithms consistently outperform the ANM algorithm in all regimes.

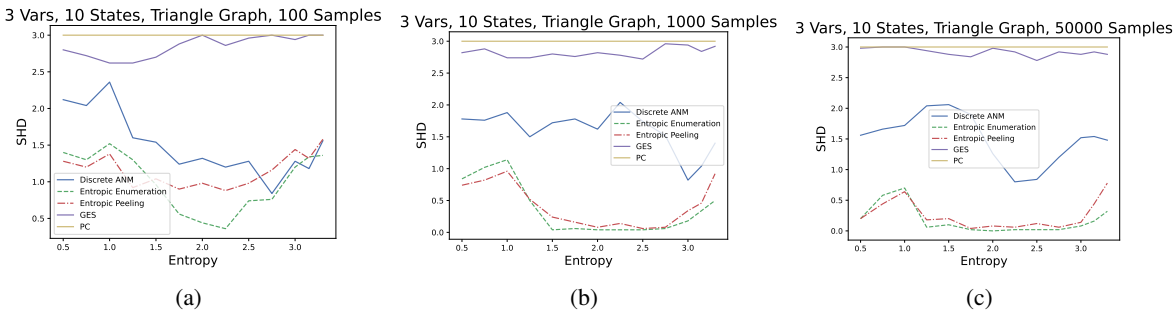


Figure 9. Performance of methods in the unconstrained setting in the triangle graph $X \rightarrow Y \rightarrow Z, X \rightarrow Z$: 50 datasets are sampled for each configuration from the unconstrained model $X = f(\text{Pa}_X, E_X)$. The x -axis shows entropy of the exogenous noise. Entropic enumeration and peeling algorithms consistently outperform the ANM algorithm in almost all regimes. Note how unlike Figure 2, we do not fix the source to have high-entropy or treat the source differently than the other nodes.

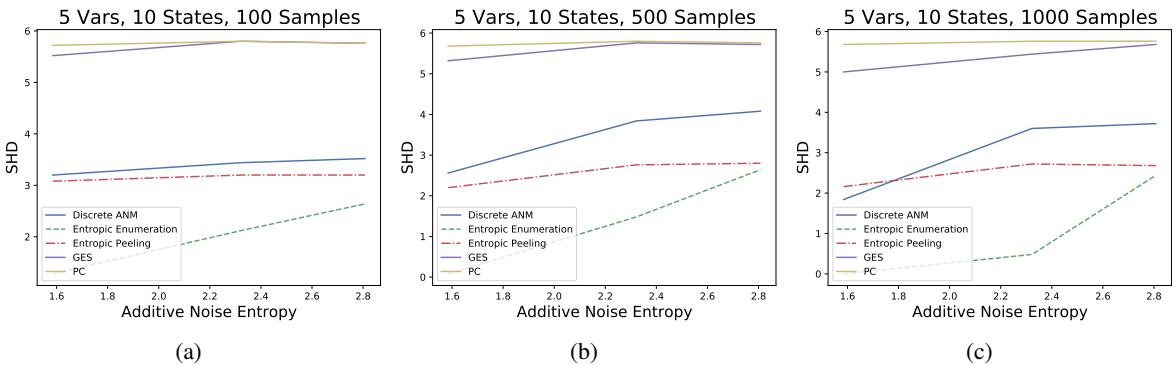


Figure 10. Performance of methods in the ANM setting on random 5-node graphs: 25 datasets are sampled for each configuration from the ANM $X = f(\text{Pa}_X) + N$. The x -axis shows entropy of the exogenous noise. Entropic enumeration outperforms consistently in all regimes.

given the graph skeleton (which can be learned using, e.g., Greedy Equivalence Search) and generated samples as input. Using these models, we can generate any number of samples and test the accuracy of our algorithms. The datasets

however are often binary which makes them less suitable for Algorithm 1. We therefore limit our use of this data to test our hypothesis that the true graph has the smallest entropy among all graphs consistent with the skeleton. The




Figure 11. Performance of methods in the unconstrained setting on random 5-node graphs: 25 datasets are sampled for each configuration from a random graph from the unconstrained model $X = f(PA_X, N)$. The x -axis shows entropy of the additive noise. Entropic enumeration and peeling algorithms consistently outperform the ANM algorithm in all regimes.

