

Applications of Information Inequalities to Database Theory Problems *

Dan Suciu
University of Washington
 Seattle, USA
 suciu@cs.washington.edu

May 2, 2023

Abstract

The paper describes several applications of information inequalities to problems in database theory. The problems discussed include: upper bounds of a query’s output, worst-case optimal join algorithms, the query domination problem, and the implication problem for approximate integrity constraints. The paper is self-contained: all required concepts and results from information inequalities are introduced here, gradually, and motivated by database problems.

1 Introduction

Notions and techniques from information theory have found a number of uses in various areas of database theory. For example, entropy and mutual information have been used to characterize database dependencies [Lee87a, Lee87b] and normal forms in relational and XML databases [AL02, AL05]. More recently, information inequalities were used with much success to obtain tight bounds on the size of the output of a query on a given database [AGM13, GLVV12, GM14, KNS16, KNS17], *and* to devise query plans for worst-case optimal join algorithms [KNS16, KNS17]. Information theory was also used to compare the sizes of the outputs of two queries, or, equivalently, to check query containment under bag semantics [KR11, KKNS21]. Finally, information theory has been used to reason about approximate integrity constraints in the data [KS22, KMP⁺20].

This paper presents some of these recent applications of information theory to databases, in a unified framework. All applications discussed here make use of information inequalities, which have been intensively studied in the information theory community [Yeu08, ZY97, ZY98, Mat07, KR13]. We will introduce gradually the concepts and results on information inequalities, motivating them with database applications.

We start by presenting in Sec. 3 a celebrated result in database theory: the AGM upper bound, which gives a tight upper bound on the query output size, given the cardinalities of the input relations. The AGM bound was first introduced by Grohe and Marx [GM06], and refined in its current form by Atserias, Grohe, and Marx [AGM13], hence the name AGM. (A related result appeared earlier in [FK98].) While the original papers already used information inequalities to prove these bounds, in this paper we provide an alternative, elementary proof, which is based on a family of inequalities due to Friedgut [Fri04], and which are of independent interest.

*Research supported in part by NSF IIS 1907997, NSF IIS 1954222, and NSF-BSF 2109922.

Next, we turn our attention in Sec. 4 to an extension of the AGM bound, by providing an upper bound on the size of the query’s output using functional dependencies and statistics on degrees, in addition to cardinality statistics. The extension to functional dependencies was first studied by Gottlob et al. [GLVV12] and then by Khamis et al. [KNS16], while the general framework was introduced by Khamis et al. [KNS17]. Here, information inequalities are a necessity, and we use this opportunity to introduce entropic vectors and polymatroids, and to define information inequalities. We show simple examples of how to compute upper bounds on the query’s output size using Shannon inequalities (monotonicity and submodularity, reviewed in Sec. 4.1).

A natural question is whether the upper bound on the query’s output size provided by information inequalities is tight: we discuss this in Sec. 5. This question is surprisingly subtle, and it requires us to dig even deeper into information theory, and discuss non-Shannon inequalities. More than 30 years ago, Pippenger [Pip86] asserted that constraints on entropies are the “*laws of information theory*” and asked whether the basic Shannon inequalities form the complete laws of information theory, i.e., whether every constraint on entropies can be derived from the Shannon’s basic inequalities. In a celebrated result published in 1998, Zhang and Yeung [ZY98] answered Pippenger’s question negatively by finding a linear inequality that is satisfied by all entropic functions with 4 variables, but cannot be derived from Shannon’s inequalities. Later, Matúš [Mat07] proved that, for 4 variables or more, there are infinitely many, independent non-Shannon inequalities. In fact, it is an open problem whether the validity of an information inequality is decidable. We provide here a short, self-contained proof of Zhang and Yeung’s result. This result has a direct consequence to our problem, computing an upper bound on the query’s output size: we prove that Shannon inequalities are insufficient to compute a tight upper bound. In contrast, we show that the upper bound derived by using general information inequalities is tight, a result related to one by Gogacz and Torunczyk [GT17] (for cardinality constraints and functional dependencies only) and another one by Khamis et al. [KNS17] (for general degree constraints). The take-away of this section is that we have two upper bounds on the query’s output size: one that uses Shannon inequalities, which is computable but not always tight, and another one that uses general information inequalities, which is tight but whose computability is an open problem.

This motivates us to look in Sec. 6 at a special case, when the two bounds coincide and, thus, are both tight and computable. This special case is when the statistics are restricted to cardinalities, and to degrees on a single variable. We call the corresponding class of information inequalities *simple inequalities*, and prove that they are valid for all entropic vectors iff they are provable using Shannon inequalities. Moreover, in this special case, the worst-case database instances (where the size of the query’s output reaches the theoretical upper bound) have a simple yet interesting structure, called *normal database instances*, which generalize the *product database instances* that are the worst case instances for the AGM bound.

In Sec. 7 we turn to the most exciting application of upper bounds to the query’s output size: the design of Worst Case Optimal Join, WCOJ, algorithms, which compute a query in a time that does not exceed the upper bound on their output size. Thus, a WCOJ algorithm is *worst-case optimal*. The vast majority of database systems today compute a conjunctive query as a sequence of binary joins, whose intermediate results may exceed the upper bound on the final output size. Therefore, database execution engines are not WCOJ algorithms. For that reason, the discovery of the first WCOJ algorithm by Ngo, Porat, Ré, and Rudra [NPRR12, NPRR18] was a highly celebrated result. While the original WCOJ algorithm was complex, some of the same authors described a very simple WCOJ, called Generic Join (GJ) in [NRR13], which, together with its refinement Leapfrog Trie Join (LFTJ) [Vel14] forms the basis of the few implementations to date [SOC16, FBS⁺20, MKS21, WWS23]. Looking back at these results, we observe that any concrete WCOJ algorithm also provides a proof of the upper bound of the query’s output size, since the size of the output cannot exceed the runtime

of the algorithm. A WCOJ algorithm can be designed in reverse: start from a *proof* of the upper bound, then convert that proof into a WCOJ algorithm. We call this paradigm *from proofs to algorithms*, and illustrate it on three different proof systems for information inequalities: we derive GJ, an algorithm we call Heavy/Light, and PANDA.

Next, in Sec. 8 we move beyond upper bounds, and consider a related problem: given two queries, check whether the size of the output of the second query is always greater than or equal to that of the first query. This problem, called the *query domination problem*, is equivalent to the query containment problem under bag semantics. The latter was introduced by Chaudhuri and Vardi [CV93], is motivated by the semantics of SQL, where queries return duplicates, hence the answer to a query is a bag rather than a set. The query containment problem is: given two queries, interpreted under bag semantics, check whether the output of the first query is always contained in that of the second query. It has been shown that the containment problem is undecidable for *unions* of conjunctive queries [IR95] and for conjunctive queries with *inequalities* [JKV06], by reduction from Hilbert’s 10th problem. However, it remains an open problem to date whether the containment of two conjunctive queries is decidable. We describe in this section a surprising finding by Kopparty and Rossman [KR11], who have reduced the containment problem to information inequalities. This result was further extended in [KKNS21], and it was shown that the containment problem under bag semantics is computationally equivalent to information inequalities with max, which are inequalities that assert that the maximum of a finite number of linear expressions is ≥ 0 . The decidability of either of these problems remains open to date.

Finally, we present in Sec. 9 another, quite distinct application of information inequalities: reasoning about approximate integrity constraints. The implication problem for integrity constraints asks whether a set of integrity constraints logically implies some other constraint: this is a problem in Logic, and consists of checking the validity of a sentence $\bigwedge_i \sigma_i \Rightarrow \sigma$. When the integrity constraints can be captured by some information measures, such as is the case for Functional Dependencies and Multivalued Dependencies, then an implication can be described as a *conditional information inequality*. The problem we study is whether the exact implication problem can be *relaxed* to an inequality between these information measures, $\sum_i h(\sigma_i) \geq h(\sigma)$. We review a result from [KS22] stating that every exact implication between FDs and MVDs relaxes to an inequality. However, in a surprising result, Kaced and Romashchenko [KR13] have given examples of conditional information inequalities that do *not* relax. In other words, the exact implication holds, but the tiniest violation of an integrity constraint in the premise may cause arbitrarily large violation of the integrity constraint in the consequence. Yet in another turn, [KS22] show that *every* conditional information inequality relaxes with some error term, which can be made arbitrarily small, at the cost of increasing the coefficients of the terms representing the premise. In particular, every conditional inequality could be derived from an unconditioned inequality, by having the error term tend to zero, since in the conditional inequality the premise is assumed to be zero, hence the magnitudes of their coefficients do not matter. This section leads us to our deepest dive into the space of entropic vectors and almost entropic vectors: we show that the set of entropic vectors is neither convex nor a cone, that its topological closure is a convex cone, called the set of *almost entropic functions*, and use the theory of closed convex cones to prove the relaxation-with-error theorem.

Acknowledgments I am deeply indebted to my collaborators, especially Hung Q. Ngo who introduced me to applications of information inequalities to databases and with whom I had wonderful collaborations, and also to Mahmoud Abo Khamis, Batya Kenig, and Phokion G. Kolaitis. I also thank Dan Olteanu and Andrei Romashchenko for commenting on an early version of this paper.

2 Basic Notations

For two natural numbers M, N we denote by $[M : N] \stackrel{\text{def}}{=} \{M, M+1, \dots, N\}$; when $M = 1$ we abbreviate $[1 : N]$ by N . We will use upper case X, Y, Z for variable names, and lower case x, y, z for values of these variables. We use boldface for tuples of variables, e.g. \mathbf{X}, \mathbf{Y} , or tuples of values, e.g. \mathbf{x}, \mathbf{y} .

A **conjunctive query**, CQ, is an expression of the form:

$$Q(\mathbf{Y}_0) = \exists \mathbf{Z} (R_1(\mathbf{Y}_1) \wedge \dots \wedge R_m(\mathbf{Y}_m)) \quad (1)$$

Each $R_j(\mathbf{Y}_j)$ is called an *atom*: R_j is a relation name, and \mathbf{Y}_j are variables. We refer to \mathbf{Y}_j interchangeably as the *variables* of R_j , or the *attributes* of R_j . The variables \mathbf{Z} are called *existential variables*, while \mathbf{Y}_0 are called *head variables*. We denote by n the total number of variables in the query, and by $\mathbf{X} = \{X_1, \dots, X_n\}$ the set of these variables. Thus $\mathbf{X} = \mathbf{Y}_0 \cup \mathbf{Z}$, and $\mathbf{Y}_j \subseteq \mathbf{X}$, $\forall j$.

Fix some infinite domain Dom . If \mathbf{X} is a set of variables, then we write $\text{Dom}^{\mathbf{X}}$ for the set of \mathbf{X} -tuples. A *database instance* is $\mathbf{D} = (R_1^D, \dots, R_m^D)$, where, for each $j = 1, m$, $R_j^D \subseteq \text{Dom}^{\mathbf{Y}_j}$, where \mathbf{Y}_j are the attributes of R_j . Unless otherwise stated, relations are assumed to be finite. When \mathbf{D} is clear from the context, then we will drop the superscript and write simply R_j for the instance R_j^D , for $j = 1, m$.

We denote by $Q(\mathbf{D}) \subseteq \text{Dom}^{\mathbf{Y}_0}$ the output, or answer to the query Q on the database \mathbf{D} . The *query evaluation problem* is: given a database instance \mathbf{D} , compute the output $Q(\mathbf{D})$. The design and analysis of efficient query evaluation algorithms is a fundamental problem in database systems and database theory. For the complexity of the query evaluation problem, we consider only the *data complexity*, where Q is fixed, and the complexity is a function of the input database \mathbf{D} .

For a simple illustration, consider:

$$Q(X) = \exists Y \exists Z (R(X, Y) \wedge S(Y, Z) \wedge T(Z, X)) \quad (2)$$

Q returns all nodes x that belong to an *RST* triangle.

A *Boolean conjunctive query* is a conjunctive query with no head variables. At the other extreme, a *full conjunctive query* is a query with no existential variables. For example, the query:

$$Q(X, Y, Z) = R(X, Y) \wedge S(Y, Z) \wedge T(Z, X) \quad (3)$$

is a full CQ computing all triangles formed by the relations R, S, T . Full conjunctive queries are of special importance because they often occur as intermediate expressions during query evaluation. Unless otherwise stated, we will assume in this paper that the query is a full conjunctive query *without self-joins*, meaning that the relation names of the atoms R_1, R_2, \dots are distinct. Such a query is also called a natural join of the relations R_1, \dots, R_m .

Fix a relation $R(\mathbf{X})$, with n attributes. A **functional dependency**, or FD, is an expression $\mathbf{U} \rightarrow \mathbf{V}$, where $\mathbf{U}, \mathbf{V} \subseteq \mathbf{X}$. An instance R^D *satisfies* the FD, and we write $R^D \models \mathbf{U} \rightarrow \mathbf{V}$, if for any two tuples $\mathbf{x}_1, \mathbf{x}_2 \in R^D$, $\mathbf{x}_1.\mathbf{U} = \mathbf{x}_2.\mathbf{U}$ implies $\mathbf{x}_1.\mathbf{V} = \mathbf{x}_2.\mathbf{V}$. A set of functional dependencies Σ *implies* a functional dependency $\mathbf{U} \rightarrow \mathbf{V}$, in notation $\Sigma \models \mathbf{U} \rightarrow \mathbf{V}$, if, for every instance R^D , if $R^D \models \Sigma$ then $R^D \models \mathbf{U} \rightarrow \mathbf{V}$. Armstrong's axioms [AD80] form a complete axiomatization of the implication problem for FDs. The *closure* of $\mathbf{U} \subseteq \mathbf{X}$, denoted \mathbf{U}^+ , is the set of all attributes X s.t. $\Sigma \models \mathbf{U} \rightarrow X$. The closure can be computed in polynomial time in the size of \mathbf{U} and Σ . A set \mathbf{U} is *closed* if $\mathbf{U}^+ = \mathbf{U}$. A *super-key* for $R(\mathbf{X})$ is a set \mathbf{U} with the property that $\mathbf{U}^+ = \mathbf{X}$, and a *key* is a minimal set of attributes that is a superkey.

A finite **lattice** is a partially ordered set (L, \preceq) where every two elements $x, y \in L$ have a least upper bound $x \vee y$, and a greatest lower bound $x \wedge y$. In particular the lattice has a smallest

and a largest element, usually denoted by $\hat{0}, \hat{1}$. Consider now a set of variables \mathbf{X} , and a set of functional dependencies, Σ , over \mathbf{X} . We denote by (L_Σ, \subseteq) the lattice consisting of the closed sets, $L_\Sigma = \{\mathbf{U} \mid \mathbf{U}^+ = \mathbf{U}\}$. One can verify that the operations in this lattice are $\mathbf{U} \wedge \mathbf{V} \stackrel{\text{def}}{=} \mathbf{U} \cap \mathbf{V}$ and $\mathbf{U} \vee \mathbf{V} \stackrel{\text{def}}{=} (\mathbf{U} \cup \mathbf{V})^+$.

The **cartesian product** of two relations $R(\mathbf{X}), S(\mathbf{Y})$ with disjoint sets of attributes is the set $R \times S \stackrel{\text{def}}{=} \{(\mathbf{x}, \mathbf{y}) \mid \mathbf{x} \in R, \mathbf{y} \in S\}$ with attributes $\mathbf{X} \cup \mathbf{Y}$; its size is $|R \times S| = |R| \cdot |S|$. Fix a set of attributes \mathbf{X} , and two \mathbf{X} -tuples $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{x}' = (x'_1, \dots, x'_n)$. Their **domain product** is the \mathbf{X} -tuple $\mathbf{x} \otimes \mathbf{x}' \stackrel{\text{def}}{=} ((x_1, x'_1), \dots, (x_n, x'_n))$; thus, the value of each attribute is a pair.

Definition 2.1. *The domain product of two relation instances R and S , with the same set of attributes \mathbf{X} , is $R \otimes S \stackrel{\text{def}}{=} \{\mathbf{x} \otimes \mathbf{x}' \mid \mathbf{x} \in R, \mathbf{x}' \in S\}$.*

We have $|R \otimes S| = |R| \cdot |S|$. If $\mathbf{D}_i = (R_1^{D_i}, \dots, R_m^{D_i})$, $i = 1, 2$, are two database instances over the same schema, then we define their domain product $\mathbf{D}_1 \otimes \mathbf{D}_2$ as $(R_1^{D_1} \otimes R_1^{D_2}, \dots, R_m^{D_1} \otimes R_m^{D_2})$. One can check that $Q(\mathbf{D}_1 \otimes \mathbf{D}_2) = Q(\mathbf{D}_1) \otimes Q(\mathbf{D}_2)$ for any conjunctive query Q . The domain product should not be confused with the cartesian product. It was first introduced by Fagin [Fag82] (under the name *direct product*) to prove the existence of an Armstrong relation for constraints defined by Horn clauses, and later used by Geiger and Pearl [GP93] to prove that Conditional Independence constraints on probability distributions also admit an Armstrong relation. The same construction appears under the name “fibered product” in [KR11].

3 Warmup: the AGM Bound

Consider a full conjunctive query:

$$Q(\mathbf{X}) = \bigwedge_{j=1, m} R_j(\mathbf{Y}_j) \quad (4)$$

where $\mathbf{X} = \{X_1, \dots, X_n\}$. Assume we have a database \mathbf{D} , and we know the cardinality of each relation R_j^D . How large could the query output be? The answer is given by an elegant result, initially formulated by Grohe and Marx [GM06] and later refined by Atserias, Grohe, and Marx [AGM13], and is called today the AGM bound of the query Q . To state this bound, we first need to review the connection between conjunctive queries and hypergraphs.

We associate Q in (4) with the hypergraph $\mathcal{H} = (\mathbf{X}, E)$, where $E = \{\mathbf{Y}_1, \dots, \mathbf{Y}_m\}$. In other words, the nodes of the hypergraph are the variables, and its hyperedges are the atoms of the query. A *fractional edge cover* of the hypergraph \mathcal{H} is a tuple of non-negative weights $\mathbf{w} = (w_j)_{j=1, m}$, such that every variable X_i is *covered*, meaning:

$$\forall i = 1, n : \sum_{j: X_i \in \mathbf{Y}_j} w_j \geq 1 \quad (5)$$

A fractional edge cover of the query Q is a fractional edge cover of its associated hypergraph. The AGM bound is the following:

Theorem 3.1 (AGM Bound). *For any fractional edge cover \mathbf{w} of the query (4), and every instance \mathbf{D} :*

$$|Q(\mathbf{D})| \leq \prod_{j=1, m} |R_j^D|^{w_j} \quad (6)$$

To reduce clutter, we will often drop \mathbf{D} from both $Q(\mathbf{D})$ and $R_j^{\mathbf{D}}$, and write the bound simply as $|Q| \leq \prod_j |R_j|^{w_j}$.

Let $\mathbf{B} = (B_j)_{j=1,m}$ be a non-negative vector, representing the cardinalities of the relations in the database. We define:

$$AGM(Q, \mathbf{B}) \stackrel{\text{def}}{=} \min_{\mathbf{w}} \prod_{j=1,m} B_j^{w_j} \quad (7)$$

where \mathbf{w} ranges over all fractional edge covers of the query's hypergraph. Then Theorem 3.1 can be restated as follows: for every instance \mathbf{D} , if $|R_j^{\mathbf{D}}| \leq B_j$ for $j = 1, m$, then $|Q(\mathbf{D})| \leq AGM(Q, \mathbf{B})$. When \mathbf{B} is clear from the context, then we write the bound simply as $AGM(Q)$.

Before we prove the bound, we illustrate it with a classic example.

Example 3.2. Consider the triangle query (3), which we repeat here: $Q(X, Y, Z) = R(X, Y) \wedge S(Y, Z) \wedge T(Z, X)$. Its associated hypergraph is a graph with three nodes X, Y, Z and three edges forming a triangle. A fractional edge cover is any non-negative tuple (w_R, w_S, w_T) satisfying:

$$\begin{array}{llll} \text{Cover } X: & w_R + & w_T & \geq 1 \\ \text{Cover } Y: & w_R + & w_S & \geq 1 \\ \text{Cover } Z: & w_S + & w_T & \geq 1 \end{array}$$

The inequality $|Q| \leq |R|^{w_R} \cdot |S|^{w_S} \cdot |T|^{w_T}$ holds for every fractional edge cover. Consider the following four fractional edge covers: $(0, 1, 1), (1, 0, 1), (1, 1, 0), (1/2, 1/2, 1/2)$: these are the extreme vertices of the edge-covering polytope. It follows that the AGM bound in (7) is achieved at one of the four extreme vertices:

$$\begin{aligned} AGM(Q) = \\ \min \left(|S| \cdot |T|, |R| \cdot |T|, |R| \cdot |S|, |R|^{1/2} \cdot |S|^{1/2} \cdot |T|^{1/2} \right) \end{aligned}$$

When $|R| = |S| = |T| = N$ then $AGM(Q) = N^{3/2}$.

In the rest of this section we will prove the AGM bound (6), then show that the bound is tight.

Friedgut's Inequalities While the original proof of the AGM bound used information inequalities, we postpone the discussion of information inequalities until Sec. 4, where we consider more general statistics. Instead, we give here a simple, elementary proof, based on an elegant family of inequalities introduced by Friedgut [Fri04].

Fix a hypergraph $\mathcal{H} = (\mathbf{X}, E)$. Let $N > 0$ be a natural number, and for each hyperedge $\mathbf{Y}_j \in E$, let $r_j \in \mathbb{R}_+^{N^{|\mathbf{Y}_j|}}$ be a non-negative, multi-dimensional vector with $|\mathbf{Y}_j|$ dimensions; we will refer to r_j as a *tensor*. In what follows, we denote by \mathbf{i} a tuple $\mathbf{i} = (i_1, \dots, i_n) \in [N]^{\mathbf{X}}$, and by \mathbf{i}_j its projection on \mathbf{Y}_j .

Theorem 3.3 (Friedgut's Inequality). [Fri04] *For every fractional edge cover \mathbf{w} of the hypergraph \mathcal{H} , the following holds:*

$$\sum_{\mathbf{i}} \prod_{j=1,m} r_j[\mathbf{i}_j] \leq \prod_{j=1,m} \left(\sum_{\mathbf{i}_j} r_j[\mathbf{i}_j]^{\frac{1}{w_j}} \right)^{w_j} \quad (8)$$

Fig. 1 illustrates several instances of (8). We invite the reader to check that Loomis–Whitney's inequality [LW49] is also an instance such an inequality. Using Theorem 3.3 we can prove the AGM




Figure 1: Examples of Friedgut’s inequalities (8). In each case we show the associated hypergraph on the right.

bound as follows. Given a relational instance $\mathbf{D} = (R_1^D, \dots, R_m^D)$ define the following tensors:

$$r_j[x_1, \dots, x_{a_j}] \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } (x_1, \dots, x_{a_j}) \in R_j^D \\ 0 & \text{otherwise} \end{cases}$$

Then the LHS of (8) is $|Q(\mathbf{D})|$ and the RHS is $\prod_j |R_j|^{w_j}$.

Proof. (of Theorem 3.3) While the original proof also used information inequalities, we give here a direct proof, by induction on the number n of vertices of the hypergraph \mathcal{H} . (This proof generalizes Loomis–Whitney’s proof in [LW49].)

We replace each tensor expression $r_j[\mathbf{i}_j]$ with $(r_j[\mathbf{i}_j])^{w_j}$, then in order to prove (8) it suffices to prove:

$$\sum_{\mathbf{i}} \prod_j r_j[\mathbf{i}_j]^{w_j} \leq \prod_j \left(\sum_{\mathbf{i}_j} r_j[\mathbf{i}_j] \right)^{w_j} \quad (9)$$

We notice that the index variables $\mathbf{i} = (i_1, \dots, i_n)$ used in the summation correspond one-to-one to the nodes of the hypergraph $\mathbf{X} = \{X_1, \dots, X_n\}$, and the subset \mathbf{i}_j contains the index variables corresponding to nodes in \mathbf{Y}_j . We now prove (9) by induction on n .

When $n = 1$ then this is Hölder’s inequality (see Fig. 1), whose proof can be found in textbooks. Assume $n > 1$ and consider the hypergraph \mathcal{H}' obtained by removing the last variable X_n : its nodes are $\{X_1, \dots, X_{n-1}\}$ and its hyperedges are $\{\mathbf{Y}_j - \{X_n\} \mid j = 1, m\}$. The weights w_1, \dots, w_m continue to be a fractional edge cover for \mathcal{H}' . Group the LHS of Eq. (9) by factoring out the sum over the variable i_n , and apply induction hypothesis to the summation over the other variables i_1, \dots, i_{n-1} , which form the hypergraph \mathcal{H}' :

$$\sum_{i_n} \left(\sum_{i_1, \dots, i_{n-1}} \prod_j r_j[\mathbf{i}_j]^{w_j} \right) \leq \sum_{i_n} \prod_j \left(\sum_{\mathbf{i}_j - \{i_n\}} r_j[\mathbf{i}_j] \right)^{w_j}$$

We factor out the products that do not depend on the variable i_n , then use the fact that $\sum_{j: X_n \in \mathbf{Y}_j} w_j \geq 1$ because X_n is covered, and apply Hölder’s inequality (Fig. 1) with $a_j[i_n] \stackrel{\text{def}}{=} r_j[\mathbf{i}_j]$.

The RHS of the expression above becomes:

$$\begin{aligned} & \prod_{j: X_n \notin Y_j} \left(\sum_{\mathbf{i}_j} r_j[\mathbf{i}_j] \right)^{w_j} \cdot \sum_{i_n} \prod_{j: X_n \in Y_j} \left(\sum_{\mathbf{i}_j - \{i_n\}} r_j[\mathbf{i}_j] \right)^{w_j} \\ & \leq \prod_{j: X_n \notin Y_j} \left(\sum_{\mathbf{i}_j} r_j[\mathbf{i}_j] \right)^{w_j} \cdot \prod_{j: X_n \in Y_j} \left(\sum_{i_n} \sum_{\mathbf{i}_j - \{i_n\}} r_j[\mathbf{i}_j] \right)^{w_j} \end{aligned}$$

This is the RHS of (9), which completes the proof. \square

The lower bound How tight is the AGM bound? One key insight in [AGM13] is that, while the upper bound is described by a linear program, a lower bound can be described using the dual linear program: tightness follows from the strong duality theorem for linear programs. They proved:

Theorem 3.4. *For any query Q with n variables, and vector \mathbf{B} there exists a database \mathbf{D} s.t. $|Q(\mathbf{D})| \geq \frac{1}{2^n} \text{AGM}(Q, \mathbf{B})$. We call such a database \mathbf{D} a worst-case instance.*

Proof. The logarithm of the AGM bound (7) is the optimal value of the following *primal linear program*:

$$\begin{aligned} & \text{minimize} && \sum_j w_j \log B_j \\ & \text{where} && \forall i : \sum_{j: X_i \in Y_j} w_j \geq 1 \\ & && \forall j : w_j \geq 0 \end{aligned}$$

The *dual linear program* is:

$$\begin{aligned} & \text{maximize} && \sum_i v_i \\ & \text{where} && \forall j : \sum_{i: X_i \in Y_j} v_i \leq \log B_j \\ & && \forall i : v_i \geq 0 \end{aligned}$$

For any two feasible solutions \mathbf{w}, \mathbf{v} of the primal and dual, *weak duality* holds: $\sum_j w_j \log B_j \geq \sum_i v_i$. If $\mathbf{w}^*, \mathbf{v}^*$ are the optimal solutions, then the *strong duality* theorem states that these two expressions are equal, therefore:

$$\text{AGM}(Q, \mathbf{B}) = 2^{\sum_j w_j^* \log B_j} = 2^{\sum_i v_i^*} = \prod_i 2^{v_i^*} \quad (10)$$

If \mathbf{v} is any dual solution, we construct the following database instance \mathbf{D} : for each variable X_i , define the domain $V_i \stackrel{\text{def}}{=} [\lfloor 2^{v_i} \rfloor] = \{1, 2, \dots, \lfloor 2^{v_i} \rfloor\}$, and set $R_j^D \stackrel{\text{def}}{=} \times_{i: X_i \in Y_j} V_i$, for $j = 1, m$. We call \mathbf{D} a *product database instance*, because each relation is a cartesian product. \mathbf{D} satisfies the cardinality constraints $|R_j^D| \leq B_j$ because

$$|R_j^D| = \prod_{i: X_i \in Y_j} \lfloor 2^{v_i} \rfloor \leq 2^{\sum_{i: X_i \in Y_j} v_i} \leq 2^{\log B_j} = B_j$$

Similarly, the output to the query is the product $Q(\mathbf{D}) = \times_i V_i$, and its size is $\prod \lfloor 2^{v_i} \rfloor$. At optimality, when $\mathbf{v} = \mathbf{v}^*$,

$$|Q(\mathbf{D})| = \prod_i \lfloor 2^{v_i^*} \rfloor \quad (11)$$

Theorem 3.4 follows from (10) and (11), and observing that $\lfloor 2^{v_i^*} \rfloor \geq \frac{1}{2} 2^{v_i^*}$. \square

Thus, one could say that the AGM bound is tight up to a “rounding error”. The original paper [AGM13] provides an extensive discussion on tightness and proves two facts. First, they construct arbitrarily large databases \mathbf{D} where the AGM bound is tight exactly. Second, they describe an example where the ratio between the lower and upper bound can be arbitrarily close to $1/2^n$, as n grow arbitrarily large; despite this example, the AGM bound is considered to be tight for practical purposes.

