
INTEGRITY CONSTRAINTS REVISITED: FROM EXACT TO APPROXIMATE IMPLICATION

BATYA KENIG^a AND DAN SUCIU^b

^a Technion
e-mail address: batyak@technion.ac.il

^b University of Washington
e-mail address: suciu@cs.washington.edu

ABSTRACT. Integrity constraints such as functional dependencies (FD) and multi-valued dependencies (MVD) are fundamental in database schema design. Likewise, probabilistic conditional independences (CI) are crucial for reasoning about multivariate probability distributions. The implication problem studies whether a set of constraints (antecedents) implies another constraint (consequent), and has been investigated in both the database and the AI literature, under the assumption that all constraints hold *exactly*. However, many applications today consider constraints that hold only *approximately*. In this paper we define an approximate implication as a linear inequality between the degree of satisfaction of the antecedents and consequent, and we study the *relaxation problem*: when does an exact implication relax to an approximate implication? We use information theory to define the degree of satisfaction, and prove several results. First, we show that any implication from a set of data dependencies (MVDs+FDs) can be relaxed to a simple linear inequality with a factor at most quadratic in the number of variables; when the consequent is an FD, the factor can be reduced to 1. Second, we prove that there exists an implication between CIs that does not admit any relaxation; however, we prove that every implication between CIs relaxes “in the limit”. Then, we show that the implication problem for differential constraints in market basket analysis also admits a relaxation with a factor equal to 1. Finally, we show how some of the results in the paper can be derived using the *I-measure* theory, which relates between information theoretic measures and set theory. Our results recover, and sometimes extend, several previously known results about the implication problem: the implication of MVDs and the implication of differential constraints for frequent item sets can be checked by considering only 2-tuple relations.

1. INTRODUCTION

Traditionally, integrity constraints are assertions about a database that are stated by the database administrator and enforced by the system during updates. However, in several applications of Big Data, integrity constraints are discovered, or mined in a database instance, as opposed to being asserted by the administrator [GR04, SBHR06, CIPY14, BBF⁺16, KN18]. For example, data cleaning can be done by first learning conditional functional dependencies in some reference data, then using them to identify inconsistencies in the test data [IC15, CIPY14]. Causal reasoning [SGS00, PE08, SGS18] and learning sum-of-product networks [PD11, FD16, MVM⁺18] repeatedly discover conditional independencies

Key words and phrases: Integrity constraints, the implication problem.

Cone	Relaxation Bounds			
	General	MVDs+FDs \Rightarrow FD	MVDs+FDs \Rightarrow any	Disjoint MVDs+FDs \Rightarrow MVD/FD
Γ_n	$(2^n)!$ (Thm. 6.1)	1 (Thm. 4.1)	$\frac{n^2}{4}$ (Thm. 4.1)	1 (Thm. 4.11)
Γ_n^*	∞ (Thm. 5.1)	1 (Thm. 4.1)	$\frac{n^2}{4}$ (Thm. 4.1)	1 (Thm. 4.11)
\mathcal{P}_n	1 (Thm. 7.4)	1 (Thm. 7.4)	1 (Thm. 7.4)	1 (Thm. 7.4)

TABLE 1. Summary of results: relaxation bounds for the implication $\Sigma \Rightarrow \tau$ for the sub-cones of Γ_n under various restrictions. (1) *General*; no restrictions to either Σ or τ . (2) Σ is a set of saturated CIs and conditional entropies (i.e., MVDs+FDs in databases), and τ is a conditional entropy. (3) Σ is a set of saturated CIs and conditional entropies, τ is *any* CI. (4) *Disjoint* integrity constraints where the terms in Σ are both saturated and *disjoint* (see Definition 4.6 in Sec. 4), and τ is saturated.

in the data. Constraints also arise in many other domains, for example in the *frequent itemset problem* [Li05, CP01], or as *measure based constraints* [SGG08] in applications like Dempster-Shafer theory, possibilistic theory, and game theory (see discussion in [SGG08]). In all these applications, quite often the constraints are learned from the data, and are not required to hold exactly, but it suffices if they hold only to a certain degree.

The classical *implication problem* asks whether a set of constraints, called the *antecedents*, logically imply another constraint called the *consequent*. In this setting, both antecedents and consequent are required to hold exactly, hence we refer to it as an *exact implication* (EI). The database literature has extensively studied the EI problem for integrity constraints and shown that the implication problem is decidable and axiomatizable for Functional Dependencies (FDs) and Multivalued Dependencies (MVDs) [Mai83, KLV13, AD80, BFH77], and undecidable for Embedded Multivalued Dependencies (EMVDs) [Her06]. The Armstrong axioms, for instance, provide a sound and complete system for inferring FDs when the antecedents are FDs [AD80], and Beeri et al. [BFH77] present a set of axioms which are sound and complete for inferring MVDs and FDs when the antecedents are a set of FDs and MVDs.

Conditional Independencies (CI) are assertions of the form $X \perp Y \mid Z$, stating that X is independent of Y conditioned on Z . A CI is *saturated* if it covers all variables in the joint distribution (i.e., $XYZ = \text{all variables}$); it is called a *conditional* if it has the form $X \rightarrow Y$ and Y is a function of X . The AI community has extensively studied the EI problem for Conditional Independencies (CI), and has shown that the implication problem is decidable and axiomatizable for saturated CIs [GP93] (where $XYZ = \text{all variables}$). On the negative side, it has been shown that in the general case, there exists no finite, sound and complete inference system for the implication of CIs [Stu90]. In the Frequent Itemset Problem, a constraint like $X \rightarrow Y \vee Z \vee U$ means that every basket that contains X also contains at least one of Y, Z, U , and the implication problem here is also decidable and axiomatizable [SVG05].

The Relaxation Problem. In this paper we consider a new problem, called the *relaxation problem*: if an exact implication holds, does an approximate implication hold too? For

example, suppose we prove that a given set of FDs implies another FD, but the input data satisfies the antecedent FDs only to some degree: to what degree does the consequent FD hold in the database? The relaxation problem asks whether we can convert an exact implication into an *approximate implication*. In other words, we ask how the error in the antecedents propagates to the error in its consequent. When relaxation holds, then the error to the consequent can be bounded, and any inference system for proving exact implication, e.g. using a set of axioms or some algorithm, can be used to infer an approximate implication.

In order to study the relaxation problem we need to measure the degree of satisfaction of a constraint. In this paper we use Information Theory. This is the natural semantics for modeling CIs of multivariate distributions, because $X \perp Y \mid Z$ iff $I(X; Y|Z) = 0$ where I is the *conditional mutual information*. FDs and MVDs are special cases of CIs [Lee87, DR00, WBW00] (reviewed in Section 2.1), and thus they are naturally modeled using the information theoretic measure $I(X; Y|Z)$ or $H(Y|X)$; in contrast, EMVDs do not appear to have a natural interpretation using information theory, and we will not discuss them further in this work. An *approximate implication* (APXI) is an inequality that (numerically) bounds the information-theoretic measure of the consequent (e.g., conditional entropy $H(\cdot|\cdot)$ if it is an FD, or the conditional mutual information $I(\cdot; \cdot|\cdot)$ if it is an MVD) by a linear combination of the information-theoretic measures of the antecedents. The link between integrity constraints and information theoretic measures is, in itself, not new. Several papers have argued that information theory is a suitable tool to express integrity constraints [Lee87, DR00, WBW00, Mal86, GR04].

An exact implication (EI) becomes an assertion of the form $(\sigma_1 = 0 \wedge \sigma_2 = 0 \wedge \dots) \Rightarrow (\tau = 0)$, while an approximate implication (APXI) is a linear inequality $\tau \leq \lambda \cdot (\sum \sigma_i)$, where $\lambda \geq 0$, and $\tau, \sigma_1, \sigma_2, \dots$ are information theoretic measures. We say that a class of constraints can be *relaxed* if every EI where the antecedents are from this class, implies the corresponding APXI; we also say that this EI admits a λ -relaxation, when we want to specify the factor λ in the APXI. By the non-negativity of the Shannon information measures (described in Section 2), we get that an APXI always implies the corresponding EI.

Results. We make several contributions, summarized in Table 1. We start by showing in Section 4 that MVDs+FDs admit an $n^2/4$ -relaxation, where n is the number of variables. When the consequent is an FD, we show that implication admits a 1-relaxation. Thus, whenever an exact implication holds between MVD+FDs, a simple linear inequality also holds between their associated information theoretic terms. In fact, we prove a stronger result that holds for CIs in general, which implies the result for MVDs+FDs. In Section 8, we further show that under some additional syntactic restrictions to the antecedents, the bound can be tightened: from a $n^2/4$ -relaxation to a 1-relaxation, even when the consequent is not an FD. We leave open the question of whether a 1-relaxation exists in general.

So far, we have restricted ourselves to saturated or conditional CIs (which correspond to MVDs or FDs). In Section 5 we remove all restrictions, and prove a negative result: there exists an EI that does not relax (Eq. (5.4), based on an example in [KR13]). Nevertheless, we show that every EI can be relaxed to its corresponding APXI plus an error term, which can be made arbitrarily small, at the cost of increasing the factor λ . This result implies that *every* EI is a consequence of its corresponding inequality APXI, plus an error term. In fact, the EI in Eq. (5.4) follows from an inequality by Matúš [Mat07], which is precisely the associated APXI plus an error term; our result shows that every EI can be proven in this style.

In Section 6 we consider approximate (and exact) implications that can be proved using *Shannon's inequalities* (monotonicity and submodularity, reviewed in Section 2.2). In general, Shannon's inequalities are sound but incomplete for proving exact and approximate implications that hold for all probability distributions [ZY97, ZY98], but they are complete for deriving inequalities that hold for all *polymatroids* (defined in Section 2.2) [Yeu08]. We also prove that every exact implication that holds for all polymatroids relaxes to an approximate implication with a finite upper bound $\lambda \leq (2^n)!$, and a lower bound $\lambda \geq 3$; the tightness of these bounds remain open.

In Section 7 we show that the Shannon inequalities are sound and complete for implication from *measure-based constraints* [SGG08] that arise in market-basket analysis. More generally, in Section 7 we restrict the class of *models* used to check an implication, to probability distributions with exactly 2 outcomes (tuples), each with probability 1/2; we justify this shortly. We prove that, under this restriction, the implication problem has a 1-relaxation. Restricting the models leads to a complete but unsound method for checking general implication, however this method is sound for saturated+conditional CIs (as we show in Section 4), and is also sound for deriving implications from Frequent Itemset constraints (as we show in Section 7).

In Section 8 we extend some of the results in this paper and provide alternative proofs using the *I-measure* theory that relates between Shannon's information measures and set-theory. Specifically, we extend the result of Section 4, and provide an alternative proof to the result of Section 7.