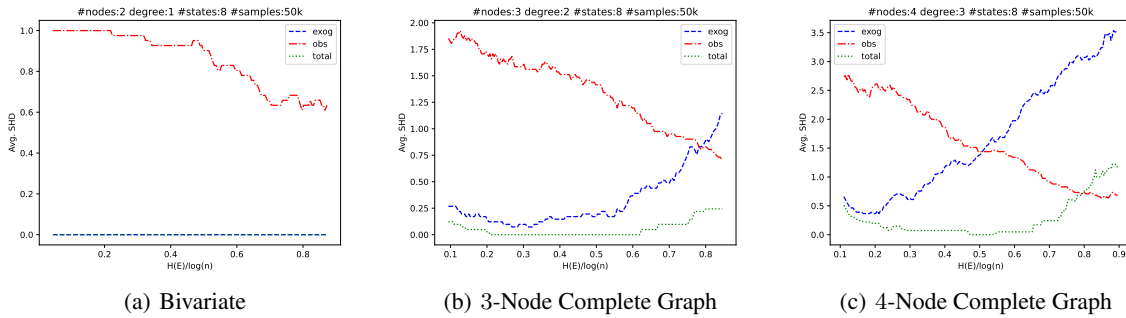


Figure 12. Average structural Hamming distance (SHD) of peeling algorithm on synthetic data for comparing *i*) exogenous entropies (blue, dashed), *ii*) entropies of observed variables (red, dotted-dashed) and *iii*) total entropies (green, dotted) in line 11 as the Oracle for Algorithm 1. Randomly orienting all edges would result in an average SHD equal to half the number of edges (0.5, 1.5, and 3.0 for (a), (b), and (c), respectively).

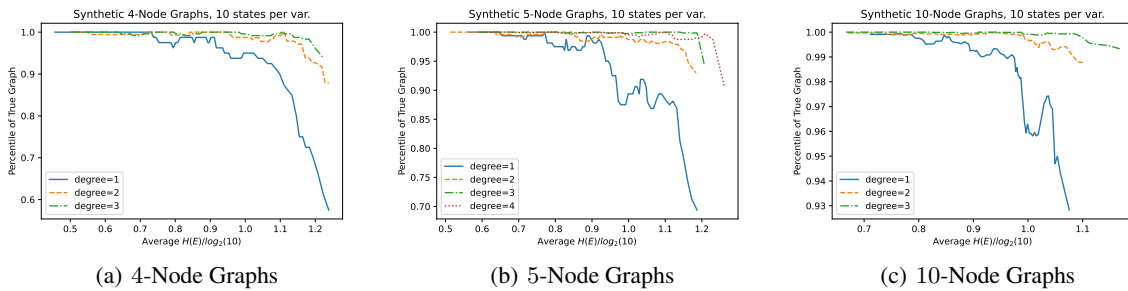


Figure 13. Percentile of the true graph's entropy compared to minimum entropy required to fit every other incorrect possible causal graph that is consistent with the skeleton (synthetic data).

results are given in Figure 14. As can be seen, for most of the small-sized graphs that we tested, the true graph has one of the smallest entropies. Indeed for *Cancer* dataset, the true graph is the unique minimizer when the number of samples is large enough. Only in *Sachs* data does the

true causal graph require one of the largest entropies. This shows that our low-entropy assumption is viable for some real datasets.

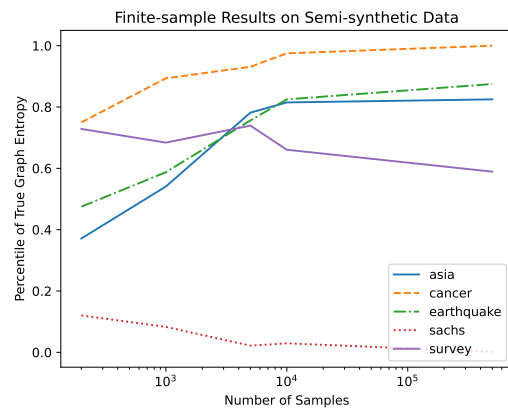


Figure 14. Percentile of true graphs entropy compared to minimum entropy required to fit wrong causal graphs in semi-synthetic data from *Bayesian Network Repository* (Scutari, 2010).