Discussion The AGM bound is elegant in that it solves completely the problem it set out to solve: find the tight upper bound when the cardinalities of all relations are known, and nothing else is known. However, the bound is limited, in that it cannot take advantage of other statistics or constraints on the input data, which are often available in practice. For example, consider the join of two relations, $Q(X, Y, Z) = R(X, Y) \wedge S(Y, Z)$, and assume that both $|R|, |S| \leq B$. The AGM bound is $|Q| \leq B^2$ (because the only fractional edge cover is $(1, 1)$), and the reader can check that this is tight, i.e. there exists relations where $|R| = |S| = B$ and $|Q| = B^2$. But, in practice, joins are often key/foreign-key joins, for example, $S.Y$ may be a key in S , and in that case $|Q| \leq B$, because every tuple in R joins with at most one tuple in S . In order to account for additional information about the data, like keys or constraints on degrees, we need to use a more powerful tool than Friedgut’s inequalities: information inequalities.

4 Max-Degree Query Bounds

We describe now the general framework for computing an upper bound for the query output size, using information inequalities. We will use the cardinalities of the input relations (like in the AGM bound), keys or, more generally, functional dependencies for individual relations, and bounds on degrees, also called maximum frequencies, which generalize keys and functional dependencies. This section is based largely on [GLVV12, KNS16, KNS17]. We start with a short review of information inequalities.

4.1 Background on Information Inequalities

Consider a finite probability distribution (D, p) , where $p : D \rightarrow [0, 1]$, $\sum_{x \in D} p(x) = 1$. We denote by X the random variable with outcomes in D , and define its *entropy* as:

$$H(X) \stackrel{\text{def}}{=} - \sum_{x \in D} p(x) \log p(x) \quad (12)$$

If $N \stackrel{\text{def}}{=} |D|$, then $0 \leq H(X) \leq \log N$, the equality $H(X) = 0$ holds iff X is deterministic (i.e. $\exists x \in D, p(x) = 1$), and the equality $H(X) = \log N$ holds iff X is uniformly distributed (i.e. $p(x) = 1/N$ for all $x \in N$).

Consider now a finite probability distribution (R, p) , where $R \subseteq \text{Dom}^{\mathbf{X}}$ is a non-empty, finite relation with n attributes $\mathbf{X} = \{X_1, \dots, X_n\}$. We will always assume w.l.o.g. that R is the support of p , otherwise we just remove from R the tuples \mathbf{x} where $p(\mathbf{x}) = 0$. For each $\alpha \subseteq [n]$, define

X	Y	Z	
0	0	0	$p = 1/4$
0	1	1	$p = 1/4$
1	0	1	$p = 1/4$
1	1	0	$p = 1/4$




Figure 2: A relation defining the parity entropy \mathbf{h} . The marginal distribution of X is $p(X = 0) = p(X = 1) = 1/2$, hence its entropy is $h(X) = 1$, and similarly for the others values.

$X_\alpha \stackrel{\text{def}}{=} (X_i)_{i \in \alpha}$ the joint random variable obtained as follows: draw randomly a tuple $\mathbf{x} \in R$ with probability $p(\mathbf{x})$, then return \mathbf{x}_α . We associate the probability space (R, p) with a vector $\mathbf{h} \in \mathbb{R}_+^{2^{[n]}}$, by defining $\mathbf{h}_\alpha \stackrel{\text{def}}{=} H(\mathbf{X}_\alpha)$, and call \mathbf{h} an *entropic vector*. For any vector \mathbf{h} (entropic or not) we will write $h(\mathbf{X}_\alpha)$ for \mathbf{h}_α . In other words, we will blur the distinction between a vector in $\mathbb{R}_+^{2^{[n]}}$, a vector in $\mathbb{R}_+^{2^{\mathbf{X}}}$, and a function $2^{\mathbf{X}} \rightarrow \mathbb{R}_+$.

In analyzing properties of queries, we often examine the entropic vector derived from a uniform distribution.

Definition 4.1. *The uniform probability space associated to a non-empty, finite relation $R \subseteq \text{Dom}^{\mathbf{X}}$ is (R, p) where $p(\mathbf{x}) = 1/|R|$ for every tuple $\mathbf{x} \in R$. We will call its entropic vector \mathbf{h} uniform and say that it is associated to R .*

If \mathbf{h} is associated to R , then $h(\mathbf{X}) = \log |R|$, and $h(\mathbf{U}) \leq \log |\Pi_{\mathbf{U}}(R)|$ for every subset $\mathbf{U} \subseteq \mathbf{X}$. Fig. 2 illustrates the entropic vector \mathbf{h} associated to a relational instance R with attributes X, Y, Z ; we call \mathbf{h} the *parity function*, because the relation R contains all triples (x, y, z) that have an even number of 1.

Any entropic vector $\mathbf{h} \in \mathbb{R}_+^{2^{\mathbf{X}}}$ satisfies the following *basic Shannon inequalities*:

$$h(\emptyset) = 0 \tag{13}$$

$$h(\mathbf{U} \cup \mathbf{V}) \geq h(\mathbf{U}) \tag{14}$$

$$h(\mathbf{U}) + h(\mathbf{V}) \geq h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V}) \tag{15}$$

The last two inequalities are called *monotonicity* and *submodularity* respectively. A *Shannon inequality* is a positive linear combination of basic Shannon inequalities.

Any vector $h : 2^{\mathbf{X}} \rightarrow \mathbb{R}_+$ that satisfies the basic Shannon inequalities is called a *polymatroid*. The set of entropic vectors is denoted by Γ_n^* and the set of polymatroids is denoted by Γ_n , where $n = |\mathbf{X}|$ is the number of variables. The following holds: $\Gamma_n^* \subsetneq \Gamma_n \subsetneq \mathbb{R}_+^{2^{[n]}}$. In particular, not every polymatroid is entropic, as we will see shortly (in Fig. 5). Fig. 3 represents these two sets, as well as other sets, defined later in this paper. In some of the literature the entropic vectors and the polymatroids are defined as $(2^n - 1)$ -dimensional vectors, by dropping the \emptyset -dimension, because, in that case, both sets Γ_n^* and Γ_n have a non-empty interior. We prefer to use 2^n dimensions since this simplifies most of the discussion, and postpone dealing with the non-empty interior until Section 9.6.

Γ_n : polymatroids
 $\bar{\Gamma}_n^*$: almost-entropic
 Γ_n^* : entropic
 Υ_n : group realizable
 N_n : normal polymatroids
 M_n : modular polymatroids




Figure 3: Landscape of polymatroids

An information inequality is an assertion stating that a linear expression of entropic terms is ≥ 0 .

Definition 4.2. We associate to any vector $\mathbf{c} \in \mathbb{R}^{2^{[n]}}$ the following information inequality:

$$\sum_{\alpha \subseteq [n]} c_\alpha h(\mathbf{X}_\alpha) \geq 0 \quad (16)$$

By using the dot-product notation, we can write the inequality as $\mathbf{c} \cdot \mathbf{h} \geq 0$. If the inequality holds for all $\mathbf{h} \in K$, where $K \subseteq \mathbb{R}_+^{2^{[n]}}$ is some set, then we say that it is valid for K , and write $K \models \mathbf{c} \cdot \mathbf{h} \geq 0$.

Thus, we will talk about inequalities valid for entropic vectors, or valid for polymatroids, and the latter are precisely the Shannon inequalities (this is implicit in the proof of Th. 5.2 below). Any Shannon inequality is also valid for entropic vectors; however, we will see in Th. 5.7 below a non-Shannon inequality, which is valid for entropic vectors, but not for polymatroids. In analogy with mathematical logic, one should think the vectors \mathbf{h} as *models*, inequalities $\mathbf{c} \cdot \mathbf{h} \geq 0$ as *formulas*, and sets $K \subseteq \mathbb{R}_+^{2^{[n]}}$ as *classes of models*.

Example 4.3. The following is a Shannon inequality, called Shearer's inequality:

$$h(XY) + h(YZ) + h(ZX) - 2h(XYZ) \geq 0 \quad (17)$$

To prove it, we apply submodularity twice, underlining the affected terms:

$$\begin{aligned}
 & \underline{h(XY)} + h(YZ) + h(ZX) \\
 & \geq h(XYZ) + \underline{h(Y)} + h(ZX) \\
 & \geq h(XYZ) + h(XYZ) + h(\emptyset) = 2h(XYZ)
 \end{aligned}$$

Equivalently, we observe that (17) is the sum of the following basic Shannon inequalities:

$$\begin{aligned}
 h(XY) + h(YZ) - h(Y) - h(XYZ) &\geq 0 \\
 h(Y) + h(ZX) - h(\emptyset) - h(XYZ) &\geq 0 \\
 h(\emptyset) &= 0
 \end{aligned}$$

We will prove shortly (Theorem 5.2 below) that one can decide in time exponential in n whether an inequality is valid for all polymatroids. In contrast, it is an open problem whether entropic validity is decidable.

4.2 The Entropic Bound and the Polymatroid Bound

The general framework for computing a bound on a query's output uses degree constraints, which, in turn, correspond to conditional entropies. We define these two notions first.

We write UV for set $U \cup V$. Given $\mathbf{h} \in \mathbb{R}_+^{2^{[n]}}$, define:

$$h(\mathbf{V}|\mathbf{U}) \stackrel{\text{def}}{=} h(UV) - h(\mathbf{U}) \quad (18)$$

\mathbf{U}, \mathbf{V} need not be disjoint, and $h(\mathbf{V}|\mathbf{U}) = h(\mathbf{V} - \mathbf{U}|\mathbf{U})$; for example, $h(XY|X) = h(Y|X)$. If $h(\mathbf{V}|\mathbf{U}) = 0$ then we say that \mathbf{h} satisfies the functional dependency $\mathbf{U} \rightarrow \mathbf{V}$, and we write $\mathbf{h} \models \mathbf{U} \rightarrow \mathbf{V}$. Lee [Lee87a] proved that, if R is a relation instance with attributes \mathbf{X} , $p : R \rightarrow [0, 1]$ is a probability distribution, and \mathbf{h} is its entropic vector, then $R \models \mathbf{U} \rightarrow \mathbf{V}$ iff $\mathbf{h} \models \mathbf{U} \rightarrow \mathbf{V}$. For a simple illustration, referring to Fig. 2, both R and its entropy \mathbf{h} satisfy the FDs $XY \rightarrow Z$, $XZ \rightarrow Y$, and $YZ \rightarrow X$: for example XY is a key (all 4 tuples have distinct values XY) and $h(Z|XY) = h(XYZ) - h(XY) = 2 - 2 = 0$.

Fix \mathbf{U} , and denote by $h(-|\mathbf{U}) : 2^{\mathbf{X}-\mathbf{U}} \rightarrow \mathbb{R}_+$ the function $\mathbf{V} \mapsto h(\mathbf{V}|\mathbf{U})$. If \mathbf{h} is a polymatroid, then $h(-|\mathbf{U})$ is also a polymatroid, called the conditional polymatroid. If \mathbf{h} is an entropic vector, then, surprisingly, $\mathbf{h}(-|\mathbf{U})$ is not necessarily entropic (as we will see later in Sec. 9.3), yet the name *conditional entropy* is justified by the following. Suppose \mathbf{h} is associated to (R, p) . Fix an outcome $\mathbf{u} \in \text{Dom}^{\mathbf{U}}$, consider the random variable \mathbf{V} conditioned on $\mathbf{U} = \mathbf{u}$, and denote its entropy by $h(\mathbf{V}|\mathbf{U} = \mathbf{u})$. Then:

$$h(\mathbf{V}|\mathbf{U}) = \mathbb{E}_{\mathbf{u}} [h(\mathbf{V}|\mathbf{U} = \mathbf{u})] \quad (19)$$

In other words, $h(\mathbf{V}|\mathbf{U})$ equals the *expectation* over the outcomes \mathbf{u} of the (standard) entropy of the random variable \mathbf{V} conditioned on $\mathbf{U} = \mathbf{u}$. The proof of identity (19) consists of applying directly the definition of the entropy given in Eq. (12).

When proving Shannon inequalities it is sometimes convenient to write the submodularity inequality as $h(\mathbf{V}|\mathbf{U}) \geq h(\mathbf{V}|\mathbf{UW})$.¹ In other words, conditioning on more variables can only decrease the entropy.

Example 4.4. *We illustrate a simple Shannon inequality with conditionals:*

$$\begin{aligned} & \underline{h(XY) + h(YZ)} + h(ZU) + h(U|XZ) + h(X|YU) \geq \\ & \geq h(XYZ) + \underline{h(Y) + h(ZU)} + h(U|XZ) + h(X|YU) \\ & \geq h(XYZ) + h(YZU) + \underline{h(U|XZ)} + \underline{h(X|YU)} \\ & \geq h(XYZ) + h(YZU) + h(U|XYZ) + h(X|YZU) \\ & = 2h(XYZU) \end{aligned}$$

Next, we define degrees of a relation instance $R \subseteq \text{Dom}^{\mathbf{X}}$. Given subsets $\mathbf{U}, \mathbf{V} \subseteq \mathbf{X}$, and $\mathbf{u} \in \text{Dom}^{\mathbf{U}}$, the *V-degree of $\mathbf{U} = \mathbf{u}$ in R* is the number of distinct values \mathbf{v} that occur in R together with \mathbf{u} ; the *max-V-degree of \mathbf{U}* is the maximum degree over all values \mathbf{u} . Formally:

$$\begin{aligned} \deg_R(\mathbf{V}|\mathbf{U} = \mathbf{u}) & \stackrel{\text{def}}{=} |\{\mathbf{v} \mid (\mathbf{u}, \mathbf{v}) \in \Pi_{UV}(R)\}| \\ \deg_R(\mathbf{V}|\mathbf{U}) & \stackrel{\text{def}}{=} \max_{\mathbf{u}} (\deg_R(\mathbf{V}|\mathbf{U} = \mathbf{u})) \end{aligned}$$

We note that $\deg_R(\mathbf{V}|\mathbf{U}) \geq 1$ (since we assumed $R \neq \emptyset$), and equality holds iff R satisfies the functional dependency $\mathbf{U} \rightarrow \mathbf{V}$.

¹This is equivalent to $h(UV) - h(\mathbf{U}) \geq h(UVW) - h(UW)$; when $\mathbf{V} \cap \mathbf{W} = \emptyset$ then this is a submodularity inequality.

Definition 4.5. Fix a relation $R(\mathbf{X})$. A degree statistics, or a statistics in short, σ , is an expression of the form $\sigma = (\mathbf{V}|\mathbf{U})$ where $\mathbf{U}, \mathbf{V} \subseteq \mathbf{X}$; when $\mathbf{U} = \emptyset$ then we call σ a cardinality statistics, and write it as (\mathbf{V}) . If Σ is a set of statistics, then we call $\mathbf{B} = (B_\sigma)_{\sigma \in \Sigma}$, where $B_\sigma \geq 1$, statistics values associated to Σ . The log-statistics values are $\mathbf{b} = \log \mathbf{B} = (b_\sigma := \log B_\sigma)_{\sigma \in \Sigma}$.

We abbreviate $h(\mathbf{V}|\mathbf{U})$ and $\deg_R(\mathbf{V}|\mathbf{U})$ with $h(\sigma)$ and $\deg_R(\sigma)$ respectively. If Σ is a set of statistics, then a Σ -inequality is an inequality of the following form:

$$\sum_{\sigma \in \Sigma} w_\sigma h(\sigma) \geq h(\mathbf{X}) \quad (20)$$

where $\mathbf{w} = (w_\sigma)_{\sigma \in \Sigma}$ are nonnegative weights.

Fix a hypergraph $\mathcal{H} = (\mathbf{X}, E)$. We say that Σ is *guarded by* \mathcal{H} if, for every $\sigma = (\mathbf{V}|\mathbf{U})$ in Σ , there exists a hyperedge $\mathbf{Y}_\sigma \in E$ such that $\mathbf{U}, \mathbf{V} \subseteq \mathbf{Y}_\sigma$; we call \mathbf{Y}_σ the *guard* of σ . When \mathcal{H} is the hypergraph of a query Q , then we say that Σ is guarded by Q , and that R_σ is the guard of σ . The following theorem establishes the key connection between information inequalities and query output size. The proof relies on a method originally introduced by Chung et al. for a combinatorial problem [CGFS86], and adapted by Grohe and Marx for constraint satisfaction [GM14], then by Atserias, Grohe, and Marx for their AGM bound [AGM13].

Theorem 4.6. Assume Σ is guarded by Q . If the Σ -inequality (20) is valid for entropic vectors then:

$$|Q| \leq \prod_{\sigma \in \Sigma} \deg_{R_\sigma}^{w_\sigma}(\sigma) \quad (21)$$

where R_σ is the guard of $\sigma \in \Sigma$.

Proof. Fix a database instance \mathbf{D} , and let \mathbf{h} be the entropic vector associated to the relation $Q(\mathbf{D})$; by uniformity, $h(\mathbf{X}) = \log |Q(\mathbf{D})|$. If $\sigma = (\mathbf{V}|\mathbf{U}) \in \Sigma$ has guard R_σ , then:

$$\begin{aligned} h(\sigma) &= \mathbb{E}_{\mathbf{u}} [h(\mathbf{V}|\mathbf{U} = \mathbf{u})] \leq \max_{\mathbf{u}} h(\mathbf{V}|\mathbf{U} = \mathbf{u}) \\ &\leq \max_{\mathbf{u}} \log \deg_{R_\sigma}(\mathbf{V}|\mathbf{U} = \mathbf{u}) \\ &= \log \deg_{R_\sigma}(\mathbf{V}|\mathbf{U}) = \log \deg_{R_\sigma}(\sigma) \end{aligned}$$

Using (20) we derive:

$$\sum_{\sigma} w_\sigma \log \deg_{R_\sigma}(\sigma) \geq \sum_{\sigma} w_\sigma h(\sigma) \geq h(\mathbf{X}) = \log |Q|$$

□

Inequality (20) is similar to the AGM inequality (6). Next, we proceed as we did for the AGM bound: fix numerical values for the statistics, then minimize the bound over all valid \mathbf{w} 's. We say that a database instance \mathbf{D} *satisfies* the statistics Σ, \mathbf{B} , in notation $\mathbf{D} \models (\Sigma, \mathbf{B})$, if $\deg_{R_\sigma}^{\mathbf{D}}(\sigma) \leq B_\sigma$ for all $\sigma \in \Sigma$. Similarly, we say that a vector \mathbf{h} satisfies the log-statistics Σ, \mathbf{b} if $h(\sigma) \leq b_\sigma$ for all σ . We define:

Definition 4.7 (Query Upper Bound). Let Σ, \mathbf{B} be statistics values, guarded by the query Q . Fix some set $K \subseteq \mathbb{R}^{2^{[n]}}$. The Upper Bound w.r.t. K of the query Q is:

$$U\text{-Bound}_K(Q, \Sigma, \mathbf{B}) \stackrel{\text{def}}{=} \inf_{\mathbf{w}: K \models \text{Eq. (20)}} \prod_{\sigma \in \Sigma} B_\sigma^{w_\sigma}$$

The entropic upper bound is $U\text{-Bound}_{\Gamma_n^*}$ and the polymatroid upper bound is $U\text{-Bound}_{\Gamma_n}$.

Sometimes it is more convenient to use the logarithm and define:

$$\text{Log-U-Bound}_K(Q, \Sigma, \mathbf{b}) \stackrel{\text{def}}{=} \inf_{\mathbf{w}: K \models \text{Eq. (20)}} \sum_{\sigma \in \Sigma} w_\sigma b_\sigma \quad (22)$$

Corollary 4.8. *The following hold:*

- $U\text{-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B}) \leq U\text{-Bound}_{\Gamma_n}(Q, \Sigma, \mathbf{B})$, and
- If $\mathbf{D} \models (\Sigma, \mathbf{B})$ then $|Q(\mathbf{D})| \leq U\text{-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B})$.

The first item is by $\Gamma_n^* \subseteq \Gamma_n$, the second is by Th. 4.6.

Let's compare these bound with the AGM bound (7). There, we had to minimize an expression where \mathbf{w} ranged over the fractional edge covers of the query's hypergraph. In our new setting, \mathbf{w} ranges over valid Σ -inequalities, a much more difficult task. To compute the polymatroid bound, \mathbf{w} ranges over Σ -inequalities valid for polymatroids, and we will show in Th. 5.2 that this bound can be computed in time exponential in n . However, in order to compute the entropic bound, \mathbf{w} needs to define a valid entropic inequality, and it is currently open whether this bound is computable. On the other hand, we will prove that the entropic bound is asymptotically tight, while the polymatroid bound is not. Thus, we are faced with a difficult choice, between an exact but non-computable bound, or a computable but inexact bound. This justifies examining non-trivial special cases of statistics Σ when these two bounds agree. We illustrate the entropic upper bound with an example.

Example 4.9. *Consider the following conjunctive query:*

$$Q(X, Y, Z, U) = R(X, Y) \wedge S(Y, Z) \wedge T(Z, U) \\ \wedge A(X, Z, U) \wedge B(X, Y, U)$$

Suppose that we are given the following set of statistics $\Sigma = \{(XY), (YZ), (ZU), (U|XZ), (X|YU)\}$. In other words, we have bounds on the cardinalities of R, S, T , but not of A, B , hence we can assume that $|A| = |B| = \infty$. Instead, we have the statistics $\deg_A(U|XZ)$ and $\deg_B(X|YU)$. The AGM bound (6) is $|Q| \leq |R| \cdot |T|$, because the only fractional edge cover whose bound is $< \infty$ is $w_R = w_T = 1$ and $w_S = w_A = w_B = 0$.

The polymatroid bound follows from these Σ -inequalities:

$$\begin{aligned} h(XY) + h(YZ) + h(ZU) + h(U|XZ) + h(X|YU) &\geq \\ &\geq 2h(XYZU) \\ h(XY) + h(ZU) &\geq h(XYZU) \\ h(XY) + h(YZ) + h(U|XZ) &\geq h(XYZU) \\ h(YZ) + h(ZU) + h(X|YU) &\geq h(XYZU) \end{aligned}$$

We proved the first inequality in Example 4.4, while the other three are immediate. They imply:

$$\begin{aligned} |Q| &\leq (|R| \cdot |S| \cdot |T| \cdot \deg_A(U|XZ) \cdot \deg_B(X|YU))^{1/2} \\ |Q| &\leq |R| \cdot |T| \\ |Q| &\leq |R| \cdot |S| \cdot \deg_B(X|YU) \\ |Q| &\leq |S| \cdot |T| \cdot \deg_A(U|XZ) \end{aligned}$$

The AGM bound is the second inequality. We show in Appendix .1 that the entropic bound is the minimum over all four expressions above. (This requires proving that there is no "better" inequality

$$R = \begin{array}{c|cccccc} X_1 \dots X_{i-1} & X_i & X_{i+1} \dots X_n & & p \\ \hline 0 & \dots & 0 & 0 & 0 & \dots & 0 & 1/2 \\ 0 & \dots & 0 & 1 & 0 & \dots & 0 & 1/2 \end{array} \quad h^{X_i}(\mathbf{U}) = \begin{cases} 1 & \text{if } X_i \in \mathbf{U} \\ 0 & \text{if } X_i \notin \mathbf{U} \end{cases}$$

Figure 4: A relation R with two tuples that agree on all attributes, except X_i . Its entropic vector is called the *basic modular function*, \mathbf{h}^{X_i} ; it is used in Th. 4.10, and discussed in more detail in Sec. 6.

that gives us a smaller bound.) Since all four inequalities are Shannon inequalities, it follows that, in this case, the entropic bound is equal to the polymatroid bound.

When XZ is a key in A , and YU is a key in B , then the polymatroid bound simplifies to:

$$|Q| \leq \min((|R| \cdot |S| \cdot |T|)^{1/2}, |R| \cdot |T|, |R| \cdot |S|, |S| \cdot |T|)$$

Special Case: Cardinality Constraints We show next that the AGM bound is a special case where the polymatroid and entropic bounds coincide; we will see a more general setting when this happens in Sec. 6. Assume Σ is restricted to cardinality constraints, and assume for simplicity that Σ has exactly one cardinality constraint (\mathbf{Y}_j) for each relation $R_j(\mathbf{Y}_j)$ in the query; $\Sigma = \{(\mathbf{Y}_1), \dots, (\mathbf{Y}_m)\}$. The Σ -inequality (20) becomes:

$$\sum_{j=1, m} w_j h(\mathbf{Y}_j) \geq h(\mathbf{X}) \quad (23)$$

When $w_1 = \dots = w_m$ then (23) is called Shearer's inequality.

Theorem 4.10. *The following are equivalent:*

- (i) *Inequality (23) is valid for polymatroids.*
- (ii) *Inequality (23) is valid for entropic functions.*
- (iii) *The weights \mathbf{w} form a fractional edge cover of the hypergraph $(\mathbf{X}, \{\mathbf{Y}_j \mid j = 1, m\})$.*

Proof. (i) \Rightarrow (ii) is immediate. For (ii) \Rightarrow (iii), assume that (23) holds for all entropic vectors, and consider any variable $X_i \in \mathbf{X}$, for $i = 1, n$. Consider the *basic modular entropic function* \mathbf{h}^{X_i} shown in Fig. 4. Since \mathbf{h}^{X_i} satisfies inequality (23), it follows that $\sum_{j=1, m: X_i \in \mathbf{Y}_j} w_j \geq 1$ (because $h^{X_i}(\mathbf{Y}_j) = 1$ iff $X_i \in \mathbf{Y}_j$), proving that \mathbf{w} is a fractional edge cover.

It remains to prove the implication (iii) \Rightarrow (i). This is a well known result, and it admits multiple proofs (we give a second proof in Sec. 7.2). Here, we will prove the inequality by induction on n :

- Partition the set of indices j into J_0 and J_1 :
 $J_0 \stackrel{\text{def}}{=} \{j \mid X_n \notin \mathbf{Y}_j\}$, $J_1 \stackrel{\text{def}}{=} \{j \mid X_n \in \mathbf{Y}_j\}$.
- If $j \in J_1$ then write $h(\mathbf{Y}_j) = h(X_n) + h(\mathbf{Y}_j - X_n | X_n)$. Note that $\sum_{j \in J_1} w_j \geq 1$ because X_n is covered.
- If $j \in J_0$ then write $h(\mathbf{Y}_j) \geq h(\mathbf{Y}_j - X_n | X_n)$.