Two Consequences. While our paper is focused on relaxation, our results have two consequences for the exact implication problem. The first is a *2-tuple model property*: an exact implication, where the antecedents are saturated or conditional CIs, holds iff it holds on all uniform probability distributions with 2 tuples. A similar result is known for MVD+FDs [SDPF81]. Geiger and Pearl [GP93], building on an earlier result by Fagin [Fag82], prove that every set of CIs has an *Armstrong model*: a discrete probability distribution that satisfies only the CIs and their consequences, and no other CI. The Armstrong model is also called a *global witness*, and, in general, can be arbitrarily large. Our result concerns a *local witness*: for any EI, if it fails on some probability distribution, then it fails on a 2-tuple uniform distribution.

The second consequence concerns the equivalence between the implication problem of saturated+conditional CIs with that of MVD+FDs. It is easy to check that the former implies the latter (Section 2). Wong et al. [WBW00] prove the other direction (i.e., if an implication from MVDs+FDs holds then the same implication holds for saturated+conditional CIs), relying on the sound and complete axiomatization of MVDs [BFH77]. Our 2-tuple model property implies the other direction almost immediately, leading to a much simpler proof in Section 4.

This article extends the conference publication by the authors [KS20]. We have added in this article all the proofs and intermediate results that were excluded from the conference paper, as well as examples illustrating the ideas and methods introduced.

2. NOTATION AND PRELIMINARIES

We denote by $[n] = \{1, 2, \dots, n\}$. If $\Omega = \{X_1, \dots, X_n\}$ denotes a set of variables and $U, V \subseteq \Omega$, then we abbreviate the union $U \cup V$ by UV .

2.1. Integrity Constraints and Conditional Independence. A *relation instance* R over signature $\Omega = \{X_1, \dots, X_n\}$ is a finite set of tuples with attributes Ω . Let $X, Y, Z \subseteq \Omega$. We say that the instance R satisfies the *functional dependency* (FD) $X \rightarrow Y$, and write $R \models X \rightarrow Y$, if for all $t_1, t_2 \in R$, $t_1[X] = t_2[X]$ implies $t_1[Y] = t_2[Y]$. We say that R satisfies the *embedded multivalued dependency* (EMVD) $X \twoheadrightarrow Y \mid Z$, and write $R \models X \twoheadrightarrow Y \mid Z$, if for all $t_1, t_2 \in R$, $t_1[X] = t_2[X]$ implies $\exists t_3 \in R$ such that $t_1[XY] = t_3[XY]$ and $t_2[XZ] = t_3[XZ]$. One can check that $X \twoheadrightarrow Y \mid Z$ iff $X \rightarrow Y$. When $XYZ = \Omega$, then we call $X \twoheadrightarrow Y \mid Z$ a *multivalued dependency*, MVD; notice that X, Y, Z are not necessarily disjoint [BFH77].

A set of constraints Σ *implies* a constraint τ , in notation $\Sigma \Rightarrow \tau$, if for every instance R , if $R \models \Sigma$ then $R \models \tau$. The implication problem has been extensively studied in the literature; Beeri et al. [BFH77] gave a complete axiomatization of FDs and MVDs along with a polynomial time procedure for deciding implication. Herrman [Her06] showed that the implication problem for EMVDs is undecidable. In the following we refer to this type of implication as *Exact Implication*, abbreviated EI.

Recall that two discrete random variables X, Y are called *independent* if $p(X = x, Y = y) = p(X = x) \cdot p(Y = y)$ for all outcomes x, y . We say that X, Y are *conditionally independent* given another discrete random variable Z if $p(X = x, Y = y \mid Z = z) = p(X = x \mid Z = z) \cdot p(Y = y \mid Z = z)$ for all outcomes x, y , and z . Fix $\Omega = \{X_1, \dots, X_n\}$ a set of n jointly distributed discrete random variables with finite domains $\mathcal{D}_1, \dots, \mathcal{D}_n$, respectively; let $p : \mathcal{D}_1 \times \dots \times \mathcal{D}_n \rightarrow [0, 1]$ be the probability mass.

Notation 2.1. For $\alpha \subseteq [n]$, denote by X_α the joint random variable $(X_i : i \in \alpha)$ with domain $\mathcal{D}_\alpha \stackrel{\text{def}}{=} \prod_{i \in \alpha} \mathcal{D}_i$.

We write $p \models X_\beta \perp X_\gamma \mid X_\alpha$ when X_β, X_γ are conditionally independent given X_α . Applying our previous definition, this means that for any triple $(x_\alpha, x_\beta, x_\gamma)$ where $x_\alpha \in \mathcal{D}_\alpha$, $x_\beta \in \mathcal{D}_\beta$ and $x_\gamma \in \mathcal{D}_\gamma$, we have that $p(x_\alpha, x_\beta, x_\gamma) = p(x_\beta \mid x_\alpha) p(x_\gamma \mid x_\alpha)$. In the special case where $\beta = \gamma$, then $p \models X_\beta \perp X_\beta \mid X_\alpha$ has the following meaning. Let $x_\alpha \in \mathcal{D}_\alpha$ such that $p(x_\alpha) > 0$, and let $x_\beta \in \mathcal{D}_\beta$. Applying the definition of conditional independence directly gives us that $p(x_\beta, x_\beta \mid x_\alpha) = p(x_\beta \mid x_\alpha) p(x_\beta \mid x_\alpha)$. On the other hand, we know that $p(x_\beta, x_\beta \mid x_\alpha) \equiv p(x_\beta, \mid x_\alpha)$. Hence, we get that $p(x_\beta, \mid x_\alpha) = p(x_\beta \mid x_\alpha) p(x_\beta \mid x_\alpha)$. This holds only in one of the following conditions: (1) $p(x_\beta \mid x_\alpha) = 0$ or (2) $p(x_\beta \mid x_\alpha) = 1$. In other words, given that $X_\alpha = x_\alpha$ there is exactly one possible assignment $X_\beta = x_\beta$ such that $p(x_\beta \mid x_\alpha) > 0$. Hence, X_α functionally determines X_β , and we denote this by $p \models X_\alpha \rightarrow X_\beta$.

An assertion $Y \perp Z \mid X$ is called a *Conditional Independence* statement, or a CI; this includes $X \rightarrow Y$ as a special case. When $XYZ = \Omega$ we call it *saturated*. When $XY \subseteq \Omega$ and $Z = \emptyset$, we call it *marginal*. A set of CIs Σ *implies* a CI τ , in notation $\Sigma \Rightarrow \tau$, if every probability distribution that satisfies Σ also satisfies τ . This implication problem has also been extensively studied: Pearl and Paz [PP86] gave a sound but incomplete set of *graphoid axioms*, Studeny [Stu90] proved that no finite axiomatization exists, while Geiger and Pearl [GP93] gave a complete axiomatization for saturated, and marginal CIs.

Lee [Lee87] observed the following connection between database constraints and CIs. The *empirical distribution* of a relation R is the uniform distribution over its tuples, in other words, $\forall t \in R, p(t) = 1/|R|$. Then:

Lemma 2.2 [Lee87]. For all $X, Y, Z \subseteq \Omega$ such that $XYZ = \Omega$.

$$R \models X \rightarrow Y \Leftrightarrow p \models X \rightarrow Y \quad \text{and} \quad R \models X \twoheadrightarrow Y \mid Z \Leftrightarrow p \models (Y \perp Z \mid X) \quad (2.1)$$

As can be seen in the Example of Table 2, the lemma no longer holds for EMVDs, and for that reason we no longer consider EMVDs in this paper. The lemma immediately implies that if Σ is a set of saturated+conditional CIs and the implication $\Sigma \Rightarrow \tau$ holds for all probability distributions, then the corresponding implication holds in databases, where the CIs are interpreted as MVDs or FDs respectively. Wong [WBW00] gave a non-trivial proof for the other direction; we will give a much shorter proof in Corollary 4.3.

X_1	X_2	X_3
a	c	e
b	c	e
a	d	e
b	d	e
b	d	f

TABLE 2. The relation $R[X_1, X_2, X_3]$ satisfies the EMVD $\emptyset \twoheadrightarrow X_1 \mid X_2$, yet for the empirical distribution, $I_h(X_1; X_2) \neq 0$ because X_1, X_2 are dependent: $p(X_1 = a) = 2/5 \neq p(X_1 = a \mid X_2 = c) = 1/2$.

2.2. Background on Information Theory. We adopt required notation from the literature on information theory [Yeu08, Cha11]. For $n > 0$, we identify vectors in \mathbb{R}^{2^n} with functions $2^{[n]} \rightarrow \mathbb{R}$.

Polymatroids. A function¹ $h \in (\mathbb{R}_+)^{2^n}$ is called a *polymatroid* if $h(\emptyset) = 0$ and it satisfies the following inequalities, called *Shannon inequalities*:

$$\text{Monotonicity: } h(A) \leq h(B) \text{ for } A \subseteq B \subseteq [n] \quad (2.2)$$

$$\text{Submodularity: } h(A \cup B) + h(A \cap B) \leq h(A) + h(B) \text{ for all } A, B \subseteq [n] \quad (2.3)$$

The set of polymatroids is denoted $\Gamma_n \subseteq (\mathbb{R}_+)^{2^n}$, and forms a polyhedral cone (reviewed in Section 5). For any polymatroid h and subsets $A, B, C \subseteq [n]$, we define²

$$h(B|A) \stackrel{\text{def}}{=} h(AB) - h(A) \quad (2.4)$$

$$I_h(B; C|A) \stackrel{\text{def}}{=} h(AB) + h(AC) - h(ABC) - h(A) \quad (2.5)$$

Then, for all $h \in \Gamma_n$, we have $I_h(B; C|A) \geq 0$ and $h(B|A) \geq 0$. The *chain rule* for entropies and mutual information, respectively, are the identities:

$$h(BC|D) = h(B|D) + h(C|BD) \quad (2.6)$$

$$I_h(B; CD|A) = I_h(B; C|A) + I_h(B; D|AC) \quad (2.7)$$

when the rule being applied is clear from the context, we will only say that we apply the chain rule.

We call $I_h(B; C|A)$ *saturated* if $ABC = [n]$, and *elemental* if $|B| = |C| = 1$. The conditional entropy $h(B|A)$ is a special case of the mutual information I_h , because $h(B|A) = I_h(B; B|A)$.

¹Most authors consider rather the space $\mathbb{R}^{2^n - 1}$, by dropping $h(\emptyset)$ because it is always 0.

²Recall that AB denotes $A \cup B$.

Entropic Functions. If X is a random variable with a finite domain \mathcal{D} and probability mass p , then $H(X)$ denotes its entropy

$$H(X) \stackrel{\text{def}}{=} \sum_{x \in \mathcal{D}} p(x) \log_2 \frac{1}{p(x)} \quad (2.8)$$

For a set of jointly distributed random variables $\Omega = \{X_1, \dots, X_n\}$ we define the function $h : 2^{[n]} \rightarrow \mathbb{R}_+$ as $h(\alpha) \stackrel{\text{def}}{=} H(X_\alpha)$ (see Notation 2.1); h is called an *entropic function*, or, with some abuse, an *entropy*. The set of entropic functions is denoted Γ_n^* . The quantities $h(B|A)$ and $I_h(B; C|A)$ are called the *conditional entropy* and *conditional mutual information* respectively. The conditional independence $p \models B \perp C \mid A$ holds iff $I_h(B; C|A) = 0$, and similarly $p \models A \rightarrow B$ iff $h(B|A) = 0$, thus, entropy provides us with an alternative characterization of CIs.

Notation 2.3. We summarize the various ways a conditional independence statement (CI) and related notions are represented in the paper, and the relationships between them. In what follows, we let p be an n -variable joint probability distribution over variable set Ω , and we let H_p denote the entropy function corresponding to p (see (2.8)). Let $X, Y, Z \subseteq \Omega$ where X and Y are non-empty.

- (1) The notation $X \perp Y \mid Z$ implies that for any set of values x, y, z in the domains of X, Y and Z respectively, it holds that $p(X = x, Y = y, Z = z) = p(X = x \mid Z = z)p(Y = y \mid Z = z)$.
- (2) The statement $X \perp Y \mid Z$ is equivalent to saying that $I_{H_p}(X; Y \mid Z) = 0$ where I_{H_p} is defined in (2.5). Note, however, that the notation $I_h(X; Y \mid Z) = 0$ is more general since we only assume that h is a polymatroid (i.e., but not necessarily the entropy function associated with a probability distribution).
- (3) For a polymatroid h (which may or may not be an entropic function), we denote by the triple $\sigma \stackrel{\text{def}}{=} (X; Y \mid Z)$ a CI statement that may hold either exactly (i.e., $I_h(X; Y \mid Z) = 0$), or approximately (i.e., $I_h(X; Y \mid Z) \leq \varepsilon$ for some threshold $\varepsilon > 0$). We then abbreviate: $h(\sigma) \stackrel{\text{def}}{=} I_h(X; Y \mid Z)$.
- (4) When $XYZ = \Omega$, and p is a uniform distribution (i.e., for every $r \in \mathcal{D}_\Omega$ either $p(r) = \frac{1}{M}$ for some $M > 0$ or $p(r) = 0$), then by Lemma 2.2, an MVD $Z \rightarrow X \mid Y$ holds in p iff $I_{H_p}(X; Y \mid Z) = 0$.

Next, we prove two simple, technical lemmas concerning CIs $(X; Y \mid Z)$ where the intersection between the variable-sets (e.g., $X \cap Y$, $X \cap Z$, or $Y \cap Z$) may be non-empty. These lemmas will be used later on.

Lemma 2.4. Let $h \in \Gamma_n$ be an n -variable polymatroid, and let X, Y and Z denote subsets of variables. Then for any CI $(X; Y \mid Z)$ where $Y \subseteq Z$, it holds that $I_h(X; Y \mid Z) = 0$.

Proof. Since $Y \subseteq Z$, we denote $Z = Z'Y$ where $Z' = Z \setminus Y$. Therefore, by (2.5):

$$I_h(X; Y \mid Z'Y) = h(XYZ') + h(Z'Y) - h(Z'Y) - h(XYZ') = 0 \quad \square$$

Lemma 2.5. Let $h \in \Gamma_n$ be an n -variable polymatroid, and let X, Y and Z denote subsets of variables. Then $I_h(X; Y \mid Z) \equiv h(B_{XY} \mid Z') + I_h(X'; Y' \mid Z')$ where $X' \stackrel{\text{def}}{=} X \setminus YZ$, $Y' \stackrel{\text{def}}{=} Y \setminus XZ$, $Z' \stackrel{\text{def}}{=} Z \setminus XY$ are pairwise disjoint, and $B_{XY} \stackrel{\text{def}}{=} X \cap Y \setminus Z$.

Proof. We define $A_{XY} \stackrel{\text{def}}{=} X \cap Y$, $A_{XZ} \stackrel{\text{def}}{=} X \cap Z$, $A_{YZ} \stackrel{\text{def}}{=} Y \cap Z$, and $A_{XYZ} \stackrel{\text{def}}{=} X \cap Y \cap Z$. We let $X' = X \setminus A_{XY}A_{XZ}$, $Y' = Y \setminus A_{XY}A_{YZ}$, and $Z' = Z \setminus A_{XZ}A_{YZ}$. Hence, we can

<table style="border-collapse: collapse; text-align: left;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 2px 10px;">X_1</th> <th style="border-right: 1px solid black; padding: 2px 10px;">X_2</th> <th style="border-right: 1px solid black; padding: 2px 10px;">U_1</th> <th style="border-right: 1px solid black; padding: 2px 10px;">U_2</th> <th style="padding: 2px 10px;">Pr</th> </tr> </thead> <tbody> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">1/2</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">1/2</td> </tr> </tbody> </table>	X_1	X_2	U_1	U_2	Pr	0	0	0	0	1/2	1	1	0	0	1/2	<table style="border-collapse: collapse; text-align: left;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 2px 10px;">X</th> <th style="border-right: 1px solid black; padding: 2px 10px;">Y</th> <th style="border-right: 1px solid black; padding: 2px 10px;">Z</th> <th style="padding: 2px 10px;">Pr</th> </tr> </thead> <tbody> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">1/4</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="padding: 2px 10px;">1/4</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="padding: 2px 10px;">1/4</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">1/4</td> </tr> </tbody> </table>	X	Y	Z	Pr	0	0	0	1/4	0	1	1	1/4	1	0	1	1/4	1	1	0	1/4	<table style="border-collapse: collapse; text-align: left;"> <thead> <tr> <th style="border-right: 1px solid black; padding: 2px 10px;">A</th> <th style="border-right: 1px solid black; padding: 2px 10px;">B</th> <th style="border-right: 1px solid black; padding: 2px 10px;">C</th> <th style="border-right: 1px solid black; padding: 2px 10px;">D</th> <th style="padding: 2px 10px;">Pr</th> </tr> </thead> <tbody> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">1/2 - ε</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="padding: 2px 10px;">1/2 - ε</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">ε</td> </tr> <tr> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">1</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="border-right: 1px solid black; padding: 2px 10px;">0</td> <td style="padding: 2px 10px;">ε</td> </tr> </tbody> </table>	A	B	C	D	Pr	0	0	0	0	1/2 - ε	0	1	0	1	1/2 - ε	1	0	1	0	ε	1	1	0	0	ε
X_1	X_2	U_1	U_2	Pr																																																										
0	0	0	0	1/2																																																										
1	1	0	0	1/2																																																										
X	Y	Z	Pr																																																											
0	0	0	1/4																																																											
0	1	1	1/4																																																											
1	0	1	1/4																																																											
1	1	0	1/4																																																											
A	B	C	D	Pr																																																										
0	0	0	0	1/2 - ε																																																										
0	1	0	1	1/2 - ε																																																										
1	0	1	0	ε																																																										
1	1	0	0	ε																																																										
(a)	(b)	(c)																																																												

FIGURE 1. Two uniform probability distributions (a),(b); a distribution from [KR13] (c).

write $I_h(X; Y|Z)$ as:

$$\begin{aligned}
& I_h(A_{XY}A_{XZ}X'; A_{XY}A_{YZ}Y'|A_{XZ}A_{YZ}Z') \\
&= I(A_{XZ}; A_{XY}A_{YZ}Y'|A_{XZ}A_{YZ}Z') + I(A_{XY}X'; A_{XY}A_{YZ}Y'|A_{XZ}A_{YZ}Z') && \text{Chain Rule 2.7} \\
&= I_h(A_{XY}X'; A_{XY}A_{YZ}Y'|A_{XZ}A_{YZ}Z') && \text{Lemma 2.4} \\
&= I_h(A_{XY}X'; A_{YZ}|A_{XZ}A_{YZ}Z') + I_h(A_{XY}X'; A_{XY}Y'|A_{XZ}A_{YZ}Z') && \text{Chain Rule 2.7} \\
&= I_h(A_{XY}X'; A_{XY}Y'|A_{XZ}A_{YZ}Z') && \text{Lemma 2.4} \\
&= I_h(A_{XY}X'; A_{XY}|Z) + I_h(A_{XY}X'; Y'|ZA_{XY}) && \text{Chain Rule 2.7} \\
&= I_h(A_{XY}; A_{XY}|Z) + I_h(X'; A_{XY}|A_{XY}Z) + I_h(X'; Y'|ZA_{XY}) + I_h(A_{XY}; Y'|X'ZA_{XY}) && \text{Chain Rule 2.7} \\
&= I_h(A_{XY}; A_{XY}|Z) + I_h(X'; Y'|ZA_{XY}) && \text{Lemma 2.4} \\
&= h(A_{XY}|Z) + I_h(X'; Y'|ZA_{XY}) \\
&= h(A_{XYZ}B_{XY}|Z) + I_h(X'; Y'|ZA_{XY}) && B_{XY} \stackrel{\text{def}}{=} A_{XY} \setminus Z \\
&= h(A_{XYZ}|Z) + h(B_{XY}|Z) + I_h(X'; Y'|ZA_{XY}) && \text{Chain Rule 2.7} \\
&= h(B_{XY}|Z) + I_h(X'; Y'|ZA_{XY}) && h(A_{XYZ}|Z) = 0
\end{aligned}$$

By definition, X' , Y' , Z and B_{XY} are pairwise disjoint, thus proving the claim. \square

2-Tuple Relations and Step Functions. 2-tuple relations play a key role for the implication problem of MVDs+FDs: if an implication fails, then there exists a witness consisting of only two tuples [SDPF81]. We define a *step function* as the entropy of the empirical distribution of a 2-tuple relation; $R = \{t_1, t_2\}$, $t_1 \neq t_2$, and $p(t_1) = p(t_2) = 1/2$. We denote the step function by h_U , where $U \subsetneq \Omega$ is the set of attributes where t_1, t_2 agree. One can check:

$$h_U(W) = \begin{cases} 0 & \text{if } W \subseteq U \\ 1 & \text{otherwise} \end{cases} \quad (2.9)$$

To see why, observe that, by definition, for every set of attributes $W \subseteq U$ we have that t_1 and t_2 agree on all attributes in W (i.e., $t_1[W] = t_2[W]$). Hence, $h_U(W) = 1 \log 1 = 0$. On the other hand, if $W \not\subseteq U$, then $t_1[W] \neq t_2[W]$, and computing the entropy gives us that $h_U(W) = \frac{1}{2} \log 2 + \frac{1}{2} \log 2 = 1$. If we set $U = \Omega$ in (2.9) then $h_\Omega \equiv 0$. That is, $h_\Omega(W) = 0$ for all $W \subseteq \Omega$, and hence, h_Ω is equivalent to the function that always returns 0. Unless otherwise stated, in this paper we do not consider h_Ω to be a step function. Thus, there are $2^n - 1$ step functions and their set is denoted \mathcal{S}_n . We will use the following fact extensively in this paper: $I_{h_U}(Y; Z|X) = 1$ if $X \subseteq U$ and $Y, Z \not\subseteq U$, and $I_{h_U}(Y; Z|X) = 0$ otherwise.

We now present two simple polymatroids that correspond to the entropic functions of the distributions in Figure 1.