Using the steps above we obtain:

$$\begin{aligned} \sum_j w_j h(\mathbf{Y}_j) &\geq h(X_n) + \sum_j w_j h(\mathbf{Y}_j - X_n | X_n) \\ &\geq h(X_n) + h(\mathbf{X} - X_n | X_n) = h(\mathbf{X}) \end{aligned}$$

The last line used induction on the polymatroid $h(-|X_n)$. \square

It follows immediately that the AGM bound, the entropic bound, and the polymatroid bound coincide in the simple case when the statistics are restricted to the cardinalities of the input relations.

Discussion Degree constraints occur often and naturally in database applications. For example, if a relation $R_j(\mathbf{Y}_j)$ has a key \mathbf{U} , then $\deg_{R_j}(\mathbf{Y}_j | \mathbf{U}) = 1$. In practice almost every relation has a key, so this case is very common. In other cases some cardinality constraints can be obtained directly from the application. For example, suppose that in a database of customers we require that no customer may have more than 10 credit cards, which naturally leads to a max-degree constraint. Such constraints are used in some modern systems, for example in *scale-independent query processing* [AFP⁺09, ACK⁺11, ALK⁺13].

5 The Worst-Case Instance

Informally, we call a database instance \mathbf{D} a *worst-case instance* if it satisfies the given statistics, and the query's output is as large as, or approaches asymptotically (in a sense to be made precise), the entropic upper bound. We will show that such a worst-case instance exists, proving that the entropic bound is *asymptotically tight*, which is a weaker notion of tightness than for the AGM bound in Th. 3.4. We will also show that, in general, the polymatroid bound is not tight, even for this weaker notion of tightness.

To construct the worst-case instance we need a *dual* definition of the entropic and polymatroid bounds. We define them directly using log-version:

Definition 5.1 (Query Lower Bound). *Fix log-statistics Σ, \mathbf{b} . For any set $K \subseteq \mathbb{R}^{2[n]}$, the Log Lower Bound w.r.t. K is:*

$$\text{Log-L-Bound}_K(Q, \Sigma, \mathbf{b}) \stackrel{\text{def}}{=} \sup_{\mathbf{h} \in K : \mathbf{h} \models (\Sigma, \mathbf{b})} h(\mathbf{X}) \quad (24)$$

As before, the entropic log-lower bound is $\text{Log-L-Bound}_{\Gamma_n^*}$, and the polymatroid log-lower bound is $\text{Log-L-Bound}_{\Gamma_n}$.

The log-lower bound asks us to find a vector \mathbf{h} that satisfies all log-statistics Σ, \mathbf{b} , and where $h(\mathbf{X})$ is as large as possible. We call \mathbf{h} a *worst case entropic vector*, or the worst case polymatroid respectively. Using \mathbf{h} , we would like to construct a worst-case database instance \mathbf{D} , that satisfies Σ, \mathbf{B} , and $\log |Q(\mathbf{D})| = h(\mathbf{X})$. The difficulty lies in the fact that, when \mathbf{h} is a polymatroid then such a database may not exist in general, and when \mathbf{h} is an entropic vector, then it may be realized by a probability space that is non-uniform, hence we cannot use it to construct \mathbf{D} . We start by observing the following, which is easy to check:

$$\text{Log-L-Bound}_K(Q, \Sigma, \mathbf{b}) \leq \text{Log-U-Bound}_K(Q, \Sigma, \mathbf{b}) \quad (25)$$

Indeed, if $\mathbf{h} \in K$ satisfies $\mathbf{h} \models (\Sigma, \mathbf{b})$, and \mathbf{w} satisfies $\forall \mathbf{h} \in K, \sum_{\sigma} w_{\sigma} h(\sigma) \geq h(\mathbf{X})$, then $h(\mathbf{X}) \leq \sum_{\sigma} w_{\sigma} h(\sigma) \leq \sum_{\sigma} w_{\sigma} b_{\sigma}$, and the claim follows from $\text{Log-L-Bound}_K = \sup_{\mathbf{h}} h(\mathbf{X}) \leq \text{Log-U-Bound}_K = \inf_{\mathbf{w}} \sum_{\sigma} w_{\sigma} b_{\sigma}$.

When $K = \Gamma_n$, then [KNS17] showed that the two bounds are equal. We prove a slightly more general statement:

Theorem 5.2. *Suppose the set K is defined by linear constraints: $K = \{\mathbf{h} \in \mathbb{R}_+^{2^n} \mid \mathbf{M} \cdot \mathbf{h} \geq 0\}$, where \mathbf{M} is some matrix.² Then, Log-L-Bound_K and Log-U-Bound_K are defined by a pair of primal/dual linear programs, with a number of variables exponential in n ; the expressions \inf, \sup in (22), (24) can be replaced by \min, \max ; and Eq. (25) becomes an equality, $h^*(\mathbf{X}) = \sum_{\sigma} w_{\sigma}^* b_{\sigma}$, where $\mathbf{h}^*, \mathbf{w}^*$ are the optimal solutions of the primal and dual program respectively.*

Proof. Denote $s = |\Sigma|$, and let \mathbf{A} be the $s \times 2^n$ matrix that maps \mathbf{h} to the vector $\mathbf{A} \cdot \mathbf{h} = (h(\sigma))_{\sigma \in \Sigma} \in \mathbb{R}^s$. Let $\mathbf{c} \in \mathbb{R}^{2^n}$ be the vector $c_{\mathbf{X}} = 1$, $c_{\mathbf{U}} = 0$ for $\mathbf{U} \neq \mathbf{X}$. The two bounds are the optimal solutions to the following pair of primal/dual linear programs:

Log-L-Bound_K Maximize $\mathbf{c}^T \cdot \mathbf{h}$ where $\mathbf{A} \cdot \mathbf{h} \leq \mathbf{b}$ $-\mathbf{M} \cdot \mathbf{h} \leq 0$	Log-U-Bound_K Minimize $\mathbf{w}^T \cdot \mathbf{b}$ where $\mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T \geq \mathbf{u}^T \cdot \mathbf{M}$
--	---

where the primal variables are $\mathbf{h} \geq 0$, and the dual variables are $\mathbf{w}, \mathbf{u} \geq 0$; the reader may check that the two programs above form indeed a primal/dual pair. Log-L-Bound_K is by definition the optimal value of the program above. We prove that Log-U-Bound_K is the value of the dual. First, observe that the Σ -inequality (20) is equivalent to $(\mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T) \cdot \mathbf{h} \geq 0$. We claim that this inequality holds $\forall \mathbf{h} \in K$ iff there exists \mathbf{u} s.t. (\mathbf{w}, \mathbf{u}) is a feasible solution to the dual. For that consider the following primal/dual programs with variables $\mathbf{h} \geq 0$ and $\mathbf{u} \geq 0$ respectively:

Minimize $(\mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T) \cdot \mathbf{h}$ where $\mathbf{M} \cdot \mathbf{h} \geq 0$	Maximize 0 where $\mathbf{u}^T \cdot \mathbf{M} \leq \mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T$
--	---

The primal (left) has optimal value 0 iff the inequality $(\mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T) \cdot \mathbf{h} \geq 0$ holds for all $\mathbf{h} \in K$; otherwise its optimal is $-\infty$. The dual (right) has optimal value 0 iff there exists a feasible solution \mathbf{u} ; otherwise its optimal is $-\infty$. Strong duality proves our claim. \square

When $K = \Gamma_n^*$, then the two terms in (25) may no longer be equal in general, but we prove that they are equal asymptotically. Call a k -*amplification* of a set of log-statistics Σ, \mathbf{b} the log-statistics $\Sigma, k\mathbf{b}$, where k is a natural number. Observe that the entropic log-upper bound increases linearly with the k -amplification:

$$\text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b}) = k \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b}) \quad (26)$$

The lower bound increases at least linearly, $\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b}) \geq k \text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b})$, because of the following proposition:

Proposition 5.3. *If $\mathbf{h}_1, \mathbf{h}_2$ are two entropic vectors, then so is $\mathbf{h}_1 + \mathbf{h}_2$.*

Proof. Suppose $\mathbf{h}_1, \mathbf{h}_2$ are realized by two finite probability spaces $(R_1, p_1), (R_2, p_2)$. Then their sum is realized by $(R_1 \otimes R_2, p)$ (see Def. 2.1), where $p(\mathbf{x}_1 \otimes \mathbf{x}_2) \stackrel{\text{def}}{=} p(\mathbf{x}_1) \cdot p(\mathbf{x}_2)$. \square

We prove:

²Equivalently: K is a polyhedral cone, reviewed in Sec. 9.3.

Theorem 5.4. *Fix any Q, Σ, \mathbf{b} . The entropic upper and lower bounds are asymptotically equal, in the following sense:*

$$\sup_k \frac{\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})}{\text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})} = 1 \quad (27)$$

The proof of this result, which appears to be novel, requires a discussion of the set of *almost-entropic functions*, and we defer this to Sec. 9.

Finally, we can answer the central question in this section: the entropic bound is tight asymptotically, while the polymatroid bound is not.

Theorem 5.5. *(1) For any query Q and statistics Σ, \mathbf{B} , the entropic bound is asymptotically tight, in the following sense:*

$$\sup_k \frac{\sup_{\mathbf{D}: \mathbf{D} \models \mathbf{B}^k} \log |Q(\mathbf{D})|}{\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})} = 1 \quad (28)$$

(2) The polymatroid bound is not asymptotically tight: there exists a query Q and statistics Σ, \mathbf{B} such that:

$$\sup_k \frac{\sup_{\mathbf{D}: \mathbf{D} \models \mathbf{B}^k} \log |Q(\mathbf{D})|}{\text{Log-L-Bound}_{\Gamma_n}(\mathbf{Q}, \Sigma, k\mathbf{b})} \leq \frac{43}{44} \quad (29)$$

Here $\mathbf{B}^k \stackrel{\text{def}}{=} (B_\sigma^k)_{\sigma \in \Sigma}$, and $\mathbf{b} \stackrel{\text{def}}{=} \log \mathbf{B}$. Moreover, this property holds even if Σ, \mathbf{B} consists only of cardinality constraints and functional dependencies, i.e. $\forall \sigma \in \Sigma$ either $\sigma = (\mathbf{V})$ or $B_\sigma = 1$.

Equivalently, Eq. (28) says that, for any statistics Σ, \mathbf{B} , if we allow a sufficiently large amplification Σ, \mathbf{B}^k , then there exists a worst-case instance \mathbf{D} satisfying the amplified statistics such that $\log |Q(\mathbf{D})|$ approaches $\text{L-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B}^k)$. Notice that this is a weaker notion of tightness than for the AGM bound in Theorem 3.4. There, tightness referred to the ratio between the lower and upper bound, while here tightness refers to the ratio of their logarithms, a weaker notion. Eq. (29) says that even this weaker tightness fails for the polymatroid bound.

The first proof of asymptotic tightness was given by Gogacz and Torunczyk [GT17], for the restricted case when the statistics are either cardinalities, or functional dependencies. The general case was proven in [KNS17]. Both results were stated slightly differently from ours, by using almost-entropic functions. We prefer to state our result in terms of the entropic functions, since it is more natural, and defer the discussion of almost entropic functions to Sec. 9, where they have a very natural justification.

In the rest of this section we prove Theorem 5.5. We use this opportunity to continue our dive into the fascinating world of entropic functions, and non-Shannon inequalities, which are needed for the proof. However, the rest of this section is rather technical, and readers not interested in this background may safely skip the rest of this section, since we do not need it, except for the short introduction of mutual information.

5.1 Background: Non-Shannon Inequalities, Lattices, Groups

Mutual Information Given a vector $\mathbf{h} \in \mathbb{R}^{2^{[n]}}$ and three disjoint sets of variables $U, V, W \subseteq X$, we denote by:

$$I_h(V; W|U) \stackrel{\text{def}}{=} h(UV) + h(UW) - h(U) - h(UVV) \quad (30)$$

When h is clear from the context, then we will drop the index h from I_h and write simply I . When h is an entropic vector, then $I_h(\mathbf{V}; \mathbf{W}|\mathbf{U})$ is called the *mutual information of \mathbf{V}, \mathbf{W} conditioned on \mathbf{U}* . In that case, $I_h(\mathbf{V}; \mathbf{W}|\mathbf{U}) = 0$ iff the probability space realizing h satisfies $\mathbf{V} \perp \mathbf{W}|\mathbf{U}$, meaning that \mathbf{V}, \mathbf{W} are independent conditioned on \mathbf{U} . With some abuse, we will call I_h a conditional mutual information even when \mathbf{h} is a polymatroid. The following properties hold:

Proposition 5.6. *For any polymatroid h :*

- $I_h(\mathbf{V}; \mathbf{W}|\mathbf{U}) \geq 0$ (this is a submodularity inequality).
- The chain rule holds:

$$I_h(\mathbf{UV}; \mathbf{W}|\mathbf{Z}) = I_h(\mathbf{U}; \mathbf{W}|\mathbf{Z}) + I_h(\mathbf{V}; \mathbf{W}|\mathbf{UZ}).$$
- An elemental mutual information term is an expression of the form $I_h(X_i; X_j|\mathbf{U})$, where $X_i \neq X_j \notin \mathbf{U}$ are single variables. Every mutual information $I(\mathbf{V}; \mathbf{W}|\mathbf{U})$ is the sum of elemental terms.
- For any subsets $\mathbf{U}_0 \subseteq \mathbf{U}$ and $\mathbf{V}_0 \subseteq \mathbf{V}$:

$$I_h(\mathbf{U}_0; \mathbf{V}_0|\mathbf{Z}) \leq I_h(\mathbf{U}; \mathbf{V}|\mathbf{Z}) \quad (31)$$

If $I_h(\mathbf{U}; \mathbf{V}|\mathbf{Z}) = 0$, then $I_h(\mathbf{U}_0; \mathbf{V}_0|\mathbf{Z}) = 0$.

Non-Shannon inequalities The first non-Shannon inequality was proven by Zhang and Yeung [ZY98]. We review it here, following the simplified presentation by Romashchenko [Rom22] (see also Csirmaz [Csi22]).

Zhang and Yeung [ZY98] proved the following:

Theorem 5.7. *The following is a non-Shannon inequality:*

$$\begin{aligned} I_h(X; Y) \leq & I_h(X; Y|A) + I_h(X; Y|B) + I_h(A; B) \\ & + I_h(X; Y|A) + I_h(A; Y|X) + I_h(A; X|Y) \end{aligned} \quad (32)$$

In other words, this inequality is valid for entropic vectors, but it cannot be proven using the basic Shannon inequalities, hence the term *non-Shannon*. The proof of this inequality, and that of several other non-Shannon inequalities proven after 1998, relies on the following *copy lemma*.

Lemma 5.8 (Copy Lemma). *Let \mathbf{X}, \mathbf{Y} be two disjoint sets of variables, and let \mathbf{h} be an entropic vector with variables \mathbf{XY} . Let \mathbf{Y}' be fresh copies of the variables \mathbf{Y} . Thus, each variable $Y \in \mathbf{Y}$ has a copy $Y' \in \mathbf{Y}'$. Then there exists an entropic vector \mathbf{h}' over variables $\mathbf{XY Y'}$ such that the following hold:*

$$\begin{aligned} I_{\mathbf{h}'}(\mathbf{Y}; \mathbf{Y}'|\mathbf{X}) &= 0 \\ \forall \mathbf{U} \subseteq \mathbf{XY} : & \mathbf{h}'(\mathbf{U}) = \mathbf{h}(\mathbf{U}) \\ \forall \mathbf{U}' \subseteq \mathbf{XY'} : & \mathbf{h}'(\mathbf{U}') = \mathbf{h}(\mathbf{U}) \end{aligned}$$

We say that \mathbf{Y}' is a copy of \mathbf{Y} over \mathbf{X} .

The first inequality asserts $\mathbf{Y} \perp \mathbf{Y}'|\mathbf{X}$. The second asserts that \mathbf{h}, \mathbf{h}' agree on \mathbf{XY} . And the last equality asserts that \mathbf{h}' on $\mathbf{XY'}$ is identical to \mathbf{h} on \mathbf{XY} up to the renaming of variables from $\mathbf{Y'}$ to \mathbf{Y} (assuming $X' = X$ for $X \in \mathbf{U}$).

Proof. (of Lemma 5.8) Let p be a probability distribution of random variables \mathbf{XY} that realizes the entropic vector \mathbf{h} . Define the following probability distribution p' of random variables $\mathbf{XY}Y'$: the domains of the variables \mathbf{Y}' is the same as that of \mathbf{Y} , and for all outcomes $(\mathbf{x}, \mathbf{y}, \mathbf{y}')$, $p'(\mathbf{X} = \mathbf{x}, \mathbf{Y} = \mathbf{y}, \mathbf{Y}' = \mathbf{y}') \stackrel{\text{def}}{=} \frac{p(\mathbf{X}=\mathbf{x}, \mathbf{Y}=\mathbf{y})p(\mathbf{X}=\mathbf{x}, \mathbf{Y}=\mathbf{y}')}{p(\mathbf{X}=\mathbf{x})}$. The claims in the lemma are easily verified. \square

In general, the copy lemma does not hold for polymatroids, as we will see shortly. We prove now Zhang and Yeung's inequality (32). Start from the following Shannon inequality over 5 variables, X, Y, A, B, A' :

$$\begin{aligned} I_h(X; Y) &\leq I_h(X; Y|A) + I_h(X; Y|B) + I_h(A; B) + \\ &+ I_h(X; Y|A') + I_h(A'; Y|X) + I_h(A'; X|Y) + \\ &+ 3I_h(A'; AB|XY) \end{aligned} \quad (33)$$

While this is “only” a Shannon inequality, it is surprisingly difficult to prove; we invite the readers to try it themselves, but, for completeness, we give the proof in Appendix .2. Consider now an entropic vector \mathbf{h} over four variables, X, Y, A, B . We apply the copy lemma, and copy AB over XY , resulting in an entropic vector h' with variables X, Y, A, B, A', B' . In particular, h' satisfies (33) (we don't use B'). Now we observe that (a) $I(A'; AB|XY) = 0$, and (b) every occurrence of A' in the second line can be replaced by A ; for example $I(X; Y|A') = I(X; Y|A)$, because $I(X; Y|A')$ is expressed in terms of $h(A'), h(XA'), h(YA'), h(XYA')$, which are equal to their copies $h(A), h(XA), h(YA), h(XYA)$. Thus, inequality (33) becomes (32), proving that (32) is valid for all entropic functions \mathbf{h} .

It remains to prove that (32) is not a Shannon inequality, and for that it suffices to describe one polymatroid that fails the inequality. To “see” this polymatroid, it is best to view it as being defined over a lattice. We take this opportunity to discuss another important concept: polymatroids on lattices.

Polymatroids on lattices A *polymatroid* on a lattice (L, \preceq) is a function $h : L \rightarrow \mathbb{R}_+$ satisfying:

$$\begin{aligned} h(\hat{0}) &= 0 \\ h(x \vee y) &\geq h(x) && \text{monotonicity} \\ h(x) + h(y) &\geq h(x \vee y) + h(x \wedge y) && \text{submodularity} \end{aligned}$$

Let \mathbf{X} be a set of n variables, and Σ a set of functional dependencies for \mathbf{X} . Recall from Sec. 2 that (L_Σ, \subseteq) is the lattice of closed sets. If a (standard) polymatroid $h \in \Gamma_n$ satisfies the functional dependencies Σ , then it is not hard to see that its restriction to L_Σ is a polymatroid on L_Σ . Conversely, any polymatroid \mathbf{h} on the lattice (L_Σ, \subseteq) can be extended to a standard polymatroid $\bar{h} : 2^{\mathbf{X}} \rightarrow \mathbb{R}_+$ by setting $\bar{h}(\mathbf{U}) \stackrel{\text{def}}{=} h(\mathbf{U}^+)$, and, furthermore, \mathbf{h} satisfies Σ . In short, there is a one-to-one correspondence between polymatroids satisfying a set of functional dependencies, and polymatroids defined on the associated lattice.

We now complete the proof of Theorem 5.7, by showing that inequality (32) does not hold for the polymatroid \mathbf{h} in Fig. 5. To read the figure, recall that $h(\mathbf{U}) = h(\mathbf{U}^+)$ for any set \mathbf{U} . For example, $ABX^+ = ABXY$, therefore $h(ABX) = h(ABX^+) = h(ABXY) = 4$, and, also, $h(AB) = h(AB^+) = h(ABXY) = 4$. We check now that \mathbf{h} violates the inequality (32):

$$\begin{aligned} I(X; Y) &= 1 \\ I(X; Y|A) + I(X; Y|B) + I(A; B) &= 0 + 0 + 0 \\ I(X; Y|A) + I(A; Y|X) + I(A; X|Y) &= 0 + 0 + 0 \end{aligned}$$

The LHS of (32) is 1, while the RHS is 0. This completes the proof of Theorem 5.7.




Figure 5: A lattice, and the polymatroid from [ZY98] defined on the lattice.

As a final comment, we note that it is instructive to check directly that the polymatroid in Fig. 5 fails to satisfy the copy lemma, without using Zhang and Yeung’s inequality; we provide a direct proof in Appendix .2.

Group-theoretic characterization of information inequalities Chan and Yeung [CY02] described an elegant characterization of information inequalities in terms of group inequalities. Given a finite group G and a subgroup $G_1 \subseteq G$, a *left coset* is a set of the form aG_1 , for some $a \in G$. By Lagrange’s theorem, the set of left cosets, denoted G/G_1 , forms a partition of G , and $|G/G_1| = |G|/|G_1|$. Fix n subgroups G_1, \dots, G_n , and consider the relational instance:

$$R = \{(aG_1, \dots, aG_n) \mid a \in G\} \quad (34)$$

whose set of attributes we identify, as usual, with $\mathbf{X} = \{X_1, \dots, X_n\}$. Notice that $|R| = |G|/|\bigcap_{i=1,n} G_i|$. The entropic vector \mathbf{h} associated to the relation R (Def. 4.1) is called a *group realizable entropic vector*, and the set of group realizable entropic vectors is denoted by $\Upsilon_n \subseteq \Gamma_n^*$, see Fig. 3. One can check that, for any subset of variables $\mathbf{U} \subseteq \mathbf{X}$, $h(\mathbf{U}) = \log |G|/|\bigcap_{X_i \in \mathbf{U}} G_i|$. The following was proven in [CY02]:

Theorem 5.9. *For any $\mathbf{h} \in \Gamma_n^*$ there exists a sequence $\mathbf{h}^{(r)} \in \Upsilon_n$, such that $\lim_{r \rightarrow \infty} \frac{1}{r} \mathbf{h}^{(r)} = \mathbf{h}$.*

It follows easily from the original proof that, if \mathbf{h} satisfies a set of functional dependencies, then so do all functions $\mathbf{h}^{(r)}$, for $r \geq 0$; for completeness, we will include the argument in Appendix .3.

Open Problems Characterizing the valid entropic information inequalities is a major open problem. Matús [Mat07] proved that, for $n \geq 4$, there are infinitely many independent non-Shannon inequalities. Currently, the only techniques known for proving such inequalities consists of repeated applications of Shannon inequalities and the Copy Lemma.

A related open problem is the complexity of deciding Shannon inequalities: what is the complexity of checking $\Gamma_n \models \mathbf{c} \cdot \mathbf{h} \geq 0$, as a function of $\|\mathbf{c}\|_1$? It is implicit in the proof of Theorem 5.2 that this can be decided in time exponential in n , but the complexity in terms of $\|\mathbf{c}\|_1$ is open. More discussion can be found in [KKNS20]

5.2 The Entropic Bound Is Asymptotically Tight

We prove here Theorem 5.5 item (1). The plan is the following. We need to find a database \mathbf{D} such that $\log |Q(\mathbf{D})|$ comes close to $\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b})$. By definition, there exists $\mathbf{h} \in \Gamma_n^*$

s.t. $h(\mathbf{X})$ is close to $\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b})$. We can't construct a database \mathbf{D} out of \mathbf{h} , because the probability distribution realizing \mathbf{h} may be non-uniform, instead we use Chan and Yeung's theorem to approximate $r\mathbf{h}$ by a group realizable vector $\mathbf{h}^{(r)}$, which is by definition associated to a relation instance. Hence, the need to amplify by the factor r . However, if we amplify, we don't know how $\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, r\mathbf{b})$ grows. Here we use Theorem 5.4, showing that $\text{Log-L-Bound}_{\Gamma_n^*}$ and $\text{Log-U-Bound}_{\Gamma_n^*}$ are asymptotically equal, then use the fact that $\text{Log-U-Bound}_{\Gamma_n^*}$ is linear, see Eq. (26). We give the details next.