Example 2.6. Consider the relational instance in Fig. 1 (a). Its entropy is the step function $h_{U_1U_2}(W)$, which is 0 for $W \subseteq U_1U_2$ and 1 otherwise. $R \models X_1 \rightarrow X_2$ because $h(X_2|X_1) = h(X_1X_2) - h(X_1) = 1 - 1 = 0$, and $R \not\models U_1 \rightarrow X_1$ because $h(X_1|U_1) = h(X_1U_1) - h(U_1) = 1 - 0 \neq 0$.

Example 2.7. The relational instance $R = \{(x, y, z) \mid x + y + z \pmod 2 = 0\}$ in Fig. 1 (b) is called the *parity function*. Its entropy is $h(X) = h(Y) = h(Z) = 1$, $h(XY) = h(XZ) = h(YZ) = h(XYZ) = 2$. We have that $R \models Y \perp Z$ because $I_h(Y; Z) = h(Y) + h(Z) - h(YZ) = 1 + 1 - 2 = 0$, but $R \not\models Y \perp Z|X$ because $I_h(Y; Z|X) = 1$. To see this, observe that:

$$\begin{aligned} I_h(Y; Z|X) &= h(XY) + h(XZ) - h(X) - h(XYZ) \\ &= 2 + 2 - 1 - 2 \\ &= 1 \end{aligned}$$

2.3. Discussion. This paper studies exact and approximate implications, expressed as equalities or inequalities of entropic functions h . For example, the augmentation axiom for MVDs [BFH77] $A \twoheadrightarrow B|CD \Rightarrow AC \twoheadrightarrow B|D$ is expressed as $I_h(B; CD|A) = 0 \Rightarrow I_h(B; D|AC) = 0$, which holds by the chain rule (2.7). Thus, our golden standard is to prove that (in)equalities hold for all entropic functions $h \in \Gamma_n^*$.

Fix a set $A \subseteq \mathbb{R}^N$ in an N -dimensional euclidean space with distance metric d . The *topological closure* of A is defined as:

$$\mathbf{cl}(A) \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathbb{R}^N \mid \forall \varepsilon > 0, \exists \mathbf{s} \in A : d(\mathbf{s}, \mathbf{x}) < \varepsilon\}$$

Equivalently, $\mathbf{cl}(A)$ consists of all limits of convergent sequences in A . The set A is called *topologically closed* if $A = \mathbf{cl}(A)$. The topological closure enjoys the following basic property: if $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is any continuous function, and $f(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in A$, then $f(\mathbf{x}) \geq 0$ for $\mathbf{x} \in \mathbf{cl}(A)$. Yeung [Yeu08] has proven that, when $n \geq 3$, then Γ_n^* is not topologically closed, in other words, $\Gamma_n^* \subsetneq \mathbf{cl}(\Gamma_n^*)$. The elements of the set $\mathbf{cl}(\Gamma_n^*)$ are called *almost entropic functions*. Equivalently, a function g is *almost entropic* if it is the limit of a sequence of entropic functions. It follows immediately from our discussion that, if an *inequality* holds for all entropic functions $h \in \Gamma_n^*$, then, by continuity, it also holds for all almost entropic functions $h \in \mathbf{cl}(\Gamma_n^*)$. However, this observation does not extend to *implications* of (in)equalities; Kaced and Romashchenko [KR13] gave an example of an exact implication that holds only for entropic functions but fails for almost entropic functions. Thus, when discussing an EI, it matters whether we assume that it holds for Γ_n^* or for $\mathbf{cl}(\Gamma_n^*)$. The only result in this paper where this distinction matters are the two main theorems in Section 5: the negative result Theorem 5.1 holds for both Γ_n^* and for $\mathbf{cl}(\Gamma_n^*)$, while the positive result Theorem 5.2 holds only for $\mathbf{cl}(\Gamma_n^*)$. The results in Section 4 apply to *any* set of polymatroids K that contains all step functions, i.e. $\mathcal{S}_n \subseteq K \subseteq \Gamma_n$, thus they apply to both Γ_n^* and $\mathbf{cl}(\Gamma_n^*)$, while those in Section 6 and Section 7 are stated only for Γ_n and only for (the conic closure of) \mathcal{S}_n respectively.

3. DEFINITION OF THE RELAXATION PROBLEM

We now formally define the relaxation problem. We fix a set of variables $\Omega = \{X_1, \dots, X_n\}$, and consider formulas of the form $\sigma = (Y; Z|X)$, where $X, Y, Z \subseteq \Omega$, which we call a *conditional independence*, CI; when $Y = Z$ then we write it as $X \rightarrow Y$ and call it a *conditional*. An *implication* is a formula $\Sigma \Rightarrow \tau$, where Σ is a set of CIs called *antecedents* and τ is a CI called *consequent*. For a CI $\sigma = (B; C|A)$, we define $h(\sigma) \stackrel{\text{def}}{=} I_h(B; C|A)$ (see (2.5)), for a set of CIs Σ , we define $h(\Sigma) \stackrel{\text{def}}{=} \sum_{\sigma \in \Sigma} h(\sigma)$. Fix a set K s.t. $\mathcal{S}_n \subseteq K \subseteq \Gamma_n$. We recall that \mathcal{S}_n is the set of step functions, and that Γ_n is the set of polymatroids.

Definition 3.1. The *exact implication* (EI) $\Sigma \Rightarrow \tau$ holds in K , denoted $K \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ if, for all $h \in K$, $h(\Sigma) = 0$ implies $h(\tau) = 0$. The λ -*approximate implication* (λ -APXI) holds in K , in notation $K \models \lambda \cdot h(\Sigma) \geq h(\tau)$, if $\forall h \in K, \lambda \cdot h(\Sigma) \geq h(\tau)$. The *approximate implication* holds, in notation $K \models_{\text{APXI}} (\Sigma \Rightarrow \tau)$, if there exist a $\lambda \geq 0$ such that the λ -APXI holds.

We will sometimes consider an equivalent definition for APXI, as $\sum_{\sigma \in \Sigma} \lambda_{\sigma} h(\sigma) \geq h(\tau)$, where $\lambda_{\sigma} \geq 0$ are coefficients, one for each $\sigma \in \Sigma$; these two definitions are equivalent, by taking $\lambda = \max_{\sigma} \lambda_{\sigma}$. Notice that both EI and APXI are preserved under subsets of K in the sense that $K_1 \subseteq K_2$ and $K_2 \models_x (\Sigma \Rightarrow \tau)$ implies $K_1 \models_x (\Sigma \Rightarrow \tau)$, for $x \in \{\text{EI}, \text{APXI}\}$.

APXI always implies EI. Indeed, $h(\tau) \leq \lambda \cdot h(\Sigma)$ and $h(\Sigma) = 0$ imply $h(\tau) \leq 0$, which further implies $h(\tau) = 0$, because $h(\tau) \geq 0$ for every CI τ , and every polymatroid h . In this paper we study the reverse.

Definition 3.2. Let \mathcal{I} be a syntactically-defined class of implication statements $(\Sigma \Rightarrow \tau)$, and let $K \subseteq \Gamma_n$. We say that \mathcal{I} *admits a relaxation* in K if every implication statement $(\Sigma \Rightarrow \tau)$ in \mathcal{I} that holds exactly also holds approximately: $K \models_{\text{EI}} \Sigma \Rightarrow \tau$ implies $K \models_{\text{APXI}} \Sigma \Rightarrow \tau$. We say that \mathcal{I} admits a λ -*relaxation* in K , if every EI admits a λ -APXI.

If $K_1 \subseteq K_2$ and \mathcal{I} admits a λ -relaxation in K_1 then, in general, it does not necessarily admit a λ -relaxation in K_2 , nor vice versa. However, the following more limited fact holds:

Fact 3.3. If \mathcal{I} admits a λ -relaxation in K , then it also admits a λ -relaxation in $\mathbf{cl}(K)$.

Proof. Let $\Sigma \Rightarrow \tau$ be a CI statement in the syntactic class \mathcal{I} , and suppose $\mathbf{cl}(K) \models_{\text{EI}} (\Sigma \Rightarrow \tau)$. Then $K \models_{\text{EI}} (\Sigma \Rightarrow \tau)$. Since \mathcal{I} admits a λ -relaxation in K , we have $K \models_{\text{APXI}} (\Sigma \Rightarrow \tau)$, or, equivalently, $\forall h \in K : \lambda \cdot h(\Sigma) \geq h(\tau)$. By continuity, $\forall h \in \mathbf{cl}(K) : \lambda h(\Sigma) \geq h(\tau)$, proving the claim. \square

Example 3.4. Let $\Sigma = \{(A; B|\emptyset), (A; C|B)\}$ and $\tau = (A; C|\emptyset)$. By the chain rule (2.7), we have that

$$\begin{aligned} I_h(A; BC) &= I_h(A; B|\emptyset) + I_h(A; C|B) && \text{and} \\ I_h(A; BC) &= I_h(A; C|\emptyset) + I_h(A; B|C) \geq I_h(A; C|\emptyset) \end{aligned}$$

Hence, $I_h(A; C|\emptyset) \leq I_h(A; B|\emptyset) + I_h(A; C|B)$. Therefore, the exact implication $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ admits a 1-APXI.

4. RELAXATION FOR FDS AND MVDs: ALWAYS POSSIBLE

In this section we consider the implication problem where the antecedents are either saturated CIs, or conditionals. This is a case of special interest in databases, because the constraints correspond to MVDs, or FDs. Recall that a CI $(B; C|A)$ is *saturated* if $ABC = \Omega$ (i.e., the set of all attributes). Our main result in this section is:

Theorem 4.1. *Assume that each formula in Σ is either saturated or a conditional (e.g., $Z \rightarrow X$), and let τ be an arbitrary CI. Assume $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$. Then:*

- (1) $\Gamma_n \models \frac{n^2}{4}h(\Sigma) \geq h(\tau)$.
- (2) *If τ is a conditional, then $\Gamma_n \models h(\Sigma) \geq h(\tau)$.*

Before we prove the theorem, we list two important consequences.

Corollary 4.2. *Let Σ consist of saturated CIs and/or conditionals, and let τ be any CI. Then $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ implies $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$.*

Proof. If $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ then $\forall h \in \Gamma_n, h(\tau) \leq \frac{n^2}{4}h(\Sigma)$, thus $h(\Sigma) = 0$ implies $h(\tau) = 0$. \square

The corollary implies that if Σ, τ are restricted to saturated CIs and/or conditionals, then the Exact Implication problem is the same for \mathcal{S}_n as it is for any other set K where $\mathcal{S}_n \subseteq K \subseteq \Gamma_n$.

The corollary has an immediate application to the inference problem in graphical models [GP93]. There, the problem is to check if every probability distribution that satisfies all CIs in Σ also satisfies the CI τ ; we have seen that this is equivalent to $\Gamma_n^* \models_{\text{EI}} \Sigma \Rightarrow \tau$. The corollary states that it is enough that this implication holds on all of the uniform 2-tuple distributions, i.e., $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$, because this implies the (even stronger!) statement $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$. Decidability (i.e., for exact implication) was already known: Geiger and Pearl [GP93] proved that the set of graphoid axioms is sound and complete for the case when both Σ and τ are saturated. Specifically, the equivalence between saturated CIs and MVDs [GP93] enables the application of the polynomial implication algorithm devised for MVDs [Bee80] which, in this setting where both Σ and τ are saturated, has a runtime complexity of $O(|\Sigma|n^2)$. Gyssens et al. [GNG14] improve this result by dropping any restrictions on τ .