By Corollary 4.8, for all $k \in \mathbb{N}$:

$$\frac{\sup_{\mathbf{D}: \mathbf{D} \models \mathbf{B}^k} \log |Q(\mathbf{D})|}{\text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})} \leq 1$$

Together with Theorem 5.4 (Eq. (27)) this implies:

$$\sup_k \frac{\sup_{\mathbf{D}: \mathbf{D} \models \mathbf{B}^k} \log |Q(\mathbf{D})|}{\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})} \leq 1$$

To prove equality, it suffices to show that, $\forall \varepsilon > 0, \exists k \in \mathbb{N}$ such that:

$$\log |Q(\mathbf{D})| \geq (1 - \varepsilon)^4 \text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b}) \quad (35)$$

Let $U \stackrel{\text{def}}{=} \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b})$. We will assume that $\text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b})$ is finite; otherwise, we let U be an arbitrarily large number and the proof below requires only minor adjustments, which we omit. We will assume w.l.o.g. that $U > 0$. Recall that $\text{Log-U-Bound}_{\Gamma_n^*}$ is linear (26). We prove:

Claim 1. *For all $\varepsilon > 0$, there exists $k \in \mathbb{N}$, and a database \mathbf{D} such that $\mathbf{D} \models (\Sigma, \mathbf{B}^k)$ and $\log |Q(\mathbf{D})| \geq (1 - \varepsilon)^4 kU$*

Eq. (35) follows from $kU = \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b}) \geq \text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})$. It remains to prove Claim 1. Since $\text{Log-L-Bound}_{\Gamma_n^*}$ and $\text{Log-U-Bound}_{\Gamma_n^*}$ are asymptotically equal (27), there exists $k_0 \in \mathbb{N}$ such that

$$\begin{aligned} \text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k_0\mathbf{b}) &\geq \\ &\geq (1 - \varepsilon) \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k_0\mathbf{b}) = (1 - \varepsilon)k_0U \end{aligned} \quad (36)$$

By the definition of $\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k_0\mathbf{b})$ in (24), there exists $\mathbf{h} \in \Gamma_n^*$ such that:

$$\begin{aligned} h(\mathbf{X}) &\geq (1 - \varepsilon) \text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k_0\mathbf{b}) \geq (1 - \varepsilon)^2 k_0U \\ h(\sigma) &\leq k_0 b_\sigma, \quad \forall \sigma \in \Sigma \end{aligned}$$

At this point we need the following Slack Lemma:

Lemma 5.10 (Slack Lemma). *For every $\mathbf{h} \in \Gamma_n^*$ and every $\varepsilon \in [0, 1]$, there exists $k \in \mathbb{N}$ and $\mathbf{h}' \in \Gamma_n^*$ such that:*

$$\begin{aligned} \mathbf{h}' &\geq (1 - \varepsilon)k\mathbf{h} \\ \forall \mathbf{U}, \mathbf{V} \subseteq \mathbf{X} : \quad h'(\mathbf{V}|\mathbf{U}) &\leq (1 - \varepsilon/2)kh(\mathbf{V}|\mathbf{U}) \end{aligned}$$

Proof. Assume w.l.o.g. that $\varepsilon > 0$, and set $k \stackrel{\text{def}}{=} \lceil \frac{1}{\varepsilon} \rceil$ and $\mathbf{h}' \stackrel{\text{def}}{=} (k-1)\mathbf{h}$. Then $\frac{1}{\varepsilon} \leq k \leq 1 + \frac{1}{\varepsilon} \leq \frac{2}{\varepsilon}$ which implies $\varepsilon/2 \leq \frac{1}{k} \leq \varepsilon$. We have:

$$\begin{aligned}\mathbf{h}' &= \left(1 - \frac{1}{k}\right) k\mathbf{h} \geq (1 - \varepsilon)k\mathbf{h} \\ \mathbf{h}'(\mathbf{V}|\mathbf{U}) &= \left(1 - \frac{1}{k}\right) k\mathbf{h}(\mathbf{V}|\mathbf{U}) \leq (1 - \varepsilon/2)k\mathbf{h}'(\mathbf{V}|\mathbf{U})\end{aligned}$$

□

We apply the Slack Lemma to \mathbf{h} and obtain a number k_1 and an entropic vector \mathbf{h}' such that:

$$\begin{aligned}h'(\mathbf{X}) &\geq (1 - \varepsilon)k_1 h(\mathbf{X}) \geq (1 - \varepsilon)^3 k_0 k_1 U \\ h'(\sigma) &\leq (1 - \varepsilon/2)k_1 h(\sigma) \leq (1 - \varepsilon/2)k_0 k_1 b_\sigma, \quad \forall \sigma \in \Sigma\end{aligned}\tag{37}$$

Let $g \stackrel{\text{def}}{=} \min_{\mathbf{U}, \mathbf{V} \subseteq \mathbf{X}: h'(\mathbf{V}|\mathbf{U}) > 0} h'(\mathbf{V}|\mathbf{U})$ be the smallest non-zero value of $h'(\mathbf{V}|\mathbf{U})$. By Chan and Yeung's theorem 5.9, there exists a group realizable entropic vector $\mathbf{h}^{(r)}$ that satisfies all the FDs satisfied by \mathbf{h}' , and $\|\mathbf{h}' - \frac{1}{r}\mathbf{h}^{(r)}\|_\infty \leq \varepsilon g/4$. Since $U > 0$ we have $h'(\mathbf{X}) > 0$ hence $h'(\mathbf{X}) \geq g$ and we derive from (37):

$$\begin{aligned}\frac{1}{r}h^{(r)}(\mathbf{X}) &\geq h'(\mathbf{X}) - \varepsilon g/4 \geq (1 - \varepsilon/4)h'(\mathbf{X}) \\ &\geq (1 - \varepsilon)^4 k_0 k_1 U\end{aligned}$$

On the other hand, $\frac{1}{r}h^{(r)}(\mathbf{V}|\mathbf{U}) \leq h(\mathbf{V}|\mathbf{U}) + \varepsilon g/2$, for all sets \mathbf{U}, \mathbf{V} . We use $\mathbf{h}' \models (1 - \varepsilon/2)k_0 k_1 \mathbf{b}$ to prove $\mathbf{h}^{(r)} \models rk_0 k_1 \mathbf{b}$. Consider a statistics $\sigma \in \Sigma$. If $h'(\sigma) = 0$, then \mathbf{h}' satisfies the FD σ , and therefore $\mathbf{h}^{(r)}$ also satisfies this FD, thus $h^{(r)}(\sigma) = 0 \leq k_0 k_1 b_\sigma$. If $h'(\sigma) > 0$ then $h'(\sigma) \geq g$ and the claim follows from:

$$\begin{aligned}\frac{1}{r}h^{(r)}(\sigma) &\leq h'(\sigma) + \varepsilon g/2 \\ &\leq h'(\sigma) + (\varepsilon/2)h'(\sigma) = (1 + \varepsilon/2)h'(\sigma) \\ &\leq (1 + \varepsilon/2)(1 - \varepsilon/2)k_0 k_1 b_\sigma \leq k_0 k_1 b_\sigma\end{aligned}$$

So far, we have:

$$h^{(r)}(\mathbf{X}) \geq (1 - \varepsilon)^4 rk_0 k_1 U \quad \mathbf{h}^{(r)} \models rk_0 k_1 \mathbf{b}\tag{38}$$

To complete the proof of Claim 1, we construct the database \mathbf{D} as follows. Let the relation R be the group realization of $\mathbf{h}^{(r)}$ (Eq. (34)). For each relation $R_j(\mathbf{Y}_j)$, define $R_j^{\mathbf{D}} \stackrel{\text{def}}{=} \Pi_{\mathbf{Y}_j}(R)$. By construction, $Q(\mathbf{D}) = R$, and $\log |Q(\mathbf{D})| = h^{(r)}(\mathbf{X}) \geq (1 - \varepsilon)^4 rk_0 k_1 U$ by (38). Furthermore, since $\mathbf{h}^{(r)}$ is group-realized, for every statistics $\sigma \in \Sigma$, with guard R_σ , we have $\log \deg_{R_\sigma}(\sigma) = h^{(r)}(\sigma) \leq rk_0 k_1 b_\sigma$; thus, $\mathbf{D} \models (\Sigma, \mathbf{B}^{rk_0 k_1})$. This implies:

$$\sup_{\mathbf{D}: \mathbf{D} \models (\Sigma, \mathbf{B}^{rk_0 k_1})} \log |Q(\mathbf{D})| \geq (1 - \varepsilon)^4 rk_0 k_1 U$$

proving Claim 1 for $k = rk_0 k_1$.

5.3 The Polymatroid Bound Is Not Asymptotically Tight

We prove now Theorem 5.5 item (2).

Proposition 5.11. *The following is a non-Shannon inequality:*

$$\begin{aligned}
11h(ABXYC) \leq & \tag{39} \\
& 3h(XY) + 3h(AX) + 3h(AY) \\
& + h(BX) + h(BY) + 5h(C) \\
& + (h(XYC|AB) + 4h(BC|AXY) + h(AC|BXY)) \\
& + (h(BXY|AC) + 2h(ABY|XC) + 2h(ABX|YC))
\end{aligned}$$

Proof. Consider the following five inequalities:

$$\begin{aligned}
0 &\leq 3h(AX) + 3h(AY) - 4h(AXY) - h(A) \\
&\quad + h(BX) + h(BY) - h(BXY) \\
&\quad - h(AB) + 3h(XY) - 2h(X) - 2h(Y) \\
0 &\leq h(A) + h(C) - h(AC) \\
0 &\leq 2(h(X) + h(C) - h(XC)) \\
0 &\leq 2(h(Y) + h(C) - h(YC)) \\
11h(ABXYC) &= 11h(ABXYC)
\end{aligned}$$

The first inequality holds because it is inequality Eq. (32), expanded and re-arranged. The next three inequalities are basic Shannon inequalities. The last line is an identity. A tedious but straightforward calculation shows that if we add the five (in)equalities above, then we obtain (39), proving the claim. \square

Consider the following query, derived from inequality (39):

$$\begin{aligned}
Q(A, B, X, Y, C) = & R_1(X, Y) \wedge R_2(A, X) \wedge R_3(A, Y) \\
& \wedge R_4(B, X) \wedge R_5(B, Y) \wedge R_6(C) \\
& \wedge R_7(A, B, X, Y, C)
\end{aligned}$$

and the following statistics:

$$\begin{aligned}
\Sigma = & \{(XY), (AX), (AY), (BX), (BY), (C), \\
& (XYC|AB), (BC|AXY), h(AC|BXY), \\
& (BXY|AC), (ABY|XC), (ABX|YC)\} \\
\mathbf{b} = & \{b_{XY} = b_{AX} = b_{AY} = b_{BX} = b_{BY} = 3, b_C = 2, \\
& b_{XYC|AB} = b_{BC|AXY} = b_{AC|BXY} = 0, \\
& b_{(BXY|AC)} = b_{ABY|XC} = b_{ABX|YC} = 0\}
\end{aligned}$$

In other words, we are given the cardinalities of R_1, \dots, R_6 , but are not given the cardinality of R_7 , instead we are told that it satisfies the 6 FD's corresponding to the 6 conditional terms in inequality (39). Consider any scale factor $k > 0$, and the scaled log-statistics $k\mathbf{b}$. Inequality (39) and the definition (22) imply:

$$\begin{aligned}
\text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b}) \leq \\
k \frac{3b_{XY} + 3b_{AX} + 3b_{AY} + b_{BX} + b_{BY} + 5b_C}{11} = \frac{43k}{11}
\end{aligned}$$




Figure 6: A polymatroid proving that the polymatroid bound is not tight.

By Corollary 4.8, for any database \mathbf{D} , if $\mathbf{D} \models (\Sigma, \mathbf{B}^k)$ then:

$$\log |Q(\mathbf{D})| \leq \frac{43k}{11}$$

On the other hand, consider the polymatroid $k\mathbf{h}$, where \mathbf{h} is the polymatroid in Fig. 6. Since $h(ABXYC) = 4$ and $\mathbf{h} \models (\Sigma, \mathbf{b})$, it follows that $kh(ABXYC) = 4k$, and $k\mathbf{h} \models k\mathbf{b}$, therefore:

$$\frac{\sup_{\mathbf{D}: \mathbf{D} \models \mathbf{B}^k} \log |Q(\mathbf{D})|}{\text{Log-L-Bound}_{\Gamma_n}(Q, \Sigma, k\mathbf{b})} \leq \frac{43}{44}$$

This implies Theorem 5.5 item (2).

6 Simple Inequalities

We have a dilemma: the entropic bound is asymptotically tight, but it is open whether it is computable, while the polymatroid bound is computable, but is provably not tight in general. We show in this section that, under a reasonable syntactic restriction on the statistics Σ , these two bounds are equal. We do this by describing a similar syntactic restriction for information inequalities, which we call *simple inequalities*. In that case validity over entropic functions coincides with validity over polymatroids, and we recover the stronger notion of tightness that we had for the AGM bound.

6.1 Background: Subclasses of Polymatroids

A polymatroid \mathbf{h} is called *modular* if the submodularity inequality (15) is an equality. Equivalently, \mathbf{h} is modular if $\mathbf{h} \geq 0$ and for every subset $\alpha \subseteq [n]$, $h(\mathbf{X}_\alpha) = \sum_{i \in \alpha} h(X_i)$. We will denote by M_n the set of modular polymatroids, see Fig. 3. For each $i = 1, n$, we call the function \mathbf{h}^{X_i} in Fig. 4 a *basic modular function*; recall that $h^{X_i}(\mathbf{U}) = 1$ when $X_i \in \mathbf{U}$ and $= 0$ otherwise. The following holds:

Proposition 6.1. (1) A function \mathbf{h} is modular iff it is a positive linear combination of basic modular functions, $\mathbf{h} = \sum_i a_i \mathbf{h}^{X_i}$, where $a_i \geq 0$ for all i . (2) Every modular function is entropic.

Proof. Item (1) is straightforward, but item (2) requires some thought. It suffices to prove that $a\mathbf{h}^{X_i}$ is entropic for all real numbers $a \geq 0$. For that purpose we need to describe one random variable

X_i , whose entropy is $h(X_i) = a$. Let N be a natural number such that $\log N \geq a$, and consider the uniform probability space where X_i has N outcomes with the same probabilities, $p_i = 1/N$, $i = 1, N$. Replace p_1 by $p_1 + \theta$, and replace each p_j with $j > 1$ by $p_j - \theta/(N-1)$, for $\theta \in [0, 1 - \frac{1}{N}]$. When $\theta = 0$ then the distribution is uniform and $h(X_j) = \log N$; when $\theta = 1 - \frac{1}{N}$ then the distribution is deterministic, $p_1 = 1, p_2 = \dots = p_N = 0$, and $h(X_j) = 0$. By continuity, there exists some θ where $h(X_j) = a$. \square

Fix a set of variables $\mathbf{W} \subseteq \mathbf{X}$. The *step function at \mathbf{W}* is:

$$h_{\mathbf{W}}(\mathbf{U}) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } \mathbf{U} \subseteq \mathbf{W} \\ 1 & \text{otherwise} \end{cases} \quad (40)$$

There are $2^n - 1$ non-zero step functions (since $h_{\mathbf{X}} = 0$). $h_{\mathbf{W}}$ is the entropy of the (uniform distribution of the) following relation with 2 tuples:

$$R_{\mathbf{W}} \stackrel{\text{def}}{=} R^{\mathbf{X}-\mathbf{W}} \stackrel{\text{def}}{=} \begin{array}{cc|c} \mathbf{W} & \mathbf{X}-\mathbf{W} & p \\ \hline 0 \dots 0 & 0 \dots 0 & 1/2 \\ 0 \dots 0 & 1 \dots 1 & 1/2 \end{array} \quad (41)$$

Sometimes it is convenient to use an alternative notation. For a set of variables $\mathbf{V} \subseteq \mathbf{X}$, define:

$$h^{\mathbf{V}}(\mathbf{U}) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } \mathbf{U} \cap \mathbf{V} = \emptyset \\ 1 & \text{otherwise} \end{cases} \quad (42)$$

Then $h^{\mathbf{V}} = h_{\mathbf{X}-\mathbf{V}}$. A basic modular function h^{X_i} is the same as the step function $h^{\{X_i\}}$; if $|\mathbf{V}| \geq 2$ then $h^{\mathbf{V}}$ is not modular.

Definition 6.2. A normal polymatroid is a positive linear combination of step functions,

$$\mathbf{h} = \sum_{\mathbf{V} \subseteq \mathbf{X}, \mathbf{V} \neq \emptyset} a_{\mathbf{V}} h^{\mathbf{V}} \quad (43)$$

where $a_{\mathbf{V}} \geq 0$ for all \mathbf{V} .

We denote by N_n the set of normal polymatroids, see Fig. 3. Normal polymatroids are the same as *polymatroids with a non-negative I-measure* described in [Yeu08, KKNS21].

Proposition 6.3. The $2^n - 1$ non-zero step functions $h^{\mathbf{V}}$, $\mathbf{V} \neq \emptyset$ form a basis of the vector space $\{\mathbf{h} \in \mathbb{R}^{2^{[n]}} \mid h(\emptyset) = 0\}$. More precisely, every such vector \mathbf{h} satisfies $\mathbf{h} = \sum_{\mathbf{V}} a_{\mathbf{V}} h^{\mathbf{V}}$, where:

$$a_{\mathbf{U}} \stackrel{\text{def}}{=} - \sum_{\mathbf{V} \subseteq \mathbf{U}} (-1)^{|\mathbf{V}|} h(\mathbf{V} \mid \mathbf{X} - \mathbf{U}) \quad (44)$$

The proof follows by solving the following system of linear equations with unknowns $a_{\mathbf{V}}$:

$$\forall \mathbf{U} \neq \emptyset: \quad h(\mathbf{U}) = \sum_{\mathbf{V}: \mathbf{V} \cap \mathbf{U} \neq \emptyset} a_{\mathbf{V}} \quad (45)$$

The solution is obtained by using Möbius' inversion formula (we prove this in Appendix .4) and consists of the expression (44). Expression (44) is called *conditional interaction information*, and denoted by $I(X_{i_1}; X_{i_2}; \dots \mid \mathbf{X} - \mathbf{U})$, where $\mathbf{U} = \{X_{i_1}, X_{i_2}, \dots\}$. The following holds (the proof is immediate and omitted):

Proposition 6.4. (1) A function \mathbf{h} is a normal polymatroid iff, for every set $U \subseteq \mathbf{X}$, $U \neq \emptyset$, the conditional interaction information (44) is ≥ 0 . (2) Every normal polymatroid is entropic.

Example 6.5. The parity function \mathbf{h} Fig. 2 is the simplest example of a polymatroid that is not normal. The coefficients can be derived using (44), or, we can check directly that:

$$\mathbf{h} = \mathbf{h}^{X,Y} + \mathbf{h}^{X,Z} + \mathbf{h}^{Y,Z} - \mathbf{h}^{X,Y,Z}$$

The coefficient of $\mathbf{h}^{X,Y,Z}$ is negative, hence \mathbf{h} is not normal.

6.2 Special Inequalities

We describe here a class of information inequalities, called *simple inequalities*, where Γ_n^* -validity coincides with Γ_n -validity. The modular and normal polymatroids turn out to be the key tools to study these inequalities.

The following is sometimes referred in the literature as the *modularization lemma*:

Lemma 6.6. For any polymatroid \mathbf{h} there exists a modular polymatroid \mathbf{h}' such that (a) $\mathbf{h}' \leq \mathbf{h}$ and (b) $\mathbf{h}'(\mathbf{X}) = \mathbf{h}(\mathbf{X})$.

Proof. Order the variables arbitrarily X_1, \dots, X_n and define $h'(X_i) \stackrel{\text{def}}{=} h(X_i | \mathbf{X}_{[1:i-1]})$, where $\mathbf{X}_{[1:i-1]} \stackrel{\text{def}}{=} \{X_1, \dots, X_{i-1}\}$. We check condition (a): for $\alpha \subseteq [n]$,

$$\begin{aligned} h'(\mathbf{X}_\alpha) &= \sum_{i \in \alpha} h(X_i | \mathbf{X}_{[1:i-1]}) \\ &\leq \sum_{i \in \alpha} h(X_i | \mathbf{X}_{[1:i-1] \cap \alpha}) = h(\mathbf{X}_\alpha) \end{aligned}$$

We check (b): $h'(\mathbf{X}) = \sum_{i=1,n} h(X_i | \mathbf{X}_{[1:i-1]}) = h(\mathbf{X})$. \square

The modularization lemma gives us an alternative, and more general proof of Theorem 4.10:

Corollary 6.7. Consider an inequality of the form $\sum_i w_i h(\mathbf{V}_i) \geq h(\mathbf{X})$, where $w_i \geq 0$ and \mathbf{V}_i are subsets of \mathbf{X} . The following conditions are equivalent:

- (1) The inequality is valid for polymatroids.
- (2) The inequality is valid for entropic functions.
- (3) The inequality is valid for modular functions.

Proof. The implications (1) \Rightarrow (2) \Rightarrow (3) are immediate. We prove (3) \Rightarrow (1), by contradiction: if the inequality fails on some polymatroid \mathbf{h} , $\sum_i w_i h(\mathbf{V}_i) < h(\mathbf{X})$, and \mathbf{h}' is the modular function in Lemma 6.6, then, $\sum_i w_i h'(\mathbf{V}_i) \leq \sum_i w_i h(\mathbf{V}_i) < h(\mathbf{X}) = h'(\mathbf{X})$ contradicting (3). \square

We prove in Appendix .5 the following extension of the Modularization Lemma:

Lemma 6.8. For any polymatroid \mathbf{h} there exists a normal polymatroid \mathbf{h}' such that (a) $\mathbf{h}' \leq \mathbf{h}$, (b) $h'(\mathbf{X}) = h(\mathbf{X})$, and (c) $h'(X_i) = h(X_i)$ for every variable $X_i \in \mathbf{X}$.

Definition 6.9. We call a set of statistics Σ simple if, for all $(\mathbf{V} | \mathbf{U}) \in \Sigma$, $|\mathbf{U}| \leq 1$. A simple information inequality is a Σ -inequality where Σ is simple:

$$\sum_{\sigma \in \Sigma} w_\sigma h(\sigma) \geq h(\mathbf{X}) \tag{46}$$

We immediately derive:

Corollary 6.10. *Given a simple inequality (46), the following are equivalent:*

- (1) *The inequality is valid for polymatroids.*
- (2) *The inequality is valid for entropic functions.*
- (3) *The inequality is valid for normal polymatroids*

The proof is identical to that of Corollary 6.7 and omitted.

6.3 Special Databases

When the statistics Σ are simple, then we show here that the polymatroid and the entropic bound coincide. We also show that the bound is tight, using a similar notion of tightness as in the AGM bound, where the ratio between the lower and upper bound depends only on the query; also, there is no need to amplify the statistics values. Moreover, like in the AGM bound, the worst-case database instance has a special structure, which we call a *normal database*. We start by showing:

Theorem 6.11. *If Σ is simple, then:*

$$\begin{aligned} \text{Log-U-Bound}_{N_n}(Q, \Sigma, \mathbf{b}) &= \\ &= \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b}) = \text{Log-U-Bound}_{\Gamma_n}(Q, \Sigma, \mathbf{b}) \end{aligned}$$

Proof. Since $N_n \subseteq \Gamma_n^* \subseteq \Gamma_n$ we have inequalities above: $\dots \leq \dots \leq \dots$ Corollary 6.10 implies $\text{Log-U-Bound}_{N_n}(Q, \Sigma, \mathbf{b}) = \text{Log-U-Bound}_{\Gamma_n}(Q, \Sigma, \mathbf{b})$, hence all three quantities are equal. \square

We describe now normal relational instances, and normal databases. Start with a single relation $R(\mathbf{X})$ with n attributes \mathbf{X} . Recall that an instance R is a *product relation* if $R = S_1 \times \dots \times S_n$, for n sets S_1, \dots, S_n : the worst-case instance of the AGM bound consisted of product relations. We generalize this concept:

Definition 6.12. *A relation instance T with n attributes is a normal relation if there exists m finite sets S_1, \dots, S_m and a function $\psi : [n] \rightarrow 2^{[m]}$ such that*

$$T = \{(\mathbf{s}_{\psi(1)}, \mathbf{s}_{\psi(2)}, \dots, \mathbf{s}_{\psi(n)}) \mid \mathbf{s} \in S_1 \times \dots \times S_m\}$$

In a normal relation the values of an attribute can be tuples themselves. Every product relation is a normal relation, but not vice versa. A database instance is *normal* if each of its relations is normal. A *basic normal relation of size N* is the following:

$$T_N^{\mathbf{V}} \stackrel{\text{def}}{=} \{(k \cdot \mathbf{1}_{X_1 \in \mathbf{V}}, \dots, k \cdot \mathbf{1}_{X_n \in \mathbf{V}}) \mid k = 0, N-1\} \quad (47)$$

Here $\mathbf{1}_{X_i \in \mathbf{V}}$ is an indicator variable that is 1 when $X_i \in \mathbf{V}$ and 0 otherwise; thus, if an attribute X_i is in \mathbf{V} then it takes the values $0, 1, \dots, N-1$ in the relation $T_N^{\mathbf{V}}$, otherwise it has constant values 0. The entropic vector of $T_N^{\mathbf{V}}$ is $(\log N)\mathbf{h}^{\mathbf{V}}$. In particular, the relation $R^{\mathbf{V}}$ in (41) is $T_2^{\mathbf{V}}$.

Example 6.13. *We give three examples of normal relations with $n = 3$ attributes:*

$$\begin{aligned} A &= \{(i, j, k) \mid i, j, k \in [0 : N-1]\} && \text{product relation} \\ B &= \{(i, i, i) \mid i \in [0 : N-1]\}, && \text{normal relation} \\ C &= \{(i, (i, j), j) \mid i, j \in [0 : N-1]\} && \text{normal relation} \end{aligned}$$

Their cardinalities are $|A| = N^3$, $|B| = N$, $|C| = N^2$. We also notice:

$$A = T_N^X \otimes T_N^Y \otimes T_N^Z \quad B = T_N^{XYZ} \quad C = T_N^{X,Y} \otimes T_N^{Y,Z}$$

We prove that the lower bound for simple statistics is tight.

Theorem 6.14. *Let Σ be a set of simple statistics for a query Q and let \mathbf{B} be statistics values. Then there exists a worst-case instance \mathbf{D} such that $|Q(\mathbf{D})| \geq \frac{1}{2^{2^n-1}} \text{U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B})$.*

Proof. We use the following, whose proof is immediate:

Proposition 6.15. *Let $R(\mathbf{X})$, $R'(\mathbf{X})$ be relations over the same attributes \mathbf{X} .*

- *If R, R' are normal relations, then $R \otimes R'$ is normal.*
- *$\deg_{R \otimes R'}(\sigma) = \deg_R(\sigma) \cdot \deg_{R'}(\sigma)$, for all $\sigma = (\mathbf{V}|\mathbf{U})$.*

Denote by $U \stackrel{\text{def}}{=} \text{U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B})$, $\mathbf{b} \stackrel{\text{def}}{=} \log \mathbf{B}$, then $\log U = \text{Log-U-Bound}_{N_n}(Q, \Sigma, \mathbf{b}) = \text{Log-L-Bound}_{N_n}(Q, \Sigma, \mathbf{b})$, by Theorem 6.11 and Theorem 5.2 respectively. Let $\mathbf{h}^* \in N_n$ be the optimal solution to the linear program defining $\text{Log-L-Bound}_{N_n}(Q, \Sigma, \mathbf{b})$ (see Theorem 5.2), then $\mathbf{h}^* \models (\Sigma, \mathbf{b})$ and $h^*(\mathbf{X}) = \log U$. Since \mathbf{h}^* is normal, it can be written as:

$$\mathbf{h}^* = \sum_{\mathbf{V} \neq \emptyset} a_{\mathbf{V}} \mathbf{h}^{\mathbf{V}}, \quad a_{\mathbf{V}} \geq 0$$

Then, $\log U = h^*(\mathbf{X}) = \sum_{\mathbf{V} \neq \emptyset} a_{\mathbf{V}}$, and $U = \prod_{\mathbf{V}} 2^{a_{\mathbf{V}}}$.

For each set \mathbf{V} , $\emptyset \neq \mathbf{V} \subseteq \mathbf{X}$ we define:

$$\begin{aligned} b_{\mathbf{V}} &\stackrel{\text{def}}{=} \lfloor 2^{a_{\mathbf{V}}} \rfloor, & P^{\mathbf{V}} &\stackrel{\text{def}}{=} T_{b_{\mathbf{V}}}^{\mathbf{V}} && \text{basic normal relation (47)} \\ R &\stackrel{\text{def}}{=} \bigotimes_{\mathbf{V}} P^{\mathbf{V}} && \text{normal relation} \end{aligned}$$

Define the worst-case instance as $\mathbf{D} = (R_1^D, \dots, R_m^D)$, where $R_j^D = \Pi_{\mathbf{Y}_j}(R)$. We first check that \mathbf{D} satisfies the constraints, and for that let $\sigma \in \Sigma$ have witness R_{σ} , then:

$$\begin{aligned} \log \deg_{R_{\sigma}}(\sigma) &= \log \deg_R(\sigma) = \sum_{\mathbf{V}} \log \deg_{P^{\mathbf{V}}}(\sigma) \\ &= \sum_{\mathbf{V}} (\log b_{\mathbf{V}}) h^{\mathbf{V}}(\sigma) \leq \sum_{\mathbf{V}} a_{\mathbf{V}} h^{\mathbf{V}}(\sigma) = h^*(\sigma) \leq b_{\sigma} \end{aligned}$$

Finally, we check the query's output size:

$$\log |Q(\mathbf{D})| = \log |R| = \log \prod_{\mathbf{V}} |P^{\mathbf{V}}| = \sum_{\mathbf{V}} (\log b_{\mathbf{V}}) h^{\mathbf{V}}(\mathbf{X})$$

Since $h^{\mathbf{V}}(\mathbf{X}) = 1$, this implies $|Q(\mathbf{D})| = \prod_{\mathbf{V}} b_{\mathbf{V}} = \prod_{\mathbf{V}} \lfloor 2^{a_{\mathbf{V}}} \rfloor \geq \frac{1}{2^{2^n-1}} U$, because $\lfloor 2^{a_{\mathbf{V}}} \rfloor \geq \frac{1}{2} 2^{a_{\mathbf{V}}}$. \square

The reader may want to check the analogy with the worst-case instance of the AGM bound: the optimal solution \mathbf{v}^* there became here \mathbf{h}^* , and the domain $V_i = \llbracket 2^{v_i^*} \rrbracket$ defined for the variable X_i became here the normal relation $P^{\mathbf{V}}$. As before, we constructed the worst-case instance \mathbf{D} without amplifying the statistics, and $|Q(\mathbf{D})|$ is within a constant, which depends only on the query, of $\text{U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{B})$.