The second consequence is the following:

Corollary 4.3. *Let Σ, τ consist of saturated CIs and/or conditionals. Then the following two statements are equivalent:*

- (1) *The implication $\Sigma \Rightarrow \tau$ holds, where we interpret Σ, τ as MVDs and/or FDs.*
- (2) $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$.

Proof. We have shown right after Lemma 2.2 that (2) implies (1). For the opposite direction, by Th. 4.1, we need only check $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$, which holds because on every uniform probability distribution a saturated CI holds iff the corresponding MVD holds, and similarly for conditionals and FDs. Since the 2-tuple relation satisfies the implication for MVDs+FDs, it also satisfies the implication for CIs, proving the claim. \square

Wong et al. [WBW00] proved that the implication for MVDs is equivalent to that of the corresponding saturated CIs (called there BMVD); they did not consider FDs. For the proof in the hard direction, they use the sound and complete axiomatization of MVDs in [BFH77]. In contrast, our proof is independent of any axiomatic system, and is also much shorter.

Finally, we notice that the corollary also implies that, in order to check an implication between MVDs and/or FDs, it suffices to check it on all 2-tuple databases: indeed, this is equivalent to checking $\mathcal{S}_n \models_{EI} \Sigma \Rightarrow \tau$, because this implies Item (2), which in turn implies item (1). This rather surprising fact was first proven in [SDPF81].

The proof of Theorem 4.1 follows from a series of lemmas, and a Theorem that is of independent interest, that we prove next. Before proceeding, we note that we can assume w.l.o.g. that Σ consists only of saturated CIs. Indeed, if Σ contains a non-saturated term, then by assumption it is a conditional, $X \rightarrow Y$, and we will replace it with two saturated terms: $(Y; Z|X)$ and $XZ \rightarrow Y$, where $Z = \Omega \setminus XY$. Denoting Σ' the new set of formulas, we have $h(\Sigma) = h(\Sigma')$, because $h(Y|X) = I_h(Y; Z|X) + h(Y|XZ)$. Thus, we will assume w.l.o.g. that all formulas in Σ are saturated.

We say that a CI $(X; Y|Z)$ is *elemental* if $|X| = |Y| = 1$.

Lemma 4.4. *Every CI $\tau = (X; Y|Z)$ can be written as a sum of $m + n_X \cdot n_Y$ elemental terms where $m = |X \cap Y \setminus Z|$, $n_X = |X \setminus YZ|$, and $n_Y = |Y \setminus XZ|$. Furthermore, it holds that $m + n_X \cdot n_Y \leq \frac{n^2}{4}$.*

Proof. By Lemma 2.5 we have that $h(\tau) \equiv h(B_{XY}|Z) + I_h(X'; Y'|Z)$ where $B_{XY} \stackrel{\text{def}}{=} X \cap Y \setminus Z$, $X' = X \setminus YZ$, and $Y' = Y \setminus XZ$. By the chain rule, $h(B_{XY}|Z)$ and $I_h(X'; Y'|Z)$ can be written as the sum of $m = |B_{XY}|$, and $|X'| \cdot |Y'| = n_X n_Y$ elemental terms respectively. Since B_{XY} , X' , and Y' are pairwise disjoint, then $m + n_X + n_Y \leq n$. Therefore:

$$\begin{aligned} m + n_X \cdot n_Y &\leq m + \frac{(n - m)^2}{4} \\ &< m + \frac{n^2 - 2m^2 + m^2}{4} && m < n \\ &= \frac{1}{4}(n^2 + m(4 - m)) \\ &\leq \frac{n^2}{4} + 1 \end{aligned}$$

Where the last transition follows by observing that $\frac{1}{4}m(4 - m) \leq 1$ for all integral m (e.g., for $m \in \{1, 3\}$ then $\frac{1}{4}m(4 - m) = \frac{3}{4}$, and for $m = 2$, $\frac{1}{4}m(4 - m) = 1$). Since the number of elemental terms is integral then $m + n_X \cdot n_Y \leq \frac{n^2}{4}$ as required. \square

Theorem 4.1 follows from the next result, which is also of independent interest. We say that σ *covers* τ if all variables in τ are contained in σ ; for example $\sigma = (abc; d|e)$ covers $\tau = (cd; be)$. Then:

Theorem 4.5. *Let τ be an elemental CI, and suppose each formula in Σ covers τ . Then $\mathcal{S}_n \models_{EI} (\Sigma \Rightarrow \tau)$ implies $\Gamma_n \models h(\tau) \leq h(\Sigma)$.*

Notice that Theorem 4.5 immediately implies Item (1) of Theorem 4.1, because by Lemma 4.4 every $\tau = (Y; Z|X)$ can be written as a sum of at most $n^2/4$ elemental terms. In what follows, we prove Theorem 4.5, then use it to prove item (2) of Theorem 4.1.

Finally, we consider whether (1) of Theorem 4.1 can be strengthened to a 1-relaxation; we give in Th. 4.11 below a sufficient condition, whose proof uses the notion of I-measure (Section 8), and leave open the question whether 1-relaxation holds in general for implications where the antecedents are saturated CIs and conditionals.

Definition 4.6. We say that two CIs $(X; Y|Z)$ and $(A; B|C)$ are *disjoint* if at least one of the following four conditions holds: (1) $X \subseteq C$, (2) $Y \subseteq C$, (3) $A \subseteq Z$, or (4) $B \subseteq Z$.

If $\tau = (X; Y|Z)$ and $\sigma = (A; B|C)$ are disjoint, then for any step function h_W , it cannot be the case that both $h_W(\tau) \neq 0$ and $h_W(\sigma) \neq 0$. Indeed, if such a W exists, then $Z, C \subseteq W$ and, assuming (1) $X \subseteq C$ (the other three cases are similar), we have $ZX \subseteq W$ thus $h_W(\tau) = 0$.

Lemma 4.7. *Suppose $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$, where $\tau = (X; Y|Z)$. Let $\sigma \in \Sigma$ such that τ, σ are disjoint (Def. 4.6). Then: $\mathcal{S}_n \models_{\text{EI}} (\Sigma \setminus \{\sigma\}) \Rightarrow \tau$.*

Proof. Let $\Sigma' \stackrel{\text{def}}{=} \Sigma \setminus \{\sigma\}$. Assume by contradiction that there exists a step function h_W such that $h_W(\Sigma') = 0$ and $h_W(\tau) = 1$. Since σ, τ are disjoint, $h_W(\sigma) = 0$. Then $h_W(\Sigma) = 0$, contradicting the assumption $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$. \square

4.1. Proof of Theorem 4.5. The following holds by the chain rule, and will be used later on.

Lemma 4.8. *Let $\sigma = (A; B|C)$ and $\tau = (X; Y|Z)$ be CIs such that $X \subseteq A$, $Y \subseteq B$, $C \subseteq Z$ and $Z \subseteq ABC$. Then, $\Gamma_n \models h(\tau) \leq h(\sigma)$.*

Proof. Since $Z \subseteq ABC$, we denote by $Z_A = A \cap Z$, $Z_B = B \cap Z$, and $Z_C = C \cap Z$. Also, we define $A' = A \setminus (Z_A \cup X)$ and $B' = B \setminus (Z_B \cup Y)$. So, we have that: $I(A; B|C) = I(Z_A A' X; Z_B B' Y|C)$. By the chain rule, we have that:

$$\begin{aligned} I(Z_A A' X; Z_B B' Y|C) &= I(Z_A; Z_B|C) + I(A' X; Z_B|C Z_A) + I(Z_A; B' Y|Z_B C) \\ &\quad + I(X; Y|C Z_A Z_B) + I(X; B'|C Z_A Z_B Y) + I(A'; B' Y|C Z_A Z_B X) \end{aligned}$$

Noting that $Z = C Z_A Z_B$, we get that $I(X; Y|Z) \leq I(A; B|C)$ as required. \square

We now prove theorem 4.5. We use lower case for single variables, thus $\tau = (x; y|Z)$ because it is elemental. We may assume w.l.o.g. that neither x nor y are in Z : $x, y \notin Z$ (otherwise $I_h(x; y|Z) = 0$ and the lemma holds trivially). The *deficit* of an elemental CI $\tau = (x; y|Z)$ is the quantity $|\Omega - Z|$. We prove by induction on the deficit of τ that $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ implies $\Gamma_n \models h(\tau) \leq h(\Sigma)$.

Assume $\mathcal{S}_n \models_{\text{EI}} (\Sigma \Rightarrow \tau)$, and consider the step function h_Z at Z . Since $h_Z(\tau) = 1$, there exists $\sigma \in \Sigma$, written $\sigma = (A; B|C)$, such that $h_Z(\sigma) = 1$; this means that $C \subseteq Z$, and $A, B \not\subseteq Z$. In particular $x, y \notin C$, therefore $x, y \in AB$, because σ covers τ . If $x \in A$ and $y \in B$ (or vice versa), then $\Gamma_n \models h(\tau) \leq h(\sigma)$ by Lemma 4.8, proving the theorem. Therefore, we assume w.l.o.g. that $x, y \in A$ and none is in B . Furthermore, since $B \not\subseteq Z$, there exists $u \in B - Z$.

Base case: τ is saturated. Then $u \notin xyZ$, contradicting the assumption that τ is saturated; in other words, in the base case, it is the case that $x \in A$ and $y \in B$.

Step: Let $Z_A = Z \cap A$, and $Z_B = Z \cap B$. Since $C \subseteq Z$, and $\sigma = (A; B|C)$ covers τ , then $Z = Z_A Z_B C$. We also write $A = xyA'Z_A$ (since $x, y \in A$) and $B = uB'Z_B$. So, we have that $\sigma = (A; B|C) = (xyA'Z_A; uB'Z_B|C)$, and we use the chain rule to define σ_1, σ_2 :

$$h(\sigma) = I_h(xyA'Z_A; uB'Z_B|C) = I_h(\underbrace{xyA'Z_A; uZ_B|C}_{\stackrel{\text{def}}{=} \sigma_1}) + I_h(\underbrace{xyA'Z_A; B'|uCZ_B}_{\stackrel{\text{def}}{=} \sigma_2})$$

We also partition Σ s.t. $h(\Sigma) = h(\sigma_1) + h(\Sigma_2)$, where $\Sigma_2 \stackrel{\text{def}}{=} (\Sigma \setminus \{\sigma\}) \cup \{\sigma_2\}$. Since σ_2 is a saturated CI (i.e., contains all variables in σ , which is saturated), then Σ_2 is saturated as well (i.e., contains only saturated CIs).