Discussion The restriction to simple statistics occurs naturally in many applications. Databases are often designed with simple keys (consisting of a single attribute), and applications that use degrees often consider only simple degrees. The restriction to simple statistics is often acceptable.

It remains open where one can extend this definition to richer classes of statistics, or inequalities, while still preserving the property that validity for entropic vectors is the same as validity for polymatroids. The set of statistics in Example 4.9 is not “simple”, yet the entropic bound coincides with the polymatroid bound. This (and other examples) suggests that other non-trivial syntactic classes may exist where these two bounds agree.

7 Query Evaluation

The query evaluation problem is: given a conjunctive query Q , evaluate it on a (usually large) database \mathbf{D} . In this paper we consider only the *data complexity*, where the query is fixed, and the runtime is given as a function of the statistics of \mathbf{D} . Database systems compute queries using a sequence of *binary joins*, of the form $C(X, Y, Z) = A(X, Y) \wedge B(Y, Z)$, which are written as $C = A \bowtie B$. Assuming all relations are pre-sorted, the time complexity of the join is $\tilde{O}(|A| + |B| + |A \bowtie B|)$. A *semi-join*, denoted $C = A \ltimes B$, is a join followed by the projection on the attributes of the first relation, meaning $C(X, Y) = \exists Z(A(X, Y) \wedge B(Y, Z))$. A semijoin can be computed in time $\tilde{O}(|A|)$.

A *Worst Case Optimal Join* (WCOJ) is an algorithm that evaluates Q in time no larger than its theoretical upper bound. A sequence of binary joins is usually not a WCOJ, because intermediate results may be larger than the theoretical upper bound of the query. For example the upper bound for the triangle query in Example 3.2 is $N^{3/2}$, but if we evaluate it as $(R \bowtie S) \bowtie T$, the join $R \bowtie S$ can have size N^2 .

Any WCOJ algorithm represents an indirect proof of the query's upper bound, since the size of the output cannot exceed the time complexity of the algorithm. For example, if we are given an algorithm for the triangle query, together with a proof that its runtime is $\tilde{O}(N^{3/2})$, then we have a proof that the size of the output is also $\tilde{O}(N^{3/2})$. This means that proving an upper bound on the query's output is *inevitable* for designing a WCOJ. We show in this section that one can proceed in reverse: given a proof of the upper bound, convert it into a WCOJ. We call this paradigm *From Proofs to Algorithms*. Thus, the question to ask in designing a WCOJ algorithm is: how do we *prove* an upper bound on the query's output? And how do we convert it into an algorithm?

7.1 Generic Join

Consider the setting of the AGM bound: we are given only cardinality statistics on the base relations. In that case, a proof of the upper bound is a proof of $\sum_j w_j h(\mathbf{Y}_j) \geq h(\mathbf{X})$, since it implies $|Q| \leq \prod_j |R_j|^{w_j}$. We gave a proof of this inequality in Theorem 4.10; the proof consists of conditioning on the last variable X_n , then applying induction on the remaining variables. We convert that proof into an algorithm: iterate X_n over its domain, and compute recursively the residual query. This algorithm is called *Generic Join*, or GJ, and was introduced by Ngo, Ré, and Rudra [NRR13]. We describe it in detail next.

Fix a full conjunctive query with variables \mathbf{X} , which we write as $Q = \bowtie_{j=1,m} R_j$. As usual, \mathbf{Y}_j are the variables of R_j . Generic Join computes Q as follows:

- Let X_n be an arbitrary variable.
- Partition the set of indices j into J_0 and J_1 :
$$J_0 \stackrel{\text{def}}{=} \{j \mid X_n \notin \mathbf{Y}_j\}, \quad J_1 \stackrel{\text{def}}{=} \{j \mid X_n \in \mathbf{Y}_j\}.$$
- Compute the set $D = \bigcap_{j \in J_1} \Pi_{X_n}(R_j)$.
- For each value $x \in D$, do:
 - Compute $R_j[x] := \Pi_{\mathbf{Y}_j - X_n}(\sigma_{X_n=x}(R_j))$, for $j \in J_1$.
 - Denote $R_j[x] := R_j$ for $j \in J_0$.
 - Compute the residual query $\bowtie_{j=1,m} R_j[x]$.

We invite the reader to check how the algorithm can be “read off” the proof of Theorem 4.10. To compute the runtime of the algorithm, assume that the relations are given in listing representation, sorted lexicographically using the attribute order X_n, X_{n-1}, \dots, X_1 . Then, the runtime, T_n , is:

$$T_n(R_1, \dots, R_m) = T_{\text{intersection}} + \sum_x T_{n-1}(R_1[x], \dots, R_m[x])$$

By induction hypothesis:

$$T_{n-1}(R_1[x], \dots, R_m[x]) = \tilde{O} \left(\prod_j |R_j[x]|^{w_j} \right)$$

which leads to:

$$\begin{aligned} T_n &= T_{\text{intersection}} + \tilde{O} \left(\prod_{j \in J_0} |R_j|^{w_j} \sum_x \prod_{j \in J_1} |R_j[x]|^{w_j} \right) \\ &\leq T_{\text{intersection}} + \tilde{O} \left(\prod_{j \in J_0} |R_j|^{w_j} \prod_{j \in J_1} \left(\sum_x |R_j[x]| \right)^{w_j} \right) \\ &= T_{\text{intersection}} + \tilde{O} \left(\prod_j |R_j|^{w_j} \right) \end{aligned}$$

We used Hölder’s inequality in Fig. 1 (since $\sum_{j \in J_1} w_j \geq 1$, because X_n is covered), and the fact that $\sum_x |R_j[x]| = |R_j|$ for $j \in J_1$. The crux of the algorithm is the intersection: its runtime should not exceed $\prod_j |R_j|^{w_j}$, and for that it suffices to iterate over the smallest set $\Pi_{X_n}(R_j)$, and probe in the others: the runtime is $\tilde{O}(\min_{j \in J_1} |R_j|) \leq \tilde{O}(\prod_{j \in J_1} |R_j|^{w_j})$, since $\sum_{j \in J_1} w_j \geq 1$.

Example 7.1. Using the variable order X, Y, Z , GJ computes the triangle query $R(X, Y) \wedge S(Y, Z) \wedge T(Z, X)$ as follows:

```
For  $x \in \Pi_X(R) \cap \Pi_X(T)$  do:
  For  $y \in \Pi_Y(R[X = x]) \cap \Pi_Y(S)$  do:
    For  $z \in \Pi_Z(S[Y = y]) \cap \Pi_Z(T[X = x])$  do:
      output( $x, y, z$ )
```

The choice of algorithm for computing the intersection is critical for GJ. To see this, consider the simplest query, $Q(X) = R(X) \wedge S(X)$, that is an intersection. The AGM bound is $\min(|R|, |S|)$, corresponding to the edge covers $(1, 0)$ and $(0, 1)$, and GJ must compute the query in time $\tilde{O}(\min(|R|, |S|))$. By assumption, R, S are already sorted, but we cannot run a standard merge algorithm, since its runtime is $O(|R| + |S|)$; instead, we iterate over the smaller relation and do a binary search in the larger.

Because of its simplicity and ease of implementation, GJ is the poster child of WCOJ algorithms. One remarkable property of GJ is that its runtime is always bounded by the AGM bound, no matter what variable order we choose. Before GJ, Veldhuizen [Vel14] described an algorithm called Leapfrog Triejoin (LFTJ), which uses a similar logic as GJ, but also specifies in the details of the required trie data structure. Several implementations of GJ/LFTJ exists today [SOC16, AtCG⁺15, FBS⁺20, MKS21, WWS23].

7.2 The Heavy/Light Algorithm

Balister and Bollobás [BB12] provided the following alternative proof of an inequality of the form (23), which we write in an equivalent form using integer coefficients:

$$E \stackrel{\text{def}}{=} \sum_{j=1,m} k_j h(\mathbf{Y}_j) \geq k_0 h(\mathbf{X}) \quad (48)$$

where $k_i \in \mathbb{N}$, for $i = 0, m$. View the expression E as a bag of terms $h(\mathbf{Y}_j)$ where each term $h(\mathbf{Y}_j)$ occurs k_j times. A *compression step* consists of the following:

- Choose two terms $h(\mathbf{U}), h(\mathbf{V}) \in E$ such that $\mathbf{U} \not\subseteq \mathbf{V}$ and $\mathbf{V} \not\subseteq \mathbf{U}$.
- Replace $h(\mathbf{U}) + h(\mathbf{V})$ with $h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V})$.

Theorem 7.2. [BB12] *Any sequence of compression steps eventually leads to:*

$$E = \ell_0 h(\mathbf{Z}_0) + \ell_1 h(\mathbf{Z}_1) + \dots, \quad \text{where } \mathbf{Z}_0 \supset \mathbf{Z}_1 \supset \dots \quad (49)$$

Furthermore, if each variable X_i is covered at least $k_0 \geq 1$ times³ by the original expression E in (48), then $\mathbf{Z}_0 = \mathbf{X}$ and $\ell_0 \geq k_0$; in particular, the inequality (48) is valid.

Proof. Each compression step strictly increases the quantity $\sum_{h(\mathbf{Z}) \in E} |\mathbf{Z}|^2$. To see this, write $\mathbf{U} = \mathbf{A} \cup \mathbf{C}$, $\mathbf{V} = \mathbf{B} \cup \mathbf{C}$ where $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are disjoint sets, then $|\mathbf{U}|^2 + |\mathbf{V}|^2 = (|\mathbf{A}| + |\mathbf{C}|)^2 + (|\mathbf{B}| + |\mathbf{C}|)^2$, while $|\mathbf{U} \cup \mathbf{V}|^2 + |\mathbf{U} \cap \mathbf{V}|^2 = (|\mathbf{A}| + |\mathbf{B}| + |\mathbf{C}|)^2 + |\mathbf{C}|^2$, and the latter is strictly larger when $|\mathbf{A}| \cdot |\mathbf{B}| > 0$. This quantity cannot exceed $(\sum_j k_j) n^2$, therefore compression needs to terminate, and this happens when for any two sets in E one contains the other. Then, E must have the form (49). Finally, we observe that compression preserves the number of times each variable X_i is covered by E , because the number of sets in $\{\mathbf{U}, \mathbf{V}\}$ containing X_i is the same as the number of sets in $\{\mathbf{U} \cup \mathbf{V}, \mathbf{U} \cap \mathbf{V}\}$ containing X_i . Therefore, if each variable is covered by E at least k_0 times, then $\mathbf{Z}_0 = \mathbf{X}$ and $\ell_0 \geq k_0$. \square

Call a sequence of compression steps that converts an expression E in (48) to (49) a *BB-proof sequence*. To derive an algorithm, we need to impose an additional restriction. Call a BB-proof sequence *divergent* if, after each compression step $h(\mathbf{U}) + h(\mathbf{V}) \rightarrow h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V})$, we can split E into $E' + E''$, such that E' contains $h(\mathbf{U} \cup \mathbf{V})$ and covers every variable at least k'_0 times, E'' contains $h(\mathbf{U} \cap \mathbf{V})$ and covers every variable at least k''_0 times, and $k_0 = k'_0 + k''_0$.

We convert a divergent BB-proof sequence into an algorithm called the *Heavy/Light Algorithm*. Let \mathbf{B} be the statistics values, $\mathbf{b} \stackrel{\text{def}}{=} \log \mathbf{B}$, and set $B \stackrel{\text{def}}{=} \max_j B_j$. Assume w.l.o.g. that the inequality $E \stackrel{\text{def}}{=} \sum_j k_j h(\mathbf{Y}_j) \geq k_0 h(\mathbf{X})$ is optimal, meaning that $\sum_j (k_j/k_0) b_j = \text{Log-U-Bound}_{\Gamma_n}(Q, \Sigma, \log(\mathbf{B}))$ (see Eq. (22)). Denote by \mathbf{h}^* an optimal solution to the dual, meaning $\mathbf{h}^* \models (\Sigma, \mathbf{b})$ and $\mathbf{h}^*(\mathbf{X}) = \text{Log-L-Bound}_{\Gamma_n}(Q, \Sigma, \mathbf{b})$ (see Eq. (24)). These two quantities are the same by Thm. 5.2: $\sum_j (k_j/k_0) b_j = \mathbf{h}^*(\mathbf{X})$. The algorithm uses a working memory which stores, for each term $h(\mathbf{Z})$ in E , a temporary relation $S(\mathbf{Z})$, called the *guard* of $h(\mathbf{Z})$, and maintains the invariant: $\log |S(\mathbf{Z})| \leq \mathbf{h}^*(\mathbf{Z})$. Initially, the working memory is $\{R_j \mid k_j > 0\}$: by complementary slackness, if $k_j > 0$ then the dual constraint is tight, $\mathbf{h}^*(\mathbf{Y}_j) = b_j$, and the invariant holds because $\log |R_j| \leq \log B_j = b_j = \mathbf{h}^*(\mathbf{Y}_j)$.

The algorithm repeatedly processes a compression step $h(\mathbf{U}) + h(\mathbf{V}) \rightarrow h(\mathbf{U} \cup \mathbf{V}) + h(\mathbf{U} \cap \mathbf{V})$ of the BB-sequence, as follows. If the two guards are $S(\mathbf{U})$ and $S'(\mathbf{V})$, let $\mathbf{C} \stackrel{\text{def}}{=} \mathbf{U} \cap \mathbf{V}$, write

³Meaning: $\sum_{j: X_i \in \mathbf{Y}_j} k_j \geq k_0$.




Figure 7: A simple polymatroid used in Example 7.3.

$U = \mathbf{AC}$, $V = \mathbf{BC}$, and define $M \stackrel{\text{def}}{=} 2^{h^*(\mathbf{A}|\mathbf{C})}$. Partition the guard $S(\mathbf{AC})$ into two subsets:

$$\begin{aligned} S_{\text{light}}(\mathbf{A}, \mathbf{C}) &= \{(\mathbf{a}, \mathbf{c}) \in S \mid \deg_S(\mathbf{A}|\mathbf{C} = \mathbf{c}) \leq M\} \\ S_{\text{heavy}}(\mathbf{C}) &= \{\mathbf{c} \in \Pi_{\mathbf{C}}(S) \mid \deg_S(\mathbf{A}|\mathbf{C} = \mathbf{c}) > M\} \end{aligned}$$

Compute new guards using a join and a semijoin:

$$\begin{aligned} S''(\mathbf{A}, \mathbf{B}, \mathbf{C}) &:= S_{\text{light}}(\mathbf{A}, \mathbf{C}) \bowtie S'(\mathbf{B}, \mathbf{C}) \\ S'''(\mathbf{C}) &:= S_{\text{heavy}}(\mathbf{C}) \ltimes S'(\mathbf{B}, \mathbf{C}) \end{aligned}$$

The invariant holds because $|S''| \leq M \cdot |S|$ implies $\log |S''| \leq h^*(\mathbf{A}|\mathbf{C}) + h^*(\mathbf{BC}) \leq h^*(\mathbf{ABC})$, and because $|S_{\text{heavy}}| \leq |S|/M$ (since every $\mathbf{c} \in S_{\text{heavy}}$ occurs $\geq M$ times in S) implies $\log |S'''| \leq \log |S| - \log M = h^*(\mathbf{AC}) - h^*(\mathbf{A}|\mathbf{C}) = h^*(\mathbf{C})$. The runtime of the join and semijoin is $\leq \tilde{O}(2^{h^*(\mathbf{X})}) = \tilde{O}(\text{U-Bound}_{\Gamma_n}(Q, \Sigma, \mathbf{B}))$. Next, the algorithm proceeds recursively, by processing independently E' and E'' , semi-joins the result of E' with the relations missing from E' , similarly semi-joins the result of E'' with the relations missing from E'' , then returns the union of these two results. Correctness is easily checked.

Example 7.3. Consider the triangle query $R(X, Y) \wedge S(Y, Z) \wedge T(Z, X)$, and the following divergent proof: $h(XY) + h(YZ) + h(ZX) \rightarrow h(XYZ) + h(Y) + h(ZX) \rightarrow h(XYZ) + h(XYZ)$. Assume for simplicity that the three relations have the same cardinalities $|R| = |S| = |T| = B$. The optimal polymatroid is $\mathbf{h}^* \stackrel{\text{def}}{=} \log B \cdot \mathbf{h}$, for \mathbf{h} in Fig. 7.

The Heavy/Light Algorithm proceeds as follows. For the first compression step it partitioning R into:

$$\begin{aligned} R_{\text{light}} &:= \{(x, y) \mid \deg_R(X|Y = y) \leq 2^{h^*(Y|X)} = B^{1/2}\} \\ R_{\text{heavy}} &:= \{y \mid \deg_R(X|Y = y) > 2^{h^*(Y|X)} = B^{1/2}\} \end{aligned}$$

then it computes:

$$\begin{aligned} \text{Temp}_1(X, Y, Z) &:= R_{\text{light}}(X, Y) \bowtie S(Y, Z) \\ \text{Temp}_2(Y) &:= R_{\text{heavy}}(Y) \ltimes S(Y, Z) \end{aligned}$$

At this point the BB-proof diverged into two branches, $h(XYZ) \geq h(XYZ)$, and $h(Y) + h(ZX) \geq h(XYZ)$, and we perform a recursive call for each branch. The first branch immediately returns

Temp_1 , which we semi-join with T : $Q'(X, Y, Z) = \text{Temp}_1(X, Y, Z) \ltimes T(Z, X)$. The other branch applies the second compression step which corresponds to the following join operation:

$$\text{Temp}_3(X, Y, Z) := \text{Temp}_2(Y) \times T(Z, X)$$

which we semi-join with R and S : $Q''(X, Y, Z) := \text{Temp}_3(X, Y, Z) \ltimes R(X, Y) \ltimes S(Y, Z)$. Finally, we return the union $Q' \cup Q''$. The reader may verify that the runtime is $\tilde{O}(B^{3/2})$.

An advantage of the Heavy/Light Algorithm over GJ is that it reuses existing join operators, which already have very efficient implementations in database systems. However, the algorithm only works for divergent BB-proofs. This raises the question: does every inequality (48) have a divergent proof? The answer is negative, as provided by the following example due to Yilei Wang [Wan22].

Example 7.4. *The following has no divergent BB-proof:*

$$E = h(XYZ) + h(ZUV) + h(VWX) + h(YUW) \geq 2h(XYZUVW)$$

Assume w.l.o.g. that we start by compressing $h(XYZ) + h(ZUV) \rightarrow h(XYZUV) + h(Z)$ (by symmetry, all other choices are equivalent). Then we need to partition E into $E' + E''$. Suppose E' contains $h(Z)$; since E' covers every variable, it must contain both remaining terms $h(VWX) + h(YUW)$, which means that E'' can only contain $h(XYZUV)$ alone, and it does not cover W .

7.3 PANDA

Both Generic Join and the Heavy/Light Algorithm are restricted to cardinality statistics, in other words they only work in the framework of the AGM bound. PANDA, introduced in [KNS17], is a WCOJ algorithm that works for general statistics. While it runs in time given by the theoretical query upper bound, it also includes a polylogarithm factor in the size of the database, with a rather large exponent. We describe PANDA here at a high level, and refer the reader to [KNS17] for details.

Let Σ be a set of statistics, and consider a Σ -inequality with integer coefficients:

$$E \stackrel{\text{def}}{=} \sum_{\sigma \in \Sigma} k_{\sigma} h(\sigma) \geq k_0 h(\mathbf{X}) \quad (50)$$

A *CD-proof sequence* for the inequality (50) is a sequence of steps that convert the LHS to the RHS, where each step is one of the following:

- Composition: $h(\mathbf{U}) + h(\mathbf{V}|\mathbf{U}) \rightarrow h(\mathbf{UV})$.
- Decomposition: $h(\mathbf{UV}) \rightarrow h(\mathbf{U}) + h(\mathbf{V}|\mathbf{U})$.
- Submodularity: $h(\mathbf{V}|\mathbf{U}) \rightarrow h(\mathbf{V}|\mathbf{UW})$.
- No-Op: $h(\mathbf{U}) \rightarrow 0$.

We say that the CD-proof sequence *proves* the inequality (50) if it starts from the LHS and ends with the RHS. The following was proven in [KNS17]:

Lemma 7.5. *Inequality (50) is valid for polymatroids iff it admits a CD-proof sequence.*

PANDA converts a CD-proof sequence into an algorithm, similarly to the way we converted a BB-sequence to an algorithm. Given statistics Σ , \mathbf{B} , guarded by the query Q (Sec. 4.2), assume that inequality (50) is the optimal solution to $\text{Log-U-Bound}_{\Gamma_n}(Q, \Sigma, \log(\mathbf{B}))$; otherwise, choose a better inequality. Denote by h^* an optimal solution to $\text{Log-L-Bound}_{\Gamma_n}(Q, \Sigma, \log(\mathbf{B}))$. The algorithm has a working memory consisting of a guard, call it $S_{\mathbf{V}|\mathbf{U}}(\mathbf{Z})$, for every term $h(\mathbf{V}|\mathbf{U})$ in E , satisfying the following invariant: $\mathbf{V} \subseteq \mathbf{Z}$ and there exists a subset $\mathbf{U}_0 \subseteq \mathbf{U} \cap \mathbf{Z}$, such that:

$$\log \deg_{S_{\mathbf{V}|\mathbf{U}}}(\mathbf{V}|\mathbf{U}_0) \leq h^*(\mathbf{V}|\mathbf{U})$$

The guard need not have all variables \mathbf{U} , but only a subset \mathbf{U}_0 that is sufficient to prove the bound on the max-degree.

Initially, the working memory consists of all guards R_σ of the statistics $\sigma = (\mathbf{V}|\mathbf{U}) \in \Sigma$, where $k_\sigma > 0$. By complementary slackness, if $k_\sigma > 0$, then the corresponding constraint on h^* is tight, $h^*(\sigma) = b_\sigma \stackrel{\text{def}}{=} \log B_\sigma$, therefore $\log \deg_{R_\sigma}(\sigma) \leq b_\sigma = h^*(\sigma)$ because the input database satisfies the statistics. PANDA performs the following action for each step of the CD-proof sequence:

Composition $h(\mathbf{U}) + h(\mathbf{V}|\mathbf{U}) \rightarrow h(\mathbf{UV})$. Compute the new guard as:

$$S_{\mathbf{UV}} := \Pi_{\mathbf{U}}(S_{\mathbf{U}}) \bowtie \Pi_{\mathbf{U}_0\mathbf{V}}(S_{\mathbf{V}|\mathbf{U}})$$

Since $|S_{\mathbf{UV}}| \leq |S_{\mathbf{U}}| \cdot \deg_{S_{\mathbf{V}|\mathbf{U}}}(\mathbf{V}|\mathbf{U}_0)$, we have:

$$\begin{aligned} \log |S_{\mathbf{UV}}| &\leq \log \deg_{S_{\mathbf{U}|\emptyset}}(\mathbf{U}|\emptyset) + \log \deg_{S_{\mathbf{V}|\mathbf{U}}}(\mathbf{V}|\mathbf{U}_0) \\ &\leq h^*(\mathbf{U}) + h^*(\mathbf{V}|\mathbf{U}) = h^*(\mathbf{UV}) \end{aligned}$$

Thus, the invariant holds, and the runtime does not exceed the polymatroid bound, whose log is $h^*(\mathbf{X})$.

Submodularity $h(\mathbf{V}|\mathbf{U}) \rightarrow h(\mathbf{V}|\mathbf{UW})$. Here PANDA only records that the new term $h(\mathbf{V}|\mathbf{UW})$ has the same guard as the old term $h(\mathbf{V}|\mathbf{U})$.

Decomposition $h(\mathbf{UV}) \rightarrow h(\mathbf{U}) + h(\mathbf{V}|\mathbf{U})$. Here PANDA first projects out the extra variables in the guard of $h(\mathbf{UV})$ and obtains a relation $S := \Pi_{\mathbf{UV}}(S_{\mathbf{UV}})$ whose size $N \stackrel{\text{def}}{=} |S|$ satisfies $\log N \leq h^*(\mathbf{UV})$. Next, it performs *regularization*: partition S into $\log N$ fragments $S = \bigcup_{i=1, \log N} S_i$, where:

$$S_i(\mathbf{U}, \mathbf{V}) \stackrel{\text{def}}{=} \{(\mathbf{u}, \mathbf{v}) \in S \mid \deg_S(\mathbf{V}|\mathbf{U} = \mathbf{u}) \in [2^{i-1}, 2^i]\}$$

PANDA then continues with $\log B$ recursive calls. The i 'th recursive call replaces S with S_i in the query, adds two new statistics $(\mathbf{V}|\mathbf{U})$ and (\mathbf{U}) to Σ , and two log-statistics values, $b_{\mathbf{V}|\mathbf{U}} \stackrel{\text{def}}{=} i$ and $b_{\mathbf{U}} \stackrel{\text{def}}{=} N/2^{i-1}$, both with guard $S_i(\mathbf{UV})$. Then, PANDA computes a *new* optimal primal/dual solutions to the polymatroid bound, resulting in a new inequality (50) and a new polymatroid h^* . It uses these to compute the residual query where S is replaced by S_i . Finally, it returns the union of all $\log N$ results from all recursive calls.

We leave out several details of PANDA, including the proof of termination, and refer the reader to [KNS17]. We also note that PANDA was extended from computing full conjunctive queries, to computing Boolean conjunctive queries, with a runtime given by the *submodular width* of the query, a notion introduced by Marx [Mar13].

8 The Domination Problem

We now move beyond the query upper bound problem, and consider a related question, called the *domination problem*: given two queries Q, Q' , check if, for any database \mathbf{D} , $|Q(\mathbf{D})| \leq |Q'(\mathbf{D})|$. The queries Q and Q' need not have the same number of variables. In this section we consider full conjunctive queries that may have self-joins, i.e. the same relation name may occur several times in the query; for example in $R(X, Y) \wedge R(Y, Z)$ the same relation R occurs twice.

Definition 8.1. *Given two conjunctive queries $Q(\mathbf{X})$, $Q'(\mathbf{Y})$ we say that Q' dominates Q , and write $Q \preceq Q'$, if for every database instance \mathbf{D} , $|Q(\mathbf{D})| \leq |Q'(\mathbf{D})|$.*

The original motivation for the domination problems comes from the query containment problem under bag semantics. Given a (not necessarily full) conjunctive query $Q(\mathbf{Y}_0)$, as in (1), its value under *bag semantics* is a bag of tuples, where each tuple \mathbf{y}_0 occurs as many times as the number of homomorphisms from Q to \mathbf{D} that map \mathbf{Y}_0 to \mathbf{y}_0 . SQL uses bag semantics. Chaudhuri and Vardi [CV93] were the first to study the *query containment problem under bag semantics*: given Q, Q' , check whether $Q(\mathbf{D}) \subseteq Q'(\mathbf{D})$ for every \mathbf{D} , where both $Q(\mathbf{D})$, $Q'(\mathbf{D})$ are bags of tuples. This problem has been intensively studied in the last thirty years. It has been shown that the containment problem under bag semantics is undecidable for *unions of conjunctive queries* [IR95], and for *conjunctive queries with inequalities* [JKV06]; both used reduction from Hilbert's 10th Problem. It should be noted that, under set semantics, the containment problem for these two classes of queries is decidable.

When $Q() = \dots$ is a Boolean query, then under standard set-semantics it returns either $\{\}$ or $\{()\}$, representing FALSE and TRUE. Under bag semantics it may return a bag $\{(), (), \dots, ()\}$, representing a number, and this number is equal to the size of the output of the full query, $Q(\mathbf{X}) = \dots$. Based on this discussion, the domination problem $Q \preceq Q'$ for full conjunctive queries is the same as the query containment problem under bag semantics for Boolean queries.

Kopparty and Rossman [KR11] were the first to establish the connection between the domination problem and information theory. We describe this connection, following their example.

Example 8.2. *This example from [KR11] is attributed to Eric Vee. Consider the following queries:*

$$\begin{aligned} Q(X, Y, Z) &= R(X, Y) \wedge R(Y, Z) \wedge R(Z, X) \\ Q'(U, V, W) &= R(U, V) \wedge R(U, W) \end{aligned}$$

We will show that $Q \preceq Q'$. Chaudhuri and Vardi [CV93] already noted that, if there exists a surjective homomorphism $Q' \rightarrow Q$, then $Q \preceq Q'$. In our example we have three homomorphisms $\varphi_1, \varphi_2, \varphi_3 : Q \rightarrow Q'$, but none of them is surjective.

Consider the following linear expression in entropic terms, defined over the variables U, V, W in Q' :

$$E \stackrel{\text{def}}{=} h(UV) + h(UW) - h(U) = h(UV) + h(W|U)$$

(The expression is derived from the tree decomposition of Q' , as we explain below.) For each of the three homomorphism φ_i , denote by $E \circ \varphi_i$ the result of substituting the variables U, V, W in E with $\varphi_i(U), \varphi_i(V), \varphi_i(W)$.

Claim 2. *The following inequality holds for all polymatroids:*

$$h(XYZ) \leq \max(E \circ \varphi_1, E \circ \varphi_2, E \circ \varphi_3) \tag{51}$$

Proof. We expand:

$$\begin{aligned}
& \max(E \circ \varphi_1, E \circ \varphi_2, E \circ \varphi_3) = \\
& = \max(h(XY) + h(Y|X), h(YZ) + h(Z|Y), h(XZ) + h(X|Z)) \\
& \geq \frac{1}{3} (h(XY) + h(YZ) + h(ZX) + h(Y|X) + h(Z|Y) + h(X|Z)) \\
& = \frac{1}{3} ((h(XY) + h(Z|Y)) + (h(YZ) + h(X|Z)) + (h(ZX) + h(Y|X))) \\
& \geq h(XYZ)
\end{aligned}$$

where the last inequality follows from $h(XY) + h(Z|Y) \geq h(XY) + h(Z|XY) = h(XYZ)$ and similarly for the other two terms. \square

To prove $Q \preceq Q'$, consider a database instance \mathbf{D} , let $N \stackrel{\text{def}}{=} |Q(\mathbf{D})|$, and consider the uniform probability distribution $(Q(\mathbf{D}), p)$. Its entropy \mathbf{h} satisfies inequality (51): assume w.l.o.g. that $h(XYZ) \leq E \circ \varphi_1 = h(XY) + h(Y|X)$ (the other two cases are similar). We use φ_1 to define a probability space $(Q'(\mathbf{D}), p')$: for every three constants u, v, w in the instance \mathbf{D} s.t. $p(X = u) \neq 0$, define

$$p'(U = u, V = v, W = w) \stackrel{\text{def}}{=} \frac{p(X = u, Y = v)p(X = u, Y = w)}{p(X = u)}$$

Thus, $V \perp W|U$, the distribution of UV is the same as that of XY , and the distribution of UW is also the same as that of XY . (This is similar to the Copy Lemma 5.8.) Denoting by \mathbf{h}' the entropic vector associated to p' , we derive:

$$\begin{aligned}
\log |Q'(\mathbf{D})| & \geq h'(UVW) = h'(VW|U) + h(U) \\
& = h'(V|U) + h'(W|U) + h'(U) \quad \text{because } V \perp W|U \\
& = h(Y|X) + h(Y|X) + h(X) = h(XY) + h(Y|X) \\
& = E \circ \varphi_1 \geq h(XYZ) = \log |Q(\mathbf{D})|
\end{aligned}$$

We generalize Example 8.2. A *tree decomposition* of a query $Q(\mathbf{X}) = \bigwedge_j R_j(\mathbf{Y}_j)$ is a pair (T, χ) , where T is a tree and $\chi : \text{Nodes}(T) \rightarrow 2^{\mathbf{X}}$ such that every atom $R_j(\mathbf{Y}_j)$ is *covered*, meaning $\exists n, \mathbf{Y}_j \subseteq \chi(n)$, and for any variable $X \in \mathbf{X}$, the set of nodes $\{n \in \text{Nodes}(T) \mid X \in \chi(n)\}$ induces a connected subgraph of T . Each set $\chi(n)$ is called a *bag*. Q is *chordal* if it admits a tree decomposition where every bag $\chi(n)$ induces a clique in the Gaifman graph of Q ; equivalently, for any two variables $X, Y \in \chi(n)$ the query has a predicate that contains both X, Y . A chordal query has a canonical tree decomposition where the bags are the maximal cliques. Q is called *acyclic*⁴ if there exists a tree decomposition where each bag is precisely one atom of the query, $\chi(n) = \mathbf{Y}_j$ for some j . An acyclic query is, in particular, chordal.

Fix a query $Q(\mathbf{U})$ with variables \mathbf{U} and a tree decomposition T . We define the following expression of entropic terms:

$$E_T \stackrel{\text{def}}{=} \sum_{n \in \text{Nodes}(T)} h(\chi(n)) - \sum_{(n, n') \in \text{Edges}(T)} h(\chi(n) \cap \chi(n')) \quad (52)$$

⁴More precisely, it is called α -acyclic [Fag83].

Equivalently, choose a root node for T and orient all edges to point away from the root. Then:

$$E_T = \sum_{n \in \text{Nodes}(T)} h(\chi(n) | \chi(n) \cap \chi(\text{Parent}(n)))$$

where we set $\chi(\text{Parent}(\text{Root})) \stackrel{\text{def}}{=} \emptyset$. The following holds:

Theorem 8.3. *Let $Q(\mathbf{X}), Q'(\mathbf{U})$ be two full conjunctive queries, over variables \mathbf{X} and \mathbf{U} respectively.*

- *[KR11, KKNS21] Let T be a tree decomposition for Q' . If the following inequality holds for all entropic vectors \mathbf{h} :*

$$h(\mathbf{X}) \leq \max_{\varphi: \varphi \in \text{hom}(Q', Q)} E_T \circ \varphi \quad (53)$$

then Q' dominates Q , $Q \preceq Q'$.

- *[KKNS21] If Q' is chordal and $Q \preceq Q'$, then inequality (53) holds for all entropic vectors, where T is the canonical tree decomposition of Q' consisting of its maximal cliques. In other words, (53) is a necessary and sufficient condition for dominance.*

Let's call a tree decomposition T *simple* if for every edge $(n, n') \in \text{Edges}(T)$, $|\chi(n) \cap \chi(n')| \leq 1$. If Q' admits a simple tree decomposition, then condition (53) is decidable; the proof follows immediately from Lemma 6.8. This implies:

Corollary 8.4. *Assume that Q' is chordal and admits a simple tree decomposition. Then it is decidable whether $Q \preceq Q'$. Moreover, if $Q \not\preceq Q'$, then there exists a normal database instance (Sec. 6.3) such that $|Q(\mathbf{D})| > |Q'(\mathbf{D})|$.*

Finally, we remark that the connection between the query domination problem and information inequalities is very tight. The following was proven in [KKNS21]:

Theorem 8.5. *The following problems are computationally equivalent. (1) Check if an inequality of the form $\max_{j=1,p} (\mathbf{c}^{(j)} \cdot \mathbf{h}) \geq 0$ is valid for all entropic vectors \mathbf{h} , where $\mathbf{c}^{(1)}, \dots, \mathbf{c}^{(p)} \in \mathbb{R}^{2^{[n]}}$ are p vectors. (2) Given two queries Q, Q' where Q' is acyclic, check whether $Q \preceq Q'$.*

It is currently open whether these problems are decidable.

9 Conditional Inequalities and Approximate Implication

Our last application of information inequalities is for the approximate implication problem, which can be described informally as follows. Let $\sigma_1, \dots, \sigma_p$ be some constraints on the database (we will define shortly what constraints we consider), and suppose we have a proof of the implication $\bigwedge_i \sigma_i \Rightarrow \sigma_0$. The question is, if the database \mathbf{D} satisfies the constraints σ_i only approximatively, is it the case that that σ_0 also holds approximatively? We will show here that this question is related to *conditional information inequalities*, whose study requires us to do another deep dive into the space of polymatroids and entropic functions. We start by defining a conditional inequality:

Definition 9.1. *A conditional information inequality is an assertion of the following form:*

$$\mathbf{c}_1 \cdot \mathbf{h} \geq 0 \wedge \dots \wedge \mathbf{c}_p \cdot \mathbf{h} \geq 0 \Rightarrow \mathbf{c}_0 \cdot \mathbf{h} \geq 0 \quad (54)$$

where $\mathbf{c}_i \in \mathbb{R}^{2^{[n]}}$ for $i = 0, p$ are vectors.

Sometimes it will be more convenient to replace \mathbf{c}_i by $-\mathbf{c}_i$, and write the implication as $\bigwedge_i \mathbf{c}_i \cdot \mathbf{h} \leq 0 \Rightarrow \mathbf{c}_0 \cdot \mathbf{h} \leq 0$. As before, the validity of a conditional inequality depends on the domain of \mathbf{h} , e.g. it can be valid for polymatroids, or entropic functions, etc.

The first non-Shannon inequality discovered was a conditional inequality [ZY97], predating the first unconditioned Shannon inequality [ZY98]. Kaced and Romashchenko [KR13] showed the first examples of essentially conditional inequalities (explained below). We start by describing the connection between conditional inequalities and the constraint implication problem in databases, then study the *relaxation problem*, a technique for transferring exact inferences to approximate judgments. We end with the proof of Theorem 5.4, which we have postponed until we developed sufficient technical machinery.

9.1 The Constraint Implication Problem

An *integrity constraint*, σ , is an assertion about a relation $R(\mathbf{X})$ that is required to hold strictly. The constraints considered here are Functional Dependencies (FD), already reviewed in Sec. 2, and Multivalued Dependencies (MVD). An *MVD* is a statement $\sigma = (\mathbf{U} \twoheadrightarrow \mathbf{V} | \mathbf{W})$ where $\mathbf{U} \cup \mathbf{V} \cup \mathbf{W}$ form a partition of \mathbf{X} . A relation instance R *satisfies* the MVD, $R \models \sigma$, if $R = \Pi_{\mathbf{UV}}(R) \bowtie \Pi_{\mathbf{UW}}(R)$.

The *implication problem* asks whether a set of FDs and/or MVDs σ_i , $i = 1, p$, implies another FD or MVD σ_0 :

$$\sigma_1 \wedge \cdots \wedge \sigma_p \Rightarrow \sigma_0 \quad (55)$$

Armstrong's axioms [AD80] are complete for the implication problem for FDs, while Beeri et al. [BFH77] gave a complete axiomatization for both FDs and MVDs, and showed that the implication problem is decidable. In contrast, Herrmann [Her06] showed that the implication problem of *Embedded MVDs* is undecidable; we do not discuss EMVDs here.

Lee [Lee87a] showed the following connection between information theory and constraints. Fix a relational instance R , and let \mathbf{h} be its associated (uniform) entropic vector. Then $R \models \mathbf{U} \rightarrow \mathbf{V}$ iff $h(\mathbf{V} | \mathbf{U}) = 0$, and $R \models \mathbf{U} \twoheadrightarrow \mathbf{V} | \mathbf{W}$ iff $I_h(\mathbf{V}; \mathbf{W} | \mathbf{U}) = 0$. Therefore, every implication problem for FDs and MVDs can be stated as a conditional information inequality. For example, the *augmentation axiom* [BFH77] states

$$(A \rightarrow B | CD) \Rightarrow (AC \rightarrow B | D)$$

and is equivalent to the following conditional inequality:⁵

$$I(B; CD | A) = 0 \Rightarrow I(B; D | AC) = 0 \quad (56)$$

This can be proven immediately by observing that the identity $I(B; CD | A) = I(B; C | A) + I(B; D | AC)$ implies:

$$I(B; D | AC) \leq I(B; CD | A)$$

Since both terms are ≥ 0 , the implication (56) follows.

Beyond database applications, Conditional Independencies (CI) are commonly used in AI, Knowledge Representation, and Machine Learning. A CI is an assertions of the form $X \perp Y | Z$, where X, Y, Z are three random variables, stating that X is independent of Y conditioned on Z . The AI community has extensively studied the implication problem for CIs. It was shown that the implication problem is decidable and finitely axiomatizable for *saturated* CIs [GP93] (where XYZ = all variables), but not finitely axiomatizable in general [Stu90].

⁵This has the form in Def. 9.1 once we write it as $-I(B; CD | A) \geq 0$ implies $-I(B; D | AC) \geq 0$.

9.2 The Relaxation Problem

How can we prove a conditional inequality (54)? One approach is as follows. Find p non-negative real numbers $\lambda_1, \dots, \lambda_p$ for which the following inequality is valid:

$$\mathbf{c} \cdot \mathbf{h} \geq \left(\sum_{i=1,p} \lambda_i \mathbf{c}_i \right) \cdot \mathbf{h} \quad (57)$$

Then, observe that (57) implies (54). A natural question is whether every conditional information inequality can be derived in this way, from an unconditional inequality: when that is the case, then we say that the conditional inequality (54) *relaxes* to (57), or that it is *essentially unconditional*. Otherwise we say that it is *essentially conditional*. For example, the augmentation axiom above can be relaxed, hence it is essentially unconditioned.

Besides offering an important proof technique, relaxation is important in modern database applications, because often the integrity constraints don't hold *exactly*, but only *approximatively*, especially when they are mined from a given dataset [GR04, SBHR06, CIPY14, BBF⁺16, KN18, SGS18, PD11, FD16]. The relaxation problem allows us to transfer proofs over exact constraints to approximate constraints.

Every implication problem for FDs and MVDs relaxes [KS22]:

Theorem 9.2. *Consider the statement (55) asserting the implication between a set of FDs/MVDs. Consider the associated conditional information inequality:*

$$\left(\bigwedge_i h(\sigma_i) = 0 \right) \Rightarrow (h(\sigma_0) = 0) \quad (58)$$

where each expression $h(\sigma_i)$ represents either $I_h(\mathbf{V}; \mathbf{W} | \mathbf{U})$ or $h(\mathbf{V} | \mathbf{U})$. Then the following are equivalent:

- The implication (55) holds for FDs/MVDs.
- The implication (58) holds for all polymatroids.
- The implication (58) holds for all entropic functions.
- The implication (58) holds for all normal polymatroids.
- The inequality $n^2/4(\sum_i h(\sigma_i)) \geq h(\sigma_0)$ holds for all polymatroids. Furthermore, if σ_0 is an FD (rather than MVD), then $n^2/4$ can be replaced by 1.

Thus, the implication problem for FDs and MVDs relaxes. A consequence of the theorem is that, if (58) fails, then there exists a relation with only two tuples falsifying the implication, namely one of the relations $R_{\mathbf{W}}$ in (41) associated to a step function.

Next, we examine whether all conditional inequalities relax. To do this, we need another (last) deep dive into the structure of entropic functions, and introduce *almost-entropic functions*.

9.3 Background: Almost Entropic Functions

We start with a brief review of cones, following [Sch03, BV04]. A set $K \subseteq \mathbb{R}^n$ is called a *cone*, if $\mathbf{x} \in K$ and $\theta \geq 0$ implies $\theta \mathbf{x} \in K$. The cone is *convex* if $\mathbf{x}_1, \mathbf{x}_2 \in K$ and $\theta \in [0, 1]$ implies $\theta \mathbf{x}_1 + (1 - \theta) \mathbf{x}_2 \in K$. For any set $K \subseteq \mathbb{R}^n$, we denote by \bar{K} its *topological closure*, and by

$K^* \stackrel{\text{def}}{=} \{\mathbf{y} \mid \forall \mathbf{x} \in K, \mathbf{y}^T \cdot \mathbf{x} \geq 0\}$, it's *dual*. The dual is always a closed, convex cone, $K \subseteq K^{**}$ and, when K is a closed, convex cone, then $K = K^{**}$.

A cone is *polyhedral* if it has the form $K = \{\mathbf{x} \mid \mathbf{M} \cdot \mathbf{x} \geq 0\}$, for some matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$. Any polyhedral cone is closed and convex, and its dual is also polyhedral.

The set of polymatroids Γ_n and of entropic vectors Γ_n^* are subsets of \mathbb{R}^{2^n} . The superscript $*$ in Γ_n^* is an unfortunate notation, since it does not represent a dual, but this notation is already widely used. Valid inequalities (Def. 4.2) are the dual cones, $(\Gamma_n^*)^*$, and $(\Gamma_n)^*$ respectively. Clearly, Γ_n is polyhedral, and therefore $\Gamma_n = \Gamma_n^{**}$. What about Γ_n^* ?

It turns out that, when $n \geq 3$, then Γ_n^* is neither a cone nor convex. This may come as a surprise, so we take the opportunity to briefly review here the elegant proof by Zhang and Yeung [ZY97]. Consider the parity function h , shown in Fig. 2, and observe that $h(X) = 1$, $h(Z|XY) = 0$, and $I_h(X; Y) = 0$. If c is a natural number, the vector $c \cdot \mathbf{h}$ is also entropic, by Prop. 5.3; but in general, $c \cdot \mathbf{h}$ is entropic only if there exists a natural number N such that $c = \log N$, which implies Γ_n^* is neither a cone, nor convex. To prove this, assume that $c \cdot \mathbf{h}$ is entropic, and realized by a probability distribution $p(X, Y, Z)$. Choose any two values x, y such that $p(X = x) > 0$ and $p(Y = y) > 0$. Since $I_{c\mathbf{h}}(X; Y) = c \cdot I_h(X; Y) = 0$, it holds that $X \perp Y$, hence $p(X = x)p(Y = y) = p(X = x, Y = y) > 0$. Furthermore, $c \cdot h(Z|XY) = 0$, therefore p satisfies the functional dependency $XY \rightarrow Z$, and there exists a unique value z s.t. $p(X = x, Y = y, Z = z) > 0$. We have obtained $p(X = x)p(Y = y) = p(X = x, Y = y, Z = z)$ and, by symmetry, it also holds that $p(X = x)p(Z = z) = p(X = x, Y = y, Z = z)$. This implies $p(Y = y) = p(Z = z)$. Since y was arbitrary, it follows that $p(Y = y) = p(Y = y')$ for all y, y' in the support of Y . Therefore, the marginal distribution of Y is uniform, and its entropy is $c \cdot h(Y) = \log N$, where N is the size of the support, proving $c = \log N$, since $h(Y) = 1$. Recall that in Sec. 4.2 we stated that the conditional entropy $h(-|U)$ is not always an entropic vector: we invite the reader to give such an example.

While Γ_n^* is neither a cone nor convex, Yeung [Yeu08] proved:

Theorem 9.3. *The topological closure $\bar{\Gamma}_n^*$ of Γ_n^* is a closed, convex cone. A vector $\mathbf{h} \in \bar{\Gamma}^*$ is called almost-entropic.*

The complete picture of all sets of polymatroids discussed in this paper is shown in Fig. 3.

If an inequality is valid for Γ_n^* , then it is also valid for $\bar{\Gamma}_n^*$, by continuity. However, this no longer holds for *conditional* information inequalities: Kaced and Romashchenko [KR13] gave an example of a conditional inequality that is valid for Γ_n^* , but not for $\bar{\Gamma}_n^*$. Since we are interested in the relaxation problem, we will consider only validity for $\bar{\Gamma}_n^*$.

When $n = 3$ then one can show that $\bar{\Gamma}_3^* = \Gamma_3$, hence it is polyhedral. However, Matúš [Mat07] showed that, for $n \geq 4$, $\bar{\Gamma}_n^*$ is not polyhedral. This explains the difficulties in understanding the non-Shannon inequalities, and also in reasoning about conditional inequalities.

9.4 A Conditional Inequality that Does Not Relax

Kaced and Romashchenko [KR13] gave four examples of essentially conditional inequalities. This is a very surprising result: it means that a proof of the implication (55) becomes useless if one of the assumptions has even a tiny violation in the data. We describe here one of their examples, following the adaptation in [KS22].

Theorem 9.4. *[KR13] The following conditional inequality is valid for $\bar{\Gamma}_n^*$, and is essentially conditional:*

$$I(X; Y|A) = I(X; Y|B) = I(A; B) = I(A; X|Y) = 0 \Rightarrow I(X; Y) = 0$$

Notice that none of the constraints is an FD or an MVD, so this does not contradict Theorem 9.2. We prove the theorem, and start by proving the conditional inequality. For that we use the following non-Shannon inequality, by Matús [Mat07]. For every $k \geq 1$,

$$I(X; Y) \leq \frac{k+3}{2} I(X; Y|A) + I(X; Y|B) + I(A; B) + \frac{k+1}{2} I(A; X|Y) + \frac{1}{k} I(A; Y|X) \quad (59)$$

When $k = 1$, this is Zhang and Yeung's inequality (32). Matús proved (59) by induction on k , by applying the Copy Lemma 5.8 at each induction step; we omit the proof.

Since (59) holds for Γ_n^* , it also holds for $\bar{\Gamma}_n^*$. To check the conditional inequality in Thm. (9.4), let $\mathbf{h} \in \bar{\Gamma}_n^*$ such that $I(X; Y|A) = I(X; Y|B) = I(A; B) = I(A; X|Y) = 0$: then inequality (59) becomes $I(X; Y) \leq \frac{1}{k} I(A; Y|X)$ and, since k is arbitrary, it follows that $I(X; Y) = 0$.

Finally, we show that the conditional inequality does not relax, by describing, for each $\lambda > 0$, an entropic vector \mathbf{h} s.t.

$$I(X; Y) \geq \lambda(I(X; Y|A) + I(X; Y|B) + I(A; B) + I(A; X|Y)) \quad (60)$$

Let \mathbf{h} be the entropy of the following distribution:

A	B	X	Y	p
0	0	0	0	$1/2 - \varepsilon$
1	0	0	1	$1/2 - \varepsilon$
0	1	1	0	ε
1	1	0	0	ε

If ε is small enough then one can check:⁶

$$I(X; Y) = \varepsilon + O(\varepsilon^2), \quad I(A; X|Y) = O(\varepsilon^2), \\ I(X; Y|A) = I(X; Y|B) = I(A; B) = 0$$

which proves (60) for ε is small enough.

9.5 Conditional Inequalities Relax with Error Terms

However, in another twist, it turns out that every conditional inequality relaxes, if we admit a small error term. The following was proven in [KS22]:

Theorem 9.5. *Suppose that the following holds:*

$$\forall \mathbf{h} \in \bar{\Gamma}_n^*, \quad \left(\bigwedge_{i=1,p} \mathbf{c}_i \cdot \mathbf{h} \leq 0 \right) \Rightarrow \mathbf{c}_0 \cdot \mathbf{h} \leq 0 \quad (61)$$

Then, for every $\varepsilon > 0$ there exists $\lambda_1, \dots, \lambda_p \geq 0$ such that:

$$\forall \mathbf{h} \in \bar{\Gamma}_n^*, \quad \mathbf{c}_0 \cdot \mathbf{h} \leq \left(\sum_{i=1,p} \lambda_i \mathbf{c}_i \cdot \mathbf{h} \right) + \varepsilon h(\mathbf{X}) \quad (62)$$

where \mathbf{X} is the set of all n variables.

⁶Complete calculations are included in [KS22]; note that here we have swapped the roles of A and B , in order to better draw the connection to Zhang and Yeung's inequality (32).

Even with the error term, condition (62) still implies (61), because, if $\mathbf{c}_i \cdot \mathbf{h} \leq 0$ for all i , then (62) implies $\mathbf{c}_0 \cdot \mathbf{h} \leq \varepsilon h(\mathbf{X})$ and, since ε is arbitrary, we obtain $\mathbf{c}_0 \cdot \mathbf{h} \leq 0$. In fact, Matús' inequality (59) can be seen as a relaxation, with an error term, of the conditional inequality in Theorem 9.4. Theorem 9.5 shows that this was not accidental: *every* conditional inequality follows from an unconditional with an error term that tends to 0. In the next section we will show that Theorem 9.5 has a surprising application, to the proof of Theorem 5.4. Before that, we prove Theorem 9.5 by showing:

Lemma 9.6. [KS22] *Let $K \subseteq \mathbb{R}^n$ be a closed, convex cone, and $\mathbf{c}_i \in \mathbb{R}^n$, $i = 0, p$ be vectors such that the following holds:*

$$\forall \mathbf{x} \in K : \left(\bigwedge_{i=1,p} \mathbf{c}_i \cdot \mathbf{x} \leq 0 \right) \Rightarrow \mathbf{c}_0 \cdot \mathbf{x} \leq 0 \quad (63)$$

Then, for every $\varepsilon > 0$, there exists $\lambda_1, \dots, \lambda_p \geq 0$ such that:

$$\forall \mathbf{x} \in K : \mathbf{c}_0 \cdot \mathbf{x} \leq \left(\sum_{i=1,p} \lambda_i \mathbf{c}_i \right) \cdot \mathbf{x} + \varepsilon \|\mathbf{x}\|_\infty \quad (64)$$

The lemma implies the theorem, because $\bar{\Gamma}_n^*$ is a closed, convex cone, and $\|\mathbf{h}\|_\infty = h(\mathbf{X})$.

Proof. (of Lemma 9.6) Let $L \stackrel{\text{def}}{=} \{-\mathbf{c}_i \mid i = 1, p\}$. Then condition (63) says that $\mathbf{x} \in K \cap L^*$ implies $\mathbf{c}_0 \cdot \mathbf{x} \leq 0$, or, equivalently, $-\mathbf{c}_0 \in (K \cap L^*)^*$. The following holds for any closed, convex cones K_1, K_2 (see [KS22, Sec.5.3]):

$$(K_1 \cap K_2)^* = \overline{\text{conhull}(K_1^* \cup K_2^*)}$$

where $\text{conhull}(A)$ is the conic hull of a set A , i.e. the set of positive, linear combinations of vectors in A . Also, if L is finite, then $L^{**} = \text{conhull}(L)$. Therefore, $-\mathbf{c}_0$ belongs to the following set:

$$\begin{aligned} (K \cap L^*)^* &= \overline{\text{conhull}(K^* \cup L^{**})} \\ &= \overline{\text{conhull}(K^* \cup \text{conhull}(L))} \\ &= \overline{\text{conhull}(K^* \cup L)} \end{aligned}$$

For any $\varepsilon > 0$, there exists $\mathbf{e} \in \mathbb{R}^n$ such that $\|\mathbf{e}\|_1 < \varepsilon$ and:

$$-\mathbf{c}_0 + \mathbf{e} \in \text{conhull}(K^* \cap L)$$

By the definition of the conic hull, and the fact that K^* is a convex cone, we obtain that there exists $\mathbf{d} \in K^*$ and $\lambda_i \geq 0$, for $i = 1, p$ such that:

$$-\mathbf{c}_0 + \mathbf{e} = \mathbf{d} - \sum_{i=1,p} \lambda_i \mathbf{c}_i$$

We prove (64). Let $\mathbf{x} \in K$, and observe that $\mathbf{d} \cdot \mathbf{x} \geq 0$, then:

$$\left(\sum_{i=1,p} \lambda_i \mathbf{c}_i \right) \cdot \mathbf{x} - \mathbf{c}_0 \cdot \mathbf{x} + \mathbf{e} \cdot \mathbf{x} = \mathbf{d} \cdot \mathbf{x} \geq 0$$

and (64) follows from $\mathbf{e} \cdot \mathbf{x} \leq \|\mathbf{e}\|_1 \cdot \|\mathbf{x}\|_\infty \leq \varepsilon \|\mathbf{x}\|_\infty$. □

We end this section with a brief discussion of why Theorem 9.5 only holds for the set of almost entropic functions $\bar{\Gamma}_n^*$, and fails for Γ_n^* . Kaced and Romashchenko [KR13] gave an example of a conditional inequality that is valid for Γ_n^* , but not for $\bar{\Gamma}_n^*$. If Theorem 9.5 were to hold for that inequality, then the relaxed inequality (62) holds for $\bar{\Gamma}_n^*$, and then we could prove that the conditional inequality also holds for $\bar{\Gamma}_n^*$.