Next, define $\tau' \stackrel{\text{def}}{=} (x; uy|Z)$ and use the chain rule to define τ_1, τ_2 :

$$h(\underbrace{x; y|Z}_{\tau}) \leq I_h(\underbrace{x; uy|Z}_{\tau'}) = I_h(\underbrace{x; u|Z}_{\stackrel{\text{def}}{=} \tau_1}) + I_h(\underbrace{x; y|uZ}_{\stackrel{\text{def}}{=} \tau_2}) \quad (4.1)$$

By Lemma 4.8, $\Gamma_n \models h(\sigma_1) \geq h(\tau_1)$. We will prove: $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$. This implies the theorem, because Σ_2 is saturated, so by the induction hypothesis, $\Gamma_n \models h(\Sigma_2) \geq h(\tau_2)$ (since the deficit of τ_2 is one less than that of τ), and the theorem follows from $h(\Sigma) = h(\sigma_1) + h(\Sigma_2) \geq h(\tau_1) + h(\tau_2) = h(\tau') \geq h(\tau)$. It remains to prove $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$, and we start with a weaker claim:

Claim 4.9. $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau_2$.

Proof. By Lemma 4.8 we have that

$$h(\sigma) = I_h(xyA'Z_A; uB'Z_B|C) \geq I_h(xy; u|Z) = I_h(y; u|Z) + I_h(x; u|yZ).$$

Therefore, $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow (x; u|yZ)$, and in particular, $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow (x; u|yZ)$. Since $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow (x; y|Z)$ (i.e., by assumption), then by the chain rule we have that $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow (x; uy|Z) = \tau'$, and the claim follows from (4.1). \square

Finally, we prove $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$. Assume otherwise, and let h_W be a step function such that $h_W(\tau_2) = I_{h_W}(x; y|uZ) = 1$, and $h_W(\Sigma_2) = 0$. This means that $uZ \subseteq W$. Therefore $uZ_B \subseteq W$, implying $I_{h_W}(xyA'Z_A; uZ_B|C) = h_W(\sigma_1) = 0$ (because $uZ_B C \subseteq uZ$). Therefore, $h_W(\Sigma) = h_W(\sigma_1) + h_W(\Sigma_2) = 0$, contradicting the fact that $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau_2$. A more direct way to see this is to observe that σ_1 and τ are disjoint (Definition 4.6). Hence, by Lemma 4.7 we have that $\mathcal{S}_n \models_{\text{EI}} \Sigma \setminus \{\sigma_1\} \Rightarrow \tau_2$, and the result follows by noting that $\Sigma \setminus \{\sigma_1\} = \Sigma_2$.

4.2. Proof of Theorem 4.1 Item 2. We now prove Item 2 of Theorem 4.1. That is, we show that if τ is a conditional $Z \rightarrow X$ and $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$, then $\Gamma_n \models h(\Sigma) \geq h(\tau)$.

For the following lemma, we recall that by the chain rule for entropies (see (2.6)), we have that for a conditional $\tau = Z \rightarrow uX$ we have that:

$$h(\tau) = h(Z \rightarrow uX) = h(uX|Z) \stackrel{(2.6)}{=} h(u|Z) + h(X|uZ) = h(\underbrace{Z \rightarrow u}_{\tau_1}) + h(\underbrace{uZ \rightarrow X}_{\tau_2}) \quad (4.2)$$

Lemma 4.10. *Let Σ be a set of saturated CIs s.t. $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$. Suppose $\tau = (Z \rightarrow uX)$, and define $\tau_1 = (Z \rightarrow u)$, $\tau_2 = (uZ \rightarrow X)$ (see (4.2)); thus, $h(\tau) = h(\tau_1) + h(\tau_2)$. Then, there exists Σ_1 and Σ_2 such that: (1) $h(\Sigma) = h(\Sigma_1) + h(\Sigma_2)$. (2) Σ_1 covers τ_1 and $\mathcal{S}_n \models_{\text{EI}} \Sigma_1 \Rightarrow \tau_1$. (3) Σ_2 is saturated and $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$.*

Proof. We partition Σ into Σ_1 and Σ_2 as follows. For every $\sigma = (A; B|C) \in \Sigma$, if $u \in C$ then we place σ in Σ_2 . Otherwise, assume w.l.o.g that $u \in A$, and we write $A = uA_ZA_XA'$ where $A_Z = A \cap Z$, $A_X = A \cap X$, and $A' = A \setminus \{uA_ZA_X\}$. We use the chain rule to define σ_1, σ_2 :

$$I_h(A; B|C) = I_h(uA_ZA_XA'; B|C) = I_h(\underbrace{uA_Z; B|C}_{\stackrel{\text{def}}{=} \sigma_1}) + I(\underbrace{A_XA'; B|uA_ZC}_{\stackrel{\text{def}}{=} \sigma_2}) \quad (4.3)$$

We place σ_1 in Σ_1 , and σ_2 in Σ_2 . We observe that σ_1 covers τ_1 (because $Z = A_ZB_ZC_Z \subseteq A_ZBC$) and σ_2 is saturated. Furthermore, $h(\Sigma_1) + h(\Sigma_2) = h(\Sigma)$. We prove $\mathcal{S}_n \models_{\text{EI}} \Sigma_1 \Rightarrow \tau_1$. By assumption, $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau_1 = (Z \rightarrow u)$. Let any $\sigma_2 = (A; B|C) \in \Sigma_2$; since $u \in C$, by Lemma 4.7 we can remove it, obtaining $\mathcal{S}_n \models_{\text{EI}} \Sigma \setminus \{\sigma_2\} \Rightarrow \tau_1$; repeating this process proves $\mathcal{S}_n \models_{\text{EI}} \Sigma_1 \Rightarrow \tau_1$. Finally, we prove $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$. By assumption, $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau_2 = (uZ \rightarrow X)$. Let any $\sigma_1 = (uA_Z; B|C) \in \Sigma_1$; since $uA_Z \subseteq uZ$, by Lemma 4.7 we can remove it, obtaining $\mathcal{S}_n \models_{\text{EI}} \Sigma \setminus \{\sigma_1\} \Rightarrow \tau_2$; repeating this process proves $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$. \square

We now complete the proof of Theorem 4.1 item 2. Let $\tau = (Z \rightarrow X)$, and Σ be saturated (recall that by our assumption Σ contains only saturated CIs). We show, by induction on $|X|$, that if $\mathcal{S}_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ then $\Gamma_n \models h(\tau) \leq h(\Sigma)$. If $|X| = 1$, then $X = \{x\}$, $h(x|Z) = I_h(x; x|Z)$ is elemental, and the claim follows from Th. 4.5. Otherwise, let u be any variable in X , write $\tau = (Z \rightarrow uX')$, and apply Lemma 4.10 to $\tau_1 = (Z \rightarrow u)$, $\tau_2 = (Zu \rightarrow X')$, which gives us a partition of Σ into Σ_1, Σ_2 such that $h(\Sigma) = h(\Sigma_1) + h(\Sigma_2)$. On the one hand, $\mathcal{S}_n \models_{\text{EI}} \Sigma_1 \Rightarrow \tau_1$, and from Th. 4.5 we derive $h(\tau_1) \leq h(\Sigma_1)$ (because τ_1 is elemental, and covered by Σ_1); on the other hand $\mathcal{S}_n \models_{\text{EI}} \Sigma_2 \Rightarrow \tau_2$ where Σ_2 is saturated, which implies, by induction, $h(\tau_2) \leq h(\Sigma_2)$. The result follows from $h(\tau) = h(\tau_1) + h(\tau_2) \leq h(\Sigma_1) + h(\Sigma_2) = h(\Sigma)$, completing the proof of Theorem 4.1.

4.3. A special case. In Theorem 4.11 we show that under the assumption that the set of CIs in Σ is pairwise disjoint (Definition 4.6), we also obtain a 1-relaxation. The rather technical proof that relies on the I-measure is deferred to Section 8, where we present the I-measure.

Theorem 4.11. *Let Σ be a set of saturated, pairwise disjoint CI terms (Def. 4.6), and τ be a saturated CI. Then, $\mathcal{S}_n \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ implies $\Gamma_n \models h(\tau) \leq h(\Sigma)$.*

5. RELAXATION FOR GENERAL CIs

We now extend our discussion from saturated CIs and conditionals (or, equivalently, FDs and MVDs), to arbitrary Conditional Independence statements. We prove two results in this section. First, we prove that relaxation fails in general, and, second, we prove that a weaker form of relaxation holds.

In both results the relaxation problem is considered in $\mathbf{cl}(\Gamma_n^*)$. Recall that our golden standard is to check whether the relaxation problem holds in Γ_n^* . As we saw, when the constraints are restricted to FDs and MVDs, then the relaxation problem is the same in \mathcal{S}_n , in Γ_n^* , in $\mathbf{cl}(\Gamma_n^*)$, and in Γ_n . But for general constraints, these relaxation problems differ. Our first result in this section (the impossibility result) also holds in Γ_n^* , by Fact. 3.3, but we leave open the question whether the second result (weak relaxation) holds in Γ_n^* .

We state formally the two results in this section:

Theorem 5.1. *There exists Σ, τ with four variables, such that $\mathbf{cl}(\Gamma_4^*) \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ and $\mathbf{cl}(\Gamma_4^*) \not\models_{\text{APXI}} (\Sigma \Rightarrow \tau)$.*

Theorem 5.2. *Let Σ, τ be arbitrary CIs, and suppose $\mathbf{cl}(\Gamma_n^*) \models_{\text{EI}} \Sigma \Rightarrow \tau$. Then, for every $\varepsilon > 0$ there exists $\lambda > 0$ such that, for all $h \in \mathbf{cl}(\Gamma_n^*)$:*

$$h(\tau) \leq \lambda \cdot h(\Sigma) + \varepsilon \cdot h(\Omega) \quad (5.1)$$

We will prove both theorems shortly. Before we do this, however, we provide some background and context for these theorems, which requires us to review the concept of *cones*.

5.1. Cones. Both theorems 5.1 and 5.2 are best understood when viewed through the lens of convex analysis, in particular *cones*. We briefly review cones here, and refer to [Sch03, Stu93, BV04] for more details. Fix some number $N > 0$. A set $K \subseteq \mathbb{R}^N$ is called a *cone*, if for every $\mathbf{x} \in K$ and $\theta \geq 0$ we have that $\theta\mathbf{x} \in K$. A set $C \subseteq \mathbb{R}^N$ is called *convex* if, for any two points $\mathbf{x}_1, \mathbf{x}_2 \in C$ and any $\theta \in [0, 1]$, $\theta\mathbf{x}_1 + (1 - \theta)\mathbf{x}_2 \in C$. Unless otherwise stated, in this paper every cone will be assumed to be convex. The intersection of a, not necessarily finite, set of convex cones is also a convex cone. The *conic hull* of a set $C \subseteq \mathbb{R}^N$, in notation $\mathbf{conhull}(C)$, is the smallest convex cone that contains C , or, equivalently, it is the set of vectors of the form $\theta_1\mathbf{x}_1 + \dots + \theta_k\mathbf{x}_k$, where $k \geq 0$, $\mathbf{x}_1, \dots, \mathbf{x}_k \in C$, and $\theta_i \geq 0, \forall i \in [k]$.