9.6 Proof of Theorem 5.4

We have now the machinery needed to prove Theorem 5.4, which states that $\text{Log-L-Bound}_{\Gamma_n^*}$ and $\text{Log-U-Bound}_{\Gamma_n^*}$ are asymptotically equal. What makes the proof a little difficult is the ill-behaved nature of the entropic functions Γ_n^* . To cope with that, use some help from the almost entropic functions $\bar{\Gamma}_n^*$, and prove that, here, the two bounds are equal (not just asymptotically). We state and prove the theorem for an arbitrary closed, convex cone K :

Theorem 9.7. *Fix Q, Σ, \mathbf{b} , and let K be a closed, convex cone s.t. $N_n \subseteq K \subseteq \Gamma_n$. (Recall that N_n is the set of normal polymatroids, Def. 6.2.) Then:*

$$\text{Log-L-Bound}_K(Q, \Sigma, \mathbf{b}) = \text{Log-U-Bound}_K(Q, \Sigma, \mathbf{b})$$

In particular, the lower and upper bounds are equal for $K = \bar{\Gamma}_n^*$, while they are not necessarily equal for Γ_n^* . It shows that the set $\bar{\Gamma}_n^*$ is better behaved than Γ_n^* , and that explains why it was used in prior work [GT17, KNS17] to study lower bounds.

Before we prove the theorem, we show how to use it to prove Theorem 5.4. Its proof follows from Theorem 9.7 and two identities, (65) and (66), which we prove below. The first is very simple:

$$\text{Log-U-Bound}_{\bar{\Gamma}_n^*}(Q, \Sigma, \mathbf{b}) = \text{Log-U-Bound}_{\Gamma_n^*}(Q, \Sigma, \mathbf{b}) \quad (65)$$

and follows directly from the fact that $\bar{\Gamma}_n^*$ and Γ_n^* have the same dual cone, $(\bar{\Gamma}_n^*)^* = (\Gamma_n^*)^*$; in other words, they define the same set of valid inequalities $\mathbf{c} \cdot \mathbf{h} \geq 0$.

The second equality requires a proof, and we state it as a lemma:

Lemma 9.8. *The following holds:*

$$\sup_k \frac{\text{Log-L-Bound}_{\Gamma_n^*}(Q, \Sigma, k\mathbf{b})}{\text{Log-L-Bound}_{\bar{\Gamma}_n^*}(Q, \Sigma, k\mathbf{b})} = 1 \quad (66)$$

Proof. The proof is similar to that of Theorem 5.5 item (1) in Sec. 5.2. The LHS is obviously ≤ 1 . To prove that it is ≥ 1 , denote by⁷ $L \stackrel{\text{def}}{=} \text{Log-L-Bound}_{\bar{\Gamma}_n^*}(Q, \Sigma, \mathbf{b})$, and observe that $\text{Log-L-Bound}_{\bar{\Gamma}_n^*}$ is linear in \mathbf{b} , $\text{Log-L-Bound}_{\bar{\Gamma}_n^*}(Q, \Sigma, k\mathbf{b}) = kL$, because $\bar{\Gamma}_n^*$ is a convex cone. It suffices to show that, for all $\varepsilon > 0$, there exists $\mathbf{h} \in \Gamma_n^*$ such that $\mathbf{h} \models (\Sigma, k\mathbf{b})$ and $h(\mathbf{X}) \geq (1 - \varepsilon)^2 kL$.

Start with some $\mathbf{h} \in \bar{\Gamma}_n^*$ satisfying:

$$\forall \sigma \in \Sigma : h(\sigma) \leq b_\sigma, \quad h(\mathbf{X}) = L$$

which exists by the definition of $\text{Log-L-Bound}_{\bar{\Gamma}_n^*}$ (Def. 5.1) and the fact that the set $\{\mathbf{h} \in \mathbb{R}^{2[n]} \mid \|\mathbf{h}\|_\infty \leq L\}$ is compact. Chan and Yeung's Theorem 5.9 proves that Υ_n is dense in Γ_n^* , and therefore it is also dense in $\bar{\Gamma}_n^*$. Assume w.l.o.g. that $h(\mathbf{X}) > 0$, then $g \stackrel{\text{def}}{=} \min_{\mathbf{U}, \mathbf{V}: h(\mathbf{V}|\mathbf{U}) > 0} h(\mathbf{V}|\mathbf{U}) > 0$ (the

⁷If $\text{Log-L-Bound}_{\bar{\Gamma}_n^*}(Q, \Sigma, \mathbf{b}) = \infty$ then we choose L an arbitrarily large number, and make minor adjustments to the proof; we omit the details.

smallest non-zero value of any expression $h(\mathbf{V}|\mathbf{U})$). Let $\delta \stackrel{\text{def}}{=} \varepsilon g/4$. Since Υ_n is dense, there exists $r \in \mathbb{N}$ and $\mathbf{h}^{(r)} \in \Upsilon_n$ such that $\|\frac{1}{r}\mathbf{h}^{(r)} - \mathbf{h}\|_\infty \leq \delta$. We have $h(\mathbf{X}) > 0$ hence $h(\mathbf{X}) \geq g$, therefore:

$$\begin{aligned} \frac{1}{r}h^{(r)}(\mathbf{X}) &\geq h(\mathbf{X}) - \delta = h(\mathbf{X}) - \varepsilon g/4 \\ &\geq h(\mathbf{X}) - (\varepsilon/4)h(\mathbf{X}) \geq (1 - \varepsilon)h(\mathbf{X}) \\ &= (1 - \varepsilon)L \end{aligned}$$

We claim that $\frac{1}{r}h^{(r)}(\mathbf{V}|\mathbf{U}) \leq (1 + \varepsilon/2)h(\mathbf{V}|\mathbf{U})$, for all \mathbf{U}, \mathbf{V} . If $h(\mathbf{V}|\mathbf{U}) = 0$, then \mathbf{h} satisfies the FD $\mathbf{U} \rightarrow \mathbf{V}$, and therefore $\mathbf{h}^{(r)}$ also satisfies this FD (see the note after Theorem 5.9), implying $h^{(r)}(\mathbf{V}|\mathbf{U}) = 0$. Otherwise, $h(\mathbf{V}|\mathbf{U}) \geq g$ and,

$$\begin{aligned} \frac{1}{r}h^{(r)}(\mathbf{V}|\mathbf{U}) &\leq h(\mathbf{V}|\mathbf{U}) + 2\delta = h(\mathbf{V}|\mathbf{U}) + \varepsilon g/2 \\ &\leq (1 + \varepsilon/2)h(\mathbf{V}|\mathbf{U}) \end{aligned}$$

Therefore, we have:

$$\forall \sigma \in \Sigma : h^{(r)}(\sigma) \leq (1 + \varepsilon/2)rb_\sigma \quad h^{(r)}(\mathbf{X}) \geq (1 - \varepsilon)rL$$

Finally, by the Slack lemma, $\exists k \in \mathbb{N}, \mathbf{h}' \in \Gamma_n^*$ such that:

$$\begin{aligned} \forall \sigma \in \Sigma : h'(\sigma) &\leq (1 - \varepsilon/2)kh^{(r)}(\sigma) \leq (1 - (\varepsilon/2)^2)krb_\sigma \leq krb_\sigma \\ h'(\mathbf{X}) &\geq (1 - \varepsilon)kh^{(r)}(\mathbf{X}) \geq (1 - \varepsilon)^2krL \end{aligned}$$

This completes the proof. \square

In the remainder of this section we prove Theorem 9.7.

Proof. (of Theorem 9.7) We will use the following definition from [BV04, Example 5.12]:

Definition 9.9. Let K be a proper cone (meaning: closed, convex, with a non-empty interior, and pointed i.e. $\mathbf{x}, -\mathbf{x} \in K$ implies $\mathbf{x} = 0$). A primal/dual cone program in standard form⁸ is the following:

Primal	Dual
Maximize $\mathbf{c}^T \cdot \mathbf{x}$	Minimize $\mathbf{y}^T \cdot \mathbf{b}$
where $\mathbf{A} \cdot \mathbf{x} = \mathbf{b}$ $\mathbf{x} \in K$	where $(\mathbf{y}^T \cdot \mathbf{A} - \mathbf{c}^T)^T \in K^*$

Denote by P^*, D^* the optimal value of the primal and dual respectively. *Weak duality* states that $P^* \leq D^*$, and is easy to prove. When *Slater's condition* holds, which says that there exists \mathbf{x} in the interior of K such that $\mathbf{A}\mathbf{x} = \mathbf{b}$, then *strong duality* holds too: $P^* = D^*$.

$\text{Log-L-Bound}_K(Q, \Sigma, \mathbf{b})$ and $\text{Log-U-Bound}_K(Q, \Sigma, \mathbf{b})$ can be expressed as a cone program, by letting \mathbf{A} and \mathbf{c} be the matrix and vector defined in the proof of Theorem 5.2 (thus, $\mathbf{A} \cdot \mathbf{h} = (h(\sigma))_{\sigma \in \Sigma}$ and $c_{\mathbf{X}} = 1, c_{\mathbf{U}} = 0$ for $\mathbf{U} \neq \mathbf{X}$):

Log-L-Bound_K Maximize $\mathbf{c}^T \cdot \mathbf{h}$ where $\mathbf{A} \cdot \mathbf{h} + \beta = \mathbf{b}$ $(\mathbf{h}, \beta) \in K \times \mathbb{R}_+^s$	Log-U-Bound_K Minimize $\mathbf{w}^T \cdot \mathbf{b}$ where $(\mathbf{w}^T \cdot \mathbf{A} - \mathbf{c}^T)^T \in K^*$ $\mathbf{w} \geq 0$
---	---

(67)

⁸We changed to the original formulation [BV04] by replacing \mathbf{c} with $-\mathbf{c}$, replacing \mathbf{y} with $-\mathbf{y}$.

Here β are slack variables that convert an inequality $h(\sigma) \leq b_\sigma$ into an equality $h(\sigma) + \beta_\sigma = b_\sigma$. We leave it to the reader to check that these two programs are indeed primal/dual as in Def. 9.9.

However, in general, Slater's condition need not hold for (67). For example, $K = \bar{\Gamma}_n^*$ lies in the hyperplane $h(\emptyset) = 0$, and thus has an empty interior. This could be addressed by removing the \emptyset dimension, but we have a bigger problem. Some of the constraints may be tight: when $b_\sigma = 0$, then $h(\sigma) = 0$, meaning that no feasible solution \mathbf{h} exists in the interior of K . Instead, we will define a different cone, K_0 , by using polymatroids on a lattice, as in Sec. 5.1.

Partition Σ into $\Sigma_0 = \{\sigma \mid b_\sigma = 0\}$ and $\Sigma_1 = \{\sigma \mid b_\sigma > 0\}$, and denote by \mathbf{b}_1 the restriction of \mathbf{b} to Σ_1 . In other words, Σ_0 defines a set of functional dependencies, while (Σ_1, \mathbf{b}_1) defines non-tight statistics. Let $(L_{\Sigma_0}, \subseteq)$ be the lattice of the closed sets of Σ_0 (defined in Sec. 2); we will drop the subscript and write simply (L, \subseteq) to reduce clutter. Define $F \subseteq \mathbb{R}^{2^{[n]}}$ the following cone⁹:

$$F \stackrel{\text{def}}{=} \{\mathbf{h} \in \mathbb{R}^{2^{[n]}} \mid \forall \sigma \in \Sigma_0 : h(\sigma) = 0\}$$

Let $L_0 \stackrel{\text{def}}{=} L - \{\emptyset\}$, $N \stackrel{\text{def}}{=} |L_0|$. Our cone $K_0 \subseteq \mathbb{R}_+^N$ is:

$$K_0 \stackrel{\text{def}}{=} \Pi_{L_0}(K \cap F) \tag{68}$$

The function Π_{L_0} projects a $2^{[n]}$ -dimensional vector $(h_{\mathbf{U}})_{\mathbf{U} \in 2^{\mathbf{X}}}$ to the N -dimensional vector $(h_{\mathbf{U}})_{\mathbf{U} \in L_0}$. Thus, K_0 not only removes the \emptyset dimension, but also removes all dimensions subject to a tight constraint. We prove the following:

- (1) K_0 is proper
- (2) $\text{Log-L-Bound}_K = \text{Log-L-Bound}_{K_0}$ and
 $\text{Log-U-Bound}_K = \text{Log-U-Bound}_{K_0}$
- (3) $\text{Log-L-Bound}_{K_0} = \text{Log-U-Bound}_{K_0}$

Theorem 9.7 follows from these three claims.

We start with item (1), and observe that K_0 is a closed, convex cone, because $K \cap F$ is a closed, convex cone, and Π_{L_0} is a linear isomorphism $K \cap F \rightarrow K_0$: indeed, Π_{L_0} is surjective by the definition of K_0 , and it is injective because, if $\mathbf{h}, \mathbf{h}' \in K \cap F$, then $\Pi_{L_0}(\mathbf{h}) = \Pi_{L_0}(\mathbf{h}')$ implies that, for every set \mathbf{U} , $h(\mathbf{U}) = h(\mathbf{U}^+) = h'(\mathbf{U}^+) = h'(\mathbf{U})$. It is immediate to check that K_0 is pointed, and we will show below that K_0 has a non-empty interior: this implies that it is proper.

Next, we prove item (2), and for that we will write \mathbf{h} for a vector in $K \cap F$ and write $\mathbf{h}^{(0)}$ for a vector in K_0 . For any statistics $\sigma = (\mathbf{V}|\mathbf{U}) \in \Sigma_1$, denote by $\sigma^+ = ((\mathbf{U}\mathbf{V})^+|\mathbf{U}^+)$. We say that a vector $\mathbf{h}^{(0)}$ satisfies the statistics (Σ_1, \mathbf{b}_1) , in notation $\mathbf{h}^{(0)} \models (\Sigma_1, \mathbf{b}_1)$, if $\mathbf{h}^{(0)}(\sigma^+) \leq \mathbf{b}_\sigma$ for all $\sigma \in \Sigma_1$. By definition, a vector $\mathbf{h}^{(0)} \in K_0$ satisfies the FDs Σ_0 .

Lemma 9.10. *The following holds:*

$$\text{Log-L-Bound}_K(Q, \Sigma, \mathbf{b}) = \text{Log-L-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1)$$

Equivalently, the lemma states:

$$\begin{aligned} & \sup_{\mathbf{h}} \{h(\mathbf{X}) \mid \mathbf{h} \in K, \mathbf{h} \models (\Sigma, \mathbf{b})\} \\ &= \sup_{\mathbf{h}^{(0)}} \left\{ h^{(0)}(\hat{1}) \mid \mathbf{h}^{(0)} \in K_0, \mathbf{h}^{(0)} \models (\Sigma_1, \mathbf{b}_1) \right\} \end{aligned}$$

⁹In fact F is even a vector space.

and the proof is immediate, because the projection of a vector $\mathbf{h} \in K$ satisfying (Σ, \mathbf{b}) is a vector $\mathbf{h}^{(0)} \in K_0$ satisfying (Σ_1, \mathbf{b}_1) , and, conversely, every such $\mathbf{h}^{(0)}$ is the projection of a vector \mathbf{h} .

Recall that we have assumed $N_n \subseteq K \subseteq \Gamma_n$:

Lemma 9.11. *The following holds:*

$$\text{Log-}U\text{-Bound}_K(Q, \Sigma, \mathbf{b}) = \text{Log-}U\text{-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1)$$

Proof. We need to prove:

$$\begin{aligned} & \inf_{\mathbf{w}} \left\{ \sum_{\sigma \in \Sigma} w_{\sigma} b_{\sigma} \mid K \models \sum_{\sigma \in \Sigma} w_{\sigma} h(\sigma) \geq h(\mathbf{X}) \right\} \\ &= \inf_{\mathbf{w}^{(0)}} \left\{ \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} b_{\sigma} \mid K_0 \models \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h^{(0)}(\sigma^+) \geq h^{(0)}(\mathbf{X}) \right\} \end{aligned} \quad (69)$$

A vector \mathbf{w} on the LHS defines an unconstrained inequality, while a vector $\mathbf{w}^{(0)}$ on the RHS defines a constrained inequality, because $\mathbf{w}^{(0)}$ satisfies

$$\forall \mathbf{h}^{(0)} \in K_0, \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h^{(0)}(\sigma^+) \geq h^{(0)}(\mathbf{X})$$

iff it satisfies

$$\forall \mathbf{h} \in K, \bigwedge_{\sigma \in \Sigma_0} h(\sigma) = 0 \Rightarrow \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h(\sigma^+) \geq h(\mathbf{X}) \quad (70)$$

which is a constrained inequality.

We start by showing that LHS \geq RHS in Eq. (69). For that it suffices observe that, if $\mathbf{w} = (w_{\sigma})_{\sigma \in \Sigma}$ defines a valid inequality $\sum_{\sigma \in \Sigma} w_{\sigma} h(\sigma) \geq h(\mathbf{X})$, then its projection $\mathbf{w}^{(0)} = (w_{\sigma^+})_{\sigma \in \Sigma_1}$ defines a valid constrained inequality $\sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h^{(0)}(\sigma^+) \geq h^{(0)}(\mathbf{X})$, because all the missing terms $h(\sigma)$ for $\sigma \in \Sigma_0$ are $= 0$. Therefore $\inf_{\mathbf{w}}(\dots) \geq \inf_{\mathbf{w}^{(0)}}(\dots)$

We prove now that LHS \leq RHS. Let $\mathbf{w}^{(0)}$ be a vector defining a valid constrained inequality (70). The objective value of the RHS of (69) is $\sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} b_{\sigma}$. By the relaxation theorem 9.5, for every $\varepsilon > 0$ there exists $\lambda_{\sigma} \geq 0$, for $\sigma \in \Sigma_0$, such that the following is a valid, unconstrained inequality:

$$K \models \sum_{\sigma \in \Sigma_0} \lambda_{\sigma} h(\sigma) + \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h(\sigma^+) + \varepsilon h(\mathbf{X}) \geq h(\mathbf{X})$$

This inequality is not yet of the form on the LHS of (69), because we have terms $h(\sigma^+)$ instead of $h(\sigma)$. For each such term, $h(\sigma^+) = h((\mathbf{UV})^+ | \mathbf{U}^+)$, where $\sigma = (\mathbf{V} | \mathbf{U})$, we use the following Shannon inequality:

$$\begin{aligned} h(\sigma^+) &= h((\mathbf{UV})^+ | \mathbf{U}^+) \leq h((\mathbf{UV})^+ | \mathbf{U}) \\ &= h(\mathbf{UV} | \mathbf{U}) + h((\mathbf{UV})^+ | \mathbf{UV}) \\ &= h(\sigma) + h((\mathbf{UV})^+ | \mathbf{UV}) \leq h(\sigma) + \sum_{\sigma \in \Sigma_0} h(\sigma) \end{aligned}$$

The last inequality, $h((\mathbf{UV})^+ | \mathbf{UV}) \leq \sum_{\sigma \in \Sigma_0} h(\sigma)$, can be checked by induction on the number of steps needed to compute the closure $(\mathbf{UV})^+$ using the FDs in Σ_0 . (It also follows from Theorem 9.2.)

This implies that there exists coefficients $\lambda'_\sigma \geq 0$ such that the following is a valid, unconstrained inequality:

$$K \models \sum_{\sigma \in \Sigma_0} \lambda'_\sigma h(\sigma) + \sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} h(\sigma) + \varepsilon h(\mathbf{X}) \geq h(\mathbf{X})$$

or, equivalently,

$$K \models \sum_{\sigma \in \Sigma_0} \frac{\lambda'_\sigma}{1 - \varepsilon} h(\sigma) + \sum_{\sigma \in \Sigma_1} \frac{w_{\sigma^+}^{(0)}}{1 - \varepsilon} h(\sigma) \geq h(\mathbf{X})$$

This is a valid inequality for the LHS of (69), and its objective value is $(\sum_{\sigma \in \Sigma_1} w_{\sigma^+}^{(0)} b_\sigma) / (1 - \varepsilon)$, because $b_\sigma = 0$ for all $\sigma \in \Sigma_0$. Since ε can be chosen arbitrarily small, it follows that $\text{LHS} \leq \text{RHS}$ in (69). \square

This completes the proof of item (2). It remains to prove item (3). For that we represent both $\text{Log-L-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1)$ and $\text{Log-U-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1)$ as the solutions to the primal/dual cone program (67) over the cone K_0 . Notice that the vector \mathbf{b} is restricted to \mathbf{b}_1 and, therefore, $b_\sigma > 0$ for all $\sigma \in \Sigma_1$: we no longer have tight constraints. Similarly, the matrix \mathbf{A} will be restricted to a matrix \mathbf{A}_1 whose rows correspond to the closed sets $\mathbf{U} \subseteq \mathbf{X}$. It remains to check Slater's condition, and, in particular, prove that K_0 has a non-empty interior. For that purpose we extended the definition of step functions from Sec. 6.1 to our lattice L . For each closed set $\mathbf{W} \in L$, s.t. $\mathbf{W} \neq \hat{1}$, we define the step function at \mathbf{W} as follows.

$$\forall \mathbf{U} \in L : \quad h_{\mathbf{W}}^{(0)}(\mathbf{U}) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } \mathbf{U} \subseteq \mathbf{W} \\ 1 & \text{otherwise} \end{cases}$$

Let $\mathbf{h}_{\mathbf{W}} \in K$ be the standard step function in Eq. (40) (we assumed $N_n \subseteq K$). $\mathbf{h}_{\mathbf{W}}$ satisfies all FDs Σ_0 , because the only FDs that it does not satisfy are of the form $\mathbf{U} \rightarrow \mathbf{V}$ where $\mathbf{U} \subseteq \mathbf{W}$, $\mathbf{V} \not\subseteq \mathbf{W}$, and none of the FDs in Σ_0 have this form because $\mathbf{W}^+ = \mathbf{W}$. It follows $\mathbf{h}_{\mathbf{W}} \in K \cap F$, and this proves $\mathbf{h}_{\mathbf{W}}^{(0)} = \Pi_{L_0}(\mathbf{h}_{\mathbf{W}}) \in K_0$. There are N step functions $\mathbf{h}_{\mathbf{W}}^{(0)} \in K_0$, and it is straightforward to check that they are independent vectors in \mathbb{R}^N . Let $\varepsilon > 0$ be small enough such that $2\varepsilon N < \min_{\sigma \in \Sigma_1} b_\sigma$. Define $\mathbf{h} \stackrel{\text{def}}{=} \sum_{\mathbf{W} \in L_0} \varepsilon \mathbf{h}_{\mathbf{W}}$, and $\mathbf{h}^{(0)} \stackrel{\text{def}}{=} \Pi_{L_0}(\mathbf{h})$. Since $K \cap F$ is a convex cone, $\mathbf{h} \in K \cap F$ and therefore $\mathbf{h}^{(0)} \in K_0$. We claim that there exists slack variables $\boldsymbol{\beta}$ such that $(\mathbf{h}^{(0)}, \boldsymbol{\beta})$ is a feasible solution to the cone program in (67) and, furthermore, $(\mathbf{h}^{(0)}, \boldsymbol{\beta})$ belongs to the interior of $K_0 \times \mathbb{R}_+^s$. Indeed, for all $\sigma \in \Sigma_1$, $h^{(0)}(\sigma) \leq \varepsilon N < b_\sigma$, hence, if we define $\beta_\sigma \stackrel{\text{def}}{=} b_\sigma - h^{(0)}(\sigma)$, the pair $(\mathbf{h}^{(0)}, \boldsymbol{\beta}) \in K_0 \times \mathbb{R}_+^s$ is a feasible solution to the primal (67). Next, we prove that $(\mathbf{h}^{(0)}, \boldsymbol{\beta})$ is in the interior of $K_0 \times \mathbb{R}_+^s$. Since $\beta_\sigma > 0$ for all $\sigma \in \Sigma_1$, it follows that $\boldsymbol{\beta}$ is in the interior of \mathbb{R}_+^s . Set $\mathbf{h}' \stackrel{\text{def}}{=} \sum_{\mathbf{W} \in L_0} (\varepsilon + \delta_{\mathbf{W}}) \mathbf{h}_{\mathbf{W}}$, where $\delta_{\mathbf{W}} \in (-\varepsilon, \varepsilon)$ are N arbitrary numbers, we have $\mathbf{h}' \in K \cap F$, hence $\Pi_{L_0}(\mathbf{h}') \in K_0$. This proves that $\mathbf{h}^{(0)}$ is in the interior of K_0 , thus, verifying Slater's condition. It follows that $\text{Log-L-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1) = \text{Log-U-Bound}_{K_0}(Q, \Sigma_1, \mathbf{b}_1)$, completing the proof of item (3). \square

10 Conclusions

Data is ultimately information, and therefore the connection between databases and information theory is no surprise. We have discussed applications of information inequalities to several database

theory problems: query upper bounds, query evaluation, query domination, and reasoning about approximate constraints. There are major open problems in information theory, for example the decidability of entropic information inequalities, the complexity of deciding Shannon inequalities, a characterization of the cone $\bar{\Gamma}_n^*$, and each such open problem has a corresponding open problem in database theory. In some cases the converse holds too, for example the query domination problem is computationally equivalent to checking validity of max-information inequalities, hence any proof of (un)-decidability of one problem carries over to the other.

A broader question is whether information theory can find wider applications in finite model theory. For example, functional dependencies and multivalued dependencies can be specified either using first order logic sentences, or using entropic terms. Are there other properties in finite model theory that can be captured using information theory? Such a connection would enable logical implications to be relaxed to approximate reasoning, with lots of potential in modern, data-driven applications that rely heavily on statistical reasoning.

References

- [ACK⁺11] Michael Armbrust, Kristal Curtis, Tim Kraska, Armando Fox, Michael J. Franklin, and David A. Patterson. PIQL: success-tolerant query processing in the cloud. *PVLDB*, 5(3):181–192, 2011.
- [AD80] William Ward Armstrong and Claude Delobel. Decomposition and functional dependencies in relations. *ACM Trans. Database Syst.*, 5(4):404–430, 1980.
- [AFP⁺09] Michael Armbrust, Armando Fox, David A. Patterson, Nick Lanham, Beth Trushkowsky, Jesse Trutna, and Haruki Oh. SCADS: scale-independent storage for social computing applications. In *CIDR 2009, Fourth Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA, January 4-7, 2009, Online Proceedings*, 2009.
- [AGM13] Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. *SIAM J. Comput.*, 42(4):1737–1767, 2013.
- [AL02] Marcelo Arenas and Leonid Libkin. A normal form for XML documents. In Lucian Popa, Serge Abiteboul, and Phokion G. Kolaitis, editors, *Proceedings of the Twenty-first ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 3-5, Madison, Wisconsin, USA*, pages 85–96. ACM, 2002.
- [AL05] Marcelo Arenas and Leonid Libkin. An information-theoretic approach to normal forms for relational and XML data. *J. ACM*, 52(2):246–283, 2005.
- [ALK⁺13] Michael Armbrust, Eric Liang, Tim Kraska, Armando Fox, Michael J. Franklin, and David A. Patterson. Generalized scale independence through incremental precomputation. In *SIGMOD 2013*, pages 625–636, 2013.
- [AtCG⁺15] Molham Aref, Balder ten Cate, Todd J. Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L. Veldhuizen, and Geoffrey Washburn. Design and implementation of the logicblox system. In Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives, editors, *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 1371–1382. ACM, 2015.

- [BB12] Paul Balister and Béla Bollobás. Projections, entropy and sumsets. *Comb.*, 32(2):125–141, 2012.
- [BBF⁺16] Tobias Bleifuß, Susanne Bülow, Johannes Frohnhofen, Julian Risch, Georg Wiese, Sebastian Kruse, Thorsten Papenbrock, and Felix Naumann. Approximate discovery of functional dependencies for large datasets. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, pages 1803–1812, 2016.
- [BFH77] Catriel Beeri, Ronald Fagin, and John H. Howard. A complete axiomatization for functional and multivalued dependencies in database relations. In *Proceedings of the 1977 ACM SIGMOD International Conference on Management of Data, Toronto, Canada, August 3-5, 1977.*, pages 47–61, 1977.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [CGFS86] Fan R. K. Chung, Ronald L. Graham, Peter Frankl, and James B. Shearer. Some intersection theorems for ordered sets and graphs. *J. Comb. Theory, Ser. A*, 43(1):23–37, 1986.
- [CIPY14] Xu Chu, Ihab F. Ilyas, Paolo Papotti, and Yin Ye. Ruleminer: Data quality rules discovery. In *IEEE 30th International Conference on Data Engineering, Chicago, ICDE 2014, IL, USA, March 31 - April 4, 2014*, pages 1222–1225, 2014.
- [Csi22] Laszlo Csirmaz. Around entropy inequalities. <https://www.lirmm.fr/~romashchen/seminar-aait.html>, 2022. Slides and video accessed March 2023.
- [CV93] Surajit Chaudhuri and Moshe Y. Vardi. Optimization of *Real* conjunctive queries. In *ACM PODS, 1993*, pages 59–70, 1993.
- [CY02] Terence H. Chan and Raymond W. Yeung. On a relation between information inequalities and group theory. *IEEE Transactions on Information Theory*, 48(7):1992–1995, 2002.
- [Fag82] Ronald Fagin. Horn clauses and database dependencies. *J. ACM*, 29(4):952–985, 1982.
- [Fag83] Ronald Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM*, 30(3):514–550, 1983.
- [FBS⁺20] Michael J. Freitag, Maximilian Bandle, Tobias Schmidt, Alfons Kemper, and Thomas Neumann. Adopting worst-case optimal joins in relational database systems. *Proc. VLDB Endow.*, 13(11):1891–1904, 2020.
- [FD16] Abram Friesen and Pedro Domingos. The sum-product theorem: A foundation for learning tractable models. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1909–1918, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [FK98] Ehud Friedgut and Jeff Kahn. On the number of copies of one hypergraph in another. *Israel Journal of Mathematics*, 105(1):251–256, 1998.

- [Fri04] Ehud Friedgut. Hypergraphs, entropy, and inequalities. *Am. Math. Mon.*, 111(9):749–760, 2004.
- [GLVV12] Georg Gottlob, Stephanie Tien Lee, Gregory Valiant, and Paul Valiant. Size and treewidth bounds for conjunctive queries. *J. ACM*, 59(3):16:1–16:35, 2012.
- [GM06] Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 289–298. ACM Press, 2006.
- [GM14] Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. *ACM Trans. Algorithms*, 11(1):4:1–4:20, 2014.
- [GP93] Dan Geiger and Judea Pearl. Logical and algorithmic properties of conditional independence and graphical models. *The Annals of Statistics*, 21(4):2001–2021, 1993.
- [GR04] Chris Giannella and Edward L. Robertson. On approximation measures for functional dependencies. *Inf. Syst.*, 29(6):483–507, 2004.
- [GT17] Tomasz Gogacz and Szymon Torunczyk. Entropy bounds for conjunctive queries with functional dependencies. In Michael Benedikt and Giorgio Orsi, editors, *20th International Conference on Database Theory, ICDT 2017, March 21-24, 2017, Venice, Italy*, volume 68 of *LIPIcs*, pages 15:1–15:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [Her06] Christian Herrmann. Corrigendum to ”on the undecidability of implications between embedded multivalued database dependencies” [inform. and comput. 122(1995) 221-235]. *Inf. Comput.*, 204(12):1847–1851, 2006.
- [IR95] Yannis E. Ioannidis and Raghu Ramakrishnan. Containment of conjunctive queries: Beyond relations as sets. *ACM Trans. Database Syst.*, 20(3):288–324, 1995.
- [JKV06] T. S. Jayram, Phokion G. Kolaitis, and Erik Vee. The containment problem for REAL conjunctive queries with inequalities. In *ACM PODS, 2006*, pages 80–89, 2006.
- [KKNS20] Mahmoud Abo Khamis, Phokion G. Kolaitis, Hung Q. Ngo, and Dan Suciu. Decision problems in information theory. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 106:1–106:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [KKNS21] Mahmoud Abo Khamis, Phokion G. Kolaitis, Hung Q. Ngo, and Dan Suciu. Bag query containment and information theory. *ACM Trans. Database Syst.*, 46(3):12:1–12:39, 2021.
- [KMP⁺20] Batya Kenig, Pranay Mundra, Guna Prasaad, Babak Salimi, and Dan Suciu. Mining approximate acyclic schemes from relations. In David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo, editors, *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, pages 297–312. ACM, 2020.

- [KN18] Sebastian Kruse and Felix Naumann. Efficient discovery of approximate dependencies. *PVLDB*, 11(7):759–772, 2018.
- [KNS16] Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. Computing join queries with functional dependencies. In Tova Milo and Wang-Chiew Tan, editors, *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, pages 327–342. ACM, 2016.
- [KNS17] Mahmoud Abo Khamis, Hung Q. Ngo, and Dan Suciu. What do shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? In Emanuel Sallinger, Jan Van den Bussche, and Floris Geerts, editors, *Proceedings of the 36th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2017, Chicago, IL, USA, May 14-19, 2017*, pages 429–444. ACM, 2017. Extended version available at <http://arxiv.org/abs/1612.02503>.
- [KR11] Swastik Kopparty and Benjamin Rossman. The homomorphism domination exponent. *Eur. J. Comb.*, 32(7):1097–1114, 2011.
- [KR13] Tarik Kaced and Andrei E. Romashchenko. Conditional information inequalities for entropic and almost entropic points. *IEEE Trans. Inf. Theory*, 59(11):7149–7167, 2013.
- [KS22] Batya Kenig and Dan Suciu. Integrity constraints revisited: From exact to approximate implication. *Log. Methods Comput. Sci.*, 18(1), 2022.
- [Lee87a] Tony T. Lee. An information-theoretic analysis of relational databases - part I: data dependencies and information metric. *IEEE Trans. Software Eng.*, 13(10):1049–1061, 1987.
- [Lee87b] Tony T. Lee. An information-theoretic analysis of relational databases - part II: information structures of database schemas. *IEEE Trans. Software Eng.*, 13(10):1061–1072, 1987.
- [LW49] L. H. Loomis and H. Whitney. An inequality related to the isoperimetric inequality. *Bull. Am. Math. Soc.*, 55:961–962, 1949.
- [Mar13] Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6):42:1–42:51, 2013.
- [Mat07] Frantisek Matús. Infinitely many information inequalities. In *IEEE International Symposium on Information Theory, ISIT 2007, Nice, France, June 24-29, 2007*, pages 41–44. IEEE, 2007.
- [MKS21] Amine Mhedhbi, Chathura Kankanamge, and Semih Salihoglu. Optimizing one-time and continuous subgraph queries using worst-case optimal joins. *ACM Trans. Database Syst.*, 46(2):6:1–6:45, 2021.
- [NPRR12] Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. Worst-case optimal join algorithms: [extended abstract]. In Michael Benedikt, Markus Krötzsch, and Maurizio Lenzerini, editors, *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2012, Scottsdale, AZ, USA, May 20-24, 2012*, pages 37–48. ACM, 2012.

- [NPRR18] Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. Worst-case optimal join algorithms. *J. ACM*, 65(3):16:1–16:40, 2018.
- [NRR13] Hung Q. Ngo, Christopher Ré, and Atri Rudra. Skew strikes back: new developments in the theory of join algorithms. *SIGMOD Rec.*, 42(4):5–16, 2013.
- [PD11] Hoifung Poon and Pedro M. Domingos. Sum-product networks: A new deep architecture. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, July 14-17, 2011*, pages 337–346, 2011.
- [Pip86] Nicholas Pippenger. What are the laws of information theory. In *1986 Special Problems on Communication and Computation Conference*, pages 3–5, 1986.
- [Rom22] Andrei E. Romashchenko. The zhang-yeung inequality: an attempt to explain. <https://www.lirmm.fr/~romashchen/seminar-aait.html>, 2022. Slides and video accessed March 2023.
- [SBHR06] Yannis Sismanis, Paul Brown, Peter J. Haas, and Berthold Reinwald. GORDIAN: efficient and scalable discovery of composite keys. In *Proceedings of the 32nd International Conference on Very Large Data Bases, Seoul, Korea, September 12-15, 2006*, pages 691–702, 2006.
- [Sch03] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [SGS18] Babak Salimi, Johannes Gehrke, and Dan Suciu. Bias in OLAP queries: Detection, explanation, and removal. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, pages 1021–1035, 2018.
- [SOC16] Maximilian Schleich, Dan Olteanu, and Radu Ciucanu. Learning linear regression models over factorized joins. In Fatma Özcan, Georgia Koutrika, and Sam Madden, editors, *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, pages 3–18. ACM, 2016.
- [Stu90] Milan Studený. Conditional independence relations have no finite complete characterization. In *11th Prague Conf. Information Theory, Statistical Decision Foundation and Random Processes*, pages 377–396. Norwell, MA, 1990.
- [Vel14] Todd L. Veldhuizen. Triejoin: A simple, worst-case optimal join algorithm. In Nicole Schweikardt, Vassilis Christophides, and Vincent Leroy, editors, *Proc. 17th International Conference on Database Theory (ICDT), Athens, Greece, March 24-28, 2014*, pages 96–106. OpenProceedings.org, 2014.
- [Wan22] Yilei Wang. Personal communication, 2022.
- [WWS23] Yisu Remy Wang, Max Willsey, and Dan Suciu. Free join: Unifying worst-case optimal and traditional joins. *CoRR*, abs/2301.10841, 2023.
- [Yeu08] Raymond W. Yeung. *Information Theory and Network Coding*. Springer Publishing Company, Incorporated, 1 edition, 2008.

- [ZY97] Zhen Zhang and Raymond W. Yeung. A non-shannon-type conditional inequality of information quantities. *IEEE Trans. Information Theory*, 43(6):1982–1986, 1997.
- [ZY98] Zhen Zhang and Raymond W. Yeung. On characterization of entropy function via information inequalities. *IEEE Transactions on Information Theory*, 44(4):1440–1452, 1998.

.1 Proof of the claim in Example 4.9

We prove that any valid Σ -information inequality is a positive linear combination of the four inequalities shown in Example 4.9. Such an inequality has the following form:

$$\begin{aligned} w_R h(XY) + w_S h(YZ) + w_T h(ZU) \\ + w_B h(U|XZ) + w_A h(X|YU) \geq h(XYZU) \end{aligned} \quad (71)$$

where $w_R, \dots, w_A \geq 0$ are non-negative real numbers. Since the inequality holds for all polymatroids, it also holds for every step function $\mathbf{h}^{\mathbf{V}}$ (see Eq. (42)), for all $\mathbf{V} \subseteq \{X, Y, Z, U\}$. There are $2^4 - 1 = 15$ step functions, but we only use 5 of them:

$$\begin{aligned} \mathbf{h}^X : & w_R + w_B \geq 1 \\ \mathbf{h}^Y : & w_R + w_S \geq 1 \\ \mathbf{h}^Z : & w_S + w_T \geq 1 \\ \mathbf{h}^U : & w_T + w_A \geq 1 \\ \mathbf{h}^{XU} : & w_R + w_T \geq 1 \end{aligned}$$

Consider the three constraints for $\mathbf{h}^Y, \mathbf{h}^Z, \mathbf{h}^{XU}$, which mention only the variables w_R, w_S, w_T . Any solution to these three constraints can be immediately extended to a solution to all 5 constraints, by setting $w_B \geq \max(0, 1 - w_R)$ and $w_A \geq \max(0, 1 - w_T)$. On the other hand, the three constraints on w_R, w_S, w_T assert that they form a fractional edge cover of a triangle. The fractional edge covering polytope of a triangle has four extreme vertices, $(0, 1, 1), (1, 0, 1), (1, 1, 0), (1/2, 1/2, 1/2)$. It follows that the extreme vertices of our polytope over all 5 variables are:

w_R	w_S	w_T	w_A	w_B
0	1	1	0	1
1	0	1	0	0
1	1	0	1	0
1/2	1/2	1/2	1/2	1/2

Each of these vectors corresponds precisely to one of the four inequalities that we listed in Example 4.9, and, conversely, any Σ -inequality of the form (71) is dominated by some convex combination of one of these four.

.2 Proof of (33), and Failure of the Copy Lemma

We start by proving Inequality (33). The proof follows immediately from the following identity:

$$\begin{aligned}
& -I(X; Y) + I(X; Y|A) + I(X; Y|B) + I(A : B) + \\
& \quad + I(X; Y|A') + I(A'; Y|X) + I(A'; X|Y) + 3I(A'; AB|XY) \\
& = I(A; B|A') + I(A; A'|Y) + I(A; A'|X) \\
& \quad + I(A; A'|BXY) + I(B; A'|Y) + I(B; A'|X) \\
& \quad + I(B; A'|AXY) + I(X; Y|BA') + I(X; Y|AA') \\
& \quad + I(X; A'|ABY) + I(Y; A'|AB)
\end{aligned}$$

Next, we prove that the polymatroid in Fig. 5 does not satisfy the Copy Lemma. Assuming otherwise, let h' denote the polymatroid over variables X, Y, A, B, A', B' . Using only basic Shannon inequalities, we derive a contradiction. We will drop the index h' from $I_{h'}(\dots)$ and write simply $I(\dots)$. By assumption, the values $h'(\mathbf{U})$ for all sets \mathbf{U} that do not contain both A and A' , or both B and A' are known, for example $h'(A'|XY) = h(A|XY) = 1$. We also know the value $h'(\mathbf{U})$ when \mathbf{U} contain XY , example $h'(AA'XY) = h'(AA'|XY) + h'(XY) = h'(A|XY) + h'(A'|XY) + h(XY) = 1 + 1 + 3 = 5$. We proceed by examining the other sets where A, A' or B, B' co-occur, and start by showing $h'(AA') \leq 3$:

$$\begin{aligned}
I(A; A'|X) &= h'(AX) + h'(A'X) - h'(X) - h'(AA'X) \\
&= 3 + 3 - 2 - h'(AA'X) \geq 0
\end{aligned}$$

and we derive $h'(AA'X) \leq 4$. Similarly (replacing X with Y) we derive $h'(AA'Y) \leq 4$. Finally, we have:

$$\begin{aligned}
I(X; Y|AA') &= \\
&= h'(AA'X) + h'(AA'Y) - h'(AA'XY) - h'(AA') \geq 0
\end{aligned}$$

and we derive that $h'(AA') \leq 3$. We repeat the argument above by replacing A with B , and derive similarly that $h'(BA') \leq 3$. Next, we show that $h'(ABA') \geq 5$, which follows from:

$$\begin{aligned}
I(XY; A'|AB) &= \\
&= h'(ABXY) + h'(ABA') - h'(AB) - h'(ABA'XY) \\
&= 4 + h'(ABA') - 4 - 5 \geq 0
\end{aligned}$$

thus $h'(ABA') \geq 5$. (We also have $h'(ABA') \leq h'(ABA'XY) = 5$, hence $h'(ABA') = 5$, but the inequality suffices for us.) Finally, we derive a contradiction:

$$\begin{aligned}
I(A; B|A') &= h'(AA') + h'(BA') - h'(A') - h'(ABA') \\
&\leq 3 + 3 - 2 - 5 = -1
\end{aligned}$$

.3 Addendum to Theorem 5.9

We briefly sketch here the proof that, if an entropic function \mathbf{h} satisfies a set of functional dependencies, then so do all $\mathbf{h}^{(r)}$, for all $r \geq 0$. For that we need to review the main argument of the proof in [CY02].

Let \mathbf{h} be an entropic function, realized by a probability distribution (R, p) . The first step is to ensure that the probabilities $p(t)$ can be assumed to be rational numbers. Assume w.l.o.g. that R is the support of p , then, by Lee's result [Lee87a], $h \models U \rightarrow V$, iff $R \models U \rightarrow V$. Consider now any sequence of probability distributions on R , $p^{(k)} : R \rightarrow [0, 1]$, of rational numbers, such that $\lim_k p^{(k)} = p$. Then $p^{(k)}$, and its entropic vector $\mathbf{h}^{(k)}$, continue to satisfy the same FDs as R and, thus, the same FDs as \mathbf{h} . Since $\mathbf{h}^{(k)}$ can be arbitrarily close to \mathbf{h} , it suffices to prove that the theorem holds for an entropic vector \mathbf{h} realized by a probability distribution (R, p) where the probabilities are rational numbers. Assume they have a common denominator $q > 0$, and let $N = |R|$.

From here on, we follow Chan and Yeung's proof [CY02]. For each $r = q, 2q, 3q, \dots$ define the following $r \times n$ matrix $\mathbf{M}_r = (m_{\rho i})_{\rho=1, r; i=1, n}$. Its rows are copies of the tuples in R , where each tuple $\mathbf{x} \in R$ occurs $r \cdot p(\mathbf{x})$ times in the matrix \mathbf{M}_r . Intuitively, \mathbf{M}_r can be viewed as a relation with n attributes and r tuples, including duplicates, whose uniform probability distribution has the same entropic vector \mathbf{h} as (R, p) . Let G be the symmetric group S_r , i.e. the group of permutation on the set $\{1, 2, \dots, r\}$; one should think of G as the group of permutations on the rows of \mathbf{M}_r . For each $i = 1, \dots, n$, let G_i the subgroup that leaves the column i invariant, in other words:

$$G_i = \{\sigma \in G \mid m_{\sigma(\rho), i} = m_{\rho, i}, \forall \rho = 1, r\}$$

Denoting similarly G_α the subgroup of permutations that leave the set of columns $\alpha \subseteq [n]$ invariant, one can check that $G_\alpha = \bigcap_{i \in \alpha} G_i$. Let $\mathbf{h}^{(r)}$ be the entropy of the uniform probability distribution on the relational instance $\{(aG_1, \dots, aG_n) \mid a \in G\}$. Using a combinatorial argument, Chan and Yeung [CY02] prove that $\lim_{r \rightarrow \infty} \frac{1}{r} \mathbf{h}^{(r)} = \mathbf{h}$. We will not repeat that argument here, but make the additional observation that, if R satisfies the FD $U \rightarrow V$, then $G_U \subseteq G_V$, which implies that $\mathbf{h}^{(r)}$ also satisfies the same FD.

.4 Proof of Equation (44)

Möbius inversion formula states that, if $f, g : 2^{\mathbf{X}} \rightarrow \mathbb{R}$ are two set functions, and one of the identities below holds, then so does the other:

$$f(U) = \sum_{V \subseteq U} g(V) \qquad g(U) = \sum_{V \subseteq U} (-1)^{|U-V|} f(V) \quad (72)$$

It is immediate to derive that the following identities are also equivalent:¹⁰

$$f(U) = \sum_{V \subseteq \mathbf{X}-U} g(V) \qquad g(U) = \sum_{V : \mathbf{X}-U \subseteq V} (-1)^{|V \cap U|} f(V)$$

To prove Equation (44), we use Equation (45), and the fact that $h(\mathbf{X}) = \sum_{V : V \subseteq \mathbf{X}} a_V$:

$$\begin{aligned} h(U) &= \sum_{V : V \cap U \neq \emptyset} a_V = h(\mathbf{X}) - \sum_{V : V \subseteq \mathbf{X}-U} a_V \\ h(\mathbf{X}|U) &= \sum_{V : V \subseteq \mathbf{X}-U} a_V \end{aligned}$$

¹⁰Define $h(U) \stackrel{\text{def}}{=} f(\mathbf{X} - U)$ then (72) becomes:

$$h(U) = \sum_{V \subseteq \mathbf{X}-U} g(V) \qquad g(U) = \sum_{V \subseteq U} (-1)^{|U-V|} h(\mathbf{X} - V)$$

The claim follows by replacing V with $\mathbf{X} - V$ in the second equation, then renaming h to f .

Möbius' inversion formula implies:

$$a_U = \sum_{V: X-U \subseteq V} (-1)^{|V \cap U|} h(\mathbf{X}|V)$$

If $\mathbf{X} - \mathbf{U} \subseteq \mathbf{V}$, then we can write \mathbf{V} uniquely as $(\mathbf{X} - \mathbf{U}) \cup \mathbf{V}_0$, where $\mathbf{V}_0 \subseteq \mathbf{U}$. After renaming \mathbf{V}_0 to \mathbf{V} we derive:

$$\begin{aligned} a_U &= \sum_{V \subseteq U} (-1)^{|V|} h(\mathbf{X}|(\mathbf{X} - \mathbf{U}) \cup V) \\ &= \sum_{V \subseteq U} (-1)^{|V|} h(\mathbf{X}) - \sum_{V \subseteq U} (-1)^{|V|} h((\mathbf{X} - \mathbf{U}) \cup V) \\ &= - \sum_{V \subseteq U} (-1)^{|V|} h((\mathbf{X} - \mathbf{U}) \cup V) \\ &= - \sum_{V \subseteq U} (-1)^{|V|} (h((\mathbf{X} - \mathbf{U}) \cup V) - h(\mathbf{X} - \mathbf{U})) \\ &= - \sum_{V \subseteq U} (-1)^{|V|} h(V|\mathbf{X} - \mathbf{U}) \end{aligned}$$

We used twice the fact that $\sum_{V \subseteq U} (-1)^{|V|} = 0$ when $U \neq \emptyset$. This completes the proof.

.5 Proof of Lemma 6.8

To prove the lemma we need to establish a simple fact:

Proposition .1. *Let $\alpha_1, \dots, \alpha_n \geq 0$ be non-negative real numbers. Define the following set function $\mathbf{h} : 2^{\mathbf{X}} \rightarrow \mathbb{R}$:*

$$h(\mathbf{U}) \stackrel{\text{def}}{=} \max_{i: X_i \in \mathbf{U}} \alpha_i$$

Then \mathbf{h} is a normal polymatroid.

Proof. Set $\alpha_0 \stackrel{\text{def}}{=} 0$ and assume w.l.o.g. that $\alpha_0 \leq \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$. Define $\delta_i \stackrel{\text{def}}{=} \alpha_i - \alpha_{i-1}$ for $i = 1, n$. We prove:

$$\mathbf{h} = \sum_{i=1, n} \delta_i \mathbf{h}^{\mathbf{X}_{[i+1:n]}} \tag{73}$$

If k is the largest index s.t. $X_k \in \mathbf{U}$, then:

$$\sum_{i=1, n} \delta_i \mathbf{h}^{\mathbf{X}_{[i+1:n]}}(\mathbf{U}) = \sum_{i=1, k} \delta_i = \alpha_k = h(\mathbf{U})$$

which proves (73). Since $\delta_i \geq 0, \forall i$, \mathbf{h} is a normal polymatroid. \square

Proof. (of Lemma 6.8) We prove the claim by induction on the number of variables $n = |\mathbf{X}|$. When $n = 0$ then the claim holds vacuously, so assume $n \geq 1$. We will use the following identity:

$$\forall \mathbf{U} \subseteq \mathbf{X} : h(\mathbf{U}) = h(\mathbf{U}|X_n) + I(\mathbf{U}; X_n)$$

Consider the following two set functions $2^{\mathbf{X}-\{X_n\}} \rightarrow \mathbb{R}_+$:

$$h_1(\mathbf{U}) \stackrel{\text{def}}{=} h(\mathbf{U}|X_n) \qquad h_2(\mathbf{U}) \stackrel{\text{def}}{=} \max_{X_i \in \mathbf{U}} I(X_i; X_n)$$

Since \mathbf{h}_1 , is a polymatroid in $n-1$ variables, by induction hypothesis we obtain a normal polymatroid \mathbf{h}'_1 satisfying properties (a),(b),(c). The second function is already a normal polymatroid, by Prop. .1. Observe that $h_2(\mathbf{U}) \leq I(\mathbf{U}; X_n)$ for any set \mathbf{U} (see Eq. 31 in Prop. 5.6). Define:

$$a \stackrel{\text{def}}{=} \min_i (h(X_n) - I(X_i; X_n)) = h(X_n) - h_2(\mathbf{X})$$

and observe that:

$$0 \leq a \leq h(X_n) - I(\mathbf{U}; X_n), \quad \forall \mathbf{U} \subseteq \mathbf{X} - \{X_n\}$$

Since both \mathbf{h}'_1 and \mathbf{h}_2 are normal polymatroids, we can write them as:

$$\mathbf{h}'_1 = \sum_{\mathbf{V} \subseteq \mathbf{X} - \{X_n\}} b_{\mathbf{V}} \mathbf{h}^{\mathbf{V}} \qquad \mathbf{h}_2 = \sum_{\mathbf{V} \subseteq \mathbf{X} - \{X_n\}} c_{\mathbf{V}} \mathbf{h}^{\mathbf{V}}$$

where $b_{\mathbf{V}}, c_{\mathbf{V}}$ are non-negative coefficients. Define \mathbf{h}' as:

$$\mathbf{h}' \stackrel{\text{def}}{=} \sum_{\mathbf{V} \subseteq \mathbf{X} - \{X_n\}} b_{\mathbf{V}} \mathbf{h}^{\mathbf{V}} + \sum_{\mathbf{V} \subseteq \mathbf{X} - \{X_n\}} c_{\mathbf{V}} \mathbf{h}^{\mathbf{V} \cup \{X_n\}} + a \mathbf{h}^{\mathbf{X}}$$

We claim that \mathbf{h}' satisfies the conditions of the lemma. Obviously, \mathbf{h}' is a normal polymatroid, it remains to check conditions (a),(b),(c). Observe that the following identities hold, for all $\mathbf{U} \subseteq \mathbf{X} - \{X_n\}$:

$$\begin{aligned} h'(\mathbf{U}) &= h'_1(\mathbf{U}) + h_2(\mathbf{U}) \\ h'(\mathbf{U} \cup \{X_n\}) &= h'_1(\mathbf{U}) + h_2(\mathbf{X} - \{X_n\}) + a \\ &= h'_1(\mathbf{U}) + h(X_n) \end{aligned}$$

We check condition (a). If \mathbf{U} does not contain X_n , then

$$h'(\mathbf{U}) = h'_1(\mathbf{U}) + h_2(\mathbf{U}) \leq h(\mathbf{U}|X_n) + I(\mathbf{U}; X_n) = h(\mathbf{U})$$

If $X_n \in \mathbf{U}$ then:

$$\begin{aligned} h'(\mathbf{U}) &= h'_1(\mathbf{U} - \{X_n\}) + h(X_n) \\ &\leq h(\mathbf{U} - \{X_n\}|X_n) + h(X_n) = h(\mathbf{U}) \end{aligned}$$

Next, we check condition (b):

$$\begin{aligned} h'(\mathbf{X}) &= h'_1(\mathbf{X} - \{X_n\}) + h(X_n) \\ &= h(\mathbf{X} - \{X_n\}|X_n) + h(X_n) = h(\mathbf{X}) \end{aligned}$$

Finally, we check condition (c) for $X_i, i < n$:

$$\begin{aligned} h'(X_i) &= h'_1(X_i) + h_2(X_i) = h_1(X_i) + h_2(X_i) \\ &= h(X_i|X_n) + I(X_i; X_n) = h(X_i) \end{aligned}$$

and finally for X_n : $h'(X_n) = h'_1(\emptyset) + h(X_n)$. □