Fix a vector $\mathbf{u} \in \mathbb{R}^N$. The set $K = \{\mathbf{x} \mid \mathbf{x} \cdot \mathbf{u} \leq 0\}$ is a convex cone called a *linear half-space*. A *polyhedral cone* is the intersection of a finite number of linear half-spaces. Equivalently, K is polyhedral if $K = \{\mathbf{x} \mid \mathbf{x} \cdot \mathbf{u}_1 \leq 0, \dots, \mathbf{x} \cdot \mathbf{u}_r \leq 0\}$, where $\mathbf{u}_1, \dots, \mathbf{u}_r \in \mathbb{R}^N$ are fixed vectors, or, also equivalently, $K = \{\mathbf{x} \mid \mathbf{x}^T A \leq 0\}$ where $A \in \mathbb{R}^{N \times r}$ is a matrix. A cone K is called *finitely generated* if $K = \mathbf{conhull}(C)$ for some finite set $C \subseteq \mathbb{R}^N$; equivalently, K is finitely generated if³ $K = \{A\mathbf{z} \mid \mathbf{z} \in \mathbb{R}^m, \mathbf{z} \geq 0\}$, for some matrix $A \in \mathbb{R}^{N \times m}$. Results by Farkas, Minkowski, and Weyl imply that a cone is finitely generated iff it is polyhedral [Sch03, pp.61].

As we discussed, an entropic function, or a polymatroid $h : 2^{[n]} \rightarrow \mathbb{R}_+$ can be seen as a vector $h \in \mathbb{R}^N$, where $N \stackrel{\text{def}}{=} 2^n$, in other words $\Gamma_n^*, \Gamma_n \subseteq \mathbb{R}^N$. By definition, Γ_n is a polyhedral cone, hence it is a finitely generated, convex cone. Yeung [Yeu08] has proven that, when $n \geq 3$, then Γ_n^* is neither convex, nor a cone, but its topological closure $\mathbf{cl}(\Gamma_n^*)$ is always a convex cone. When $n \leq 3$, $\mathbf{cl}(\Gamma_n^*) = \Gamma_n$ and thus $\mathbf{cl}(\Gamma_n^*)$ is finitely generated. For $n \geq 4$, $\mathbf{cl}(\Gamma_n^*)$ is not finitely generated [Mat07]. A conditional entropy $h(B|A) = h(AB) - h(A)$ is equal to $\mathbf{u} \cdot h$ where $\mathbf{u} \in \mathbb{R}^N$ is the vector having +1 on the dimension AB , -1 on the dimension A , and 0 everywhere else. Similarly, the mutual information $I_h(B; C|A) = h(AB) + h(AC) - h(ABC) - h(A)$ equals $\mathbf{v} \cdot h$ where \mathbf{v} is a vector with two +1's corresponding to dimensions AB and AC , two -1's corresponding to dimensions ABC and A , and the rest 0.

This discussion justifies phrasing the relaxation problem as follows. Fix a convex cone $K \subseteq \mathbb{R}^N$, and let $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_m$ be $m + 1$ vectors in \mathbb{R}^N . Relaxation asks whether statement (5.2) below implies statement (5.3):

$$\forall \mathbf{x} \in K : \quad \mathbf{x} \cdot \mathbf{y}_1 \leq 0 \wedge \dots \wedge \mathbf{x} \cdot \mathbf{y}_m \leq 0 \Rightarrow \mathbf{x} \cdot \mathbf{y}_0 \leq 0 \quad (5.2)$$

$$\exists \theta_1, \dots, \theta_m, \forall \mathbf{x} \in K : \quad \mathbf{x} \cdot \mathbf{y}_0 \leq \theta_1 \mathbf{x} \cdot \mathbf{y}_1 + \dots + \theta_m \mathbf{x} \cdot \mathbf{y}_m \quad (5.3)$$

³Here, and throughout the paper, the inequality $\mathbf{z} \geq 0$ means component wise inequalities.

When each \mathbf{y}_i has the property $\mathbf{x} \cdot \mathbf{y}_i \geq 0$ for all $\mathbf{x} \in K$ (as is the case with the vectors defining $h(B|A)$ and $I_h(B; C|A)$), then the condition $\mathbf{x} \cdot \mathbf{y}_i \leq 0$ is equivalent to $\mathbf{x} \cdot \mathbf{y}_i = 0$, and statement (5.2) is an Exact Implication. Furthermore, we can set all θ_i 's in statement (5.3) to be equal to $\lambda \stackrel{\text{def}}{=} \max_i \theta_i$, because $\mathbf{x} \cdot \mathbf{y}_i \geq 0$ implies $\sum_i \theta_i \mathbf{x} \cdot \mathbf{y}_i \leq \lambda \sum_i \mathbf{x} \cdot \mathbf{y}_i$, and statement (5.3) is an Approximate Implication. If we view relaxation at this level of generality, then it is easy to find cases where relaxation holds, and where relaxation fails:

Theorem 5.3. (1) If K is finitely generated, then statement (5.2) implies statement (5.3). (2) There exists a convex cone $K \subseteq \mathbb{R}^3$ where statement (5.2) does not imply statement (5.3).

Proof. (1) We give here a quick and simple proof based on Farkas lemma; in the next section we give in Theorem 6.1 a slightly more elaborate proof, based on the strong duality property (which is a consequence of Farkas' lemma), in order to obtain an upper bound on the relaxation coefficient.

More precisely, we use here the following version of Farkas lemma [Sch03, pp.61, Corollary 5.3a]. For any matrix $A \in \mathbb{R}^{N \times M}$ and vector $\mathbf{y} \in \mathbb{R}^N$ the following two statements are equivalent:⁴

- $\forall \mathbf{x} \in \mathbb{R}^N, \mathbf{x}^T A \leq 0$ implies $\mathbf{x}^T \mathbf{y} \leq 0$.
- There exists $\mathbf{z} \in \mathbb{R}^N, \mathbf{z} \geq 0$ such that $A\mathbf{z} = \mathbf{y}$.

If K is finitely generated, then it is polyhedral, hence $K = \{\mathbf{x} \mid \mathbf{x}^T \mathbf{u}_1 \leq 0 \wedge \dots \wedge \mathbf{x}^T \mathbf{u}_r \leq 0\}$. Let A be the $N \times (r + m)$ matrix whose columns are the vectors $\mathbf{u}_1, \dots, \mathbf{u}_r, \mathbf{y}_1, \dots, \mathbf{y}_m$. Then, statement (5.2) can be written equivalently as:

$$\forall \mathbf{x} \in \mathbb{R}^N : \quad \mathbf{x}^T A \leq 0 \Rightarrow \mathbf{x}^T \mathbf{y} \leq 0$$

Farkas lemma implies the existence of a vector $\mathbf{z} \in \mathbb{R}^{r+m}$ such that $\mathbf{z} \geq 0$ and $A\mathbf{z} = \mathbf{y}$. Denoting its components as $\mathbf{z} = (\gamma_1, \dots, \gamma_r, \theta_1, \dots, \theta_m)$, we have $\mathbf{y} = (\sum_j \gamma_j \mathbf{u}_j) + (\sum_i \theta_i \mathbf{y}_i)$ and therefore:

$$\forall \mathbf{x} \in K : \quad \mathbf{x} \cdot \mathbf{y} = \left(\sum_j \gamma_j \mathbf{x} \cdot \mathbf{u}_j \right) + \left(\sum_i \theta_i \mathbf{x} \cdot \mathbf{y}_i \right) \leq \sum_i \theta_i \mathbf{x} \cdot \mathbf{y}_i$$

since $\mathbf{x} \cdot \mathbf{u}_j \leq 0$, proving the claim.

(2) Let $K \subseteq \mathbb{R}^3$ be the following cone:

$$K = \{(x_1, x_2, x_3) \mid x_1 \geq 0, x_3 \geq 0, x_1 x_3 \geq x_2^2\}$$

Equivalently, K is the *positive semidefinite cone* [BV04], i.e. the set of semi-positive definite 2×2 matrices:

$$K \stackrel{\text{def}}{=} \left\{ A \stackrel{\text{def}}{=} \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \mid \forall \mathbf{u} \in \mathbb{R}^2 : \mathbf{u}^T A \mathbf{u} \geq 0 \right\}$$

It is immediate to check that K is a convex cone. Then K satisfies the following Exact Implication:

$$\forall (x_1, x_2, x_3) \in K : \quad x_1 \leq 0 \Rightarrow x_2 \leq 0$$

because $x_1 \leq 0$ is equivalent to $x_1 = 0$, implying $x_2^2 \leq 0$ thus $x_2 = 0$. However, K does not satisfy the corresponding Approximate Implication, more precisely the following is false:

$$\exists \lambda > 0, \forall (x_1, x_2, x_3) \in K : \quad x_2 \leq \lambda x_1 \quad (\text{this is false})$$

⁴The first statement is normally given with ≥ 0 instead of ≤ 0 . It is easy to revert the inequality by replacing A, \mathbf{y} with $-A, -\mathbf{y}$.

Indeed, for every choice of $\lambda > 0$, choose $0 < x_1 < 1/\lambda$, and let $x_2 = 1$, $x_3 = 1/x_1$. Then $(x_1, x_2, x_3) \in K$, yet $x_2 > \lambda x_1$. \square

5.2. Proof of Theorem 5.1. For $n \leq 3$, the set $\mathbf{cl}(\Gamma_n^*)$ is a polyhedral cone, and relaxation holds, by Theorem 5.3 (1). Thus, we need a counterexample with $n = 4$ jointly distributed random variables. The cone $\mathbf{cl}(\Gamma_4^*)$ is a subset of \mathbb{R}^{16} , hence the counterexample needed to prove Theorem 5.1 will be more complex than that used to prove Theorem 5.3 (2). For that purpose, we adapt an example by Kaced and Romashchenko [KR13, Inequality ($\mathcal{I}5'$) and Claim 5], built upon an earlier example by Matúš [Mat07].

Let Σ and τ be the following:

$$\Sigma = \{(C; D|A), (C; D|B), (A; B), (B; C|D)\} \quad \tau = (C; D) \quad (5.4)$$

We first prove that, for any $\lambda \geq 0$, there exists an entropic function h such that:

$$I_h(C; D) > \lambda \cdot (I_h(C; D|A) + I_h(C; D|B) + I_h(A; B) + I_h(B; C|D)) \quad (5.5)$$

Indeed, consider the distribution shown in Fig. 1 (c) (from [KR13]), where $0 < \varepsilon < 1/2$. We will prove that for this distribution, the following identities hold:

$$I_h(C; D) = \varepsilon + O(\varepsilon^2) \quad (5.6)$$

$$I_h(C; D|A) = I_h(C; D|B) = I_h(A; B) = 0 \quad (5.7)$$

$$I_h(B; C|D) = O(\varepsilon^2) \quad (5.8)$$

These three identities prove Eq.(5.5), by choosing ε small enough. These equalities were stated in [KR13, Claim 5], but no proof was provided; for completeness, we include here their proofs. In general, if X, Y are two joint random variables, then the conditional entropy $h(Y|X)$ satisfies the following equality: $h(Y|X) = \sum_x p(X = x)h(Y|X = x)$, where x ranges over the outcomes of the random variable X , $h(Y|X = x)$ is the standard entropy of the random variable Y conditioned on $X = x$, and, by convention, $p(X = x)h(Y|X = x) = 0$ when $p(X = x) = 0$. Similarly, $I_h(Y; Z|X) = \sum_x p(X = x)I_h(Y; Z|X = x)$. Furthermore, Y, Z are independent random variables iff $I_h(Y; Z) = 0$. Therefore, the expressions in Eq. (5.7) become:

$$I_h(C; D|A) = (1 - 2\varepsilon)I_h(C; D|A = 0) + 2\varepsilon I_h(C; D|A = 1) = 0 + 0 = 0$$

$$I_h(C; D|B) = \frac{1}{2}I_h(C; D|B = 0) + \frac{1}{2}I_h(C; D|B = 1) = 0 + 0 = 0$$

$$I_h(A; B) = 0$$

In the first line we used the fact that when $A = 0$ then C is constant and thus C, D are independent, and when $A = 1$ then D is constant. The same argument applies to the second line. For the last line it suffices to check that A, B are independent: $p(A = 1) = 2\varepsilon$, $p(B = 1) = 1/2$, and $p(A = 1, B = 1) = \varepsilon = p(A = 1) \cdot p(B = 1)$. (It is standard that the other combinations of values for A and B follow, for example $p(A = 1, B = 0) = p(A = 1) - p(A = 1, B = 1) = p(A = 1) - p(A = 1)p(B = 1) = p(A = 1)(1 - p(B = 1)) = p(A = 1)p(B = 0)$.) This proves Eq. (5.7).

For the proof of identities (5.6) and (5.8), we start by listing explicitly the probability distribution of C, D , then expand $I_h(B; C|D) = p(D = 0)I_h(B; C|D = 0) + p(B; C|D = 1)$ and list explicitly the probability distribution of B, C conditioned on $D = 0$ (since $I_h(B; C|D = 1) = 0$):

C	D		B	C	
0	0	1/2	0	0	$(1/2 - \varepsilon)/(1/2 + \varepsilon) \stackrel{\text{def}}{=} 1 - \varepsilon_1 - \varepsilon_2$
0	1	$1/2 - \varepsilon$	0	1	$\varepsilon/(1/2 + \varepsilon) \stackrel{\text{def}}{=} \varepsilon_1$
1	0	ε	1	0	$\varepsilon/(1/2 + \varepsilon) \stackrel{\text{def}}{=} \varepsilon_2$

where $\varepsilon_1 = \varepsilon_2 = O(\varepsilon)$ (we use different symbols $\varepsilon_1, \varepsilon_2$, even though they are equal, to make it easier to follow the calculations below). These two distributions are quite different (the first has the probability mass split almost 1/2 and 1/2, the second has the entire mass concentrated on the first outcome) and we need different techniques to compute the mutual information $I_h(\dots)$. We begin with $I_h(C; D)$:

$$I_h(C; D) = h(D) - h(D|C) = h(D) - p(C=0)h(D|C=0) = h(D) - (1-\varepsilon)h(D|C=0)$$

For any number $x \in (0, 1)$, denote by $B(x)$ a Bernoulli random variable X , with outcomes $p(X=0) = 1-x$ and $p(X=1) = x$. We denote by $f(x)$ its entropy, and $f'(x), f''(x)$ its derivatives:

$$\begin{aligned} f(x) &\stackrel{\text{def}}{=} -(x \log x + (1-x) \log(1-x)) = -(x \ln x + (1-x) \ln(1-x)) \log e \\ f'(x) &= -(\ln x - \ln(1-x)) \log e \\ f''(x) &= -\left(\frac{1}{x} + \frac{1}{1-x}\right) \log e \end{aligned}$$

Notice that $f'(1/2) = 0$, and $f''(1/2) = -4 \log e \approx -5.77$. We use Taylor's expansion to compute f in a vicinity of 1/2. For $0 < \theta < 1/2$, there exists $\tau \in (0, \theta)$ such that:

$$f(1/2 - \theta) = f(1/2) + \theta f'(1/2) + \frac{\theta^2}{2} f''(1/2 - \tau) = 1 - O(\theta^2)$$

The distribution of D is simply $B(1/2 - \varepsilon)$, and the distribution of $D|_{C=0}$ is $B\left(\frac{1/2-\varepsilon}{1-\varepsilon}\right) = B(1/2 - \delta)$ where $\delta = O(\varepsilon)$, therefore:

$$\begin{aligned} I_h(C; D) &= f(1/2 - \varepsilon) - (1 - \varepsilon)f(1/2 - \delta) \\ &= 1 - O(\varepsilon^2) - (1 - \varepsilon)(1 - O(\delta^2)) \\ &= \varepsilon + O(\delta^2) - O(\varepsilon^2) \end{aligned}$$

which proves (5.6).

Next, we compute $I_h(B; C|D=0)$, for which we apply the formula

$$I(X; Y) = \sum_{x,y} p(X=x, Y=y) \log \frac{p(X=x, Y=y)}{p(X=x)p(Y=y)}$$

We denote by $p_0(-) \stackrel{\text{def}}{=} p(-|D=0)$, therefore:

$$\begin{aligned} I_h(B; C|D=0) &= (1 - \varepsilon_1 - \varepsilon_2) \log \frac{1 - \varepsilon_1 - \varepsilon_2}{(1 - \varepsilon_2)(1 - \varepsilon_1)} + \varepsilon_1 \log \frac{\varepsilon_1}{(1 - \varepsilon_2)\varepsilon_1} + \varepsilon_2 \log \frac{\varepsilon_2}{\varepsilon_2(1 - \varepsilon_1)} \\ &= (1 - \varepsilon_1 - \varepsilon_2) \log \frac{1 - \varepsilon_1 - \varepsilon_2}{(1 - \varepsilon_2)(1 - \varepsilon_1)} + \varepsilon_1 \log \frac{1}{(1 - \varepsilon_2)} + \varepsilon_2 \log \frac{1}{(1 - \varepsilon_1)} \\ &= -((1 - \varepsilon_1 - \varepsilon_2) \ln(1 - \varepsilon_0) + \varepsilon_1 \ln(1 - \varepsilon_2) + \varepsilon_2 \ln(1 - \varepsilon_1)) \log e \end{aligned}$$

where:

$$\varepsilon_0 = 1 - \frac{(1 - \varepsilon_2)(1 - \varepsilon_1)}{1 - \varepsilon_1 - \varepsilon_2} = \frac{\varepsilon_1 \varepsilon_2}{1 - \varepsilon_1 - \varepsilon_2} = O(\varepsilon_1 \varepsilon_2)$$

Finally, we use the fact that $-\ln(1 - x) = O(x)$ and obtain:

$$I_h(B; C|D = 0) = (1 - \varepsilon_1 - \varepsilon_2)O(\varepsilon_0) + \varepsilon_1 O(\varepsilon_2) + \varepsilon_2 O(\varepsilon_1) = O(\varepsilon_1 \varepsilon_2)$$

which proves equation (5.8)

Next, we prove $\mathbf{cl}(\Gamma_n^*) \models_{\text{EI}} (\Sigma \Rightarrow \tau)$. This follows from an inequality initially proven by Matúš [Mat07], then adapted by [KR13]. For completeness, we review here that inequality, starting with the statement of Theorem 2 in [Mat07], which asserts that for every entropic vector $h \in \Gamma_5^*$ over 5 variables, and every natural number $k \geq 1$:

$$k \left(\square_{13,24} + \Delta_{34|5} + \Delta_{45|3} \right) + \Delta_{35|4} + \frac{k(k-1)}{2} \left(\Delta_{24|3} + \Delta_{34|2} \right) \cdot h \geq 0$$

where, using our notation:

$$\Delta_{YZ|X} \cdot h \stackrel{\text{def}}{=} I_h(Y; Z|X) \quad \square_{AB,CD} \cdot h \stackrel{\text{def}}{=} I_h(C; D|A) + I_h(C; D|B) + I_h(A; B) - I_h(C; D)$$

Substituting $A = 1, B = 3, C = 4, D = 2, E = 5$ in Matúš's inequality and dividing by k , we obtain the following (which is Eq.(ii) in Theorem 2 of [KR13]):

$$\begin{aligned} & (I_h(C; D|A) + I_h(C; D|B) + I_h(A; B) - I_h(C; D)) \\ & + I_h(B; C|E) + I_h(C; E|B) + \frac{1}{k} I_h(B; E|C) + \frac{k-1}{2} (I_h(C; D|B) + I_h(B; C|D)) \geq 0 \end{aligned}$$

Finally, we set $E = D$ to obtain the following inequality, for all $h \in \Gamma_n^*$ and $k \geq 1$:

$$I_h(C; D) \leq I_h(C; D|A) + \frac{k+3}{2} I_h(C; D|B) + I_h(A; B) + \frac{k+1}{2} I_h(B; C|D) + \frac{1}{k} I_h(B; D|C) \quad (5.9)$$

By continuity, the inequality also holds for $\mathbf{cl}(\Gamma_n^*)$ too. We can prove now that the Exact Implication $\Sigma \Rightarrow \tau$ holds in $\mathbf{cl}(\Gamma_n^*)$. Assume $I_h(C; D|A) = I_h(C; D|B) = I_h(A; B) = I_h(B; C|D) = 0$. Then inequality (5.9) implies $0 \leq I_h(C; D) \leq \frac{1}{k} I_h(B; D|C)$, for any $k \geq 1$. Letting $k \rightarrow \infty$, implies $I_h(C; D) = 0$.

It is interesting to observe that inequality (5.9) is almost a relaxation of the implication (5.4): the only extra term is the last term, which can be made arbitrarily small by increasing k . Our second result generalizes this.

5.3. Proof and Discussion of Theorem 5.2. The proof of Theorem 5.2 follows from a more general statement about cones:

Theorem 5.4. *Let $K \subseteq \mathbb{R}^N$ be a topologically closed, convex cone, and let $\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_m$ be $m + 1$ vectors in \mathbb{R}^N . The following are equivalent:*

$$\forall \mathbf{x} \in K : \quad \mathbf{x} \cdot \mathbf{y}_1 \leq 0, \dots, \mathbf{x} \cdot \mathbf{y}_m \leq 0 \Rightarrow \mathbf{x} \cdot \mathbf{y}_0 \leq 0 \quad (5.10)$$

$$\forall \varepsilon > 0, \exists \theta_1, \dots, \theta_m \geq 0, \forall \mathbf{x} \in K, \quad \mathbf{x} \cdot \mathbf{y}_0 \leq \theta_1 \mathbf{x} \cdot \mathbf{y}_1 + \dots + \theta_m \mathbf{x} \cdot \mathbf{y}_m + \varepsilon \|\mathbf{x}\|_\infty \quad (5.11)$$

We first show that Theorem 5.4 implies Theorem 5.2. For this purpose we take $K \stackrel{\text{def}}{=} \mathbf{cl}(\Gamma_n^*)$, which is a closed, convex cone [Yeu08]. Let $\Sigma = \{(B_1; C_1|A_1), \dots, (B_m; C_m|A_m)\}$, and $\tau = (B_0; C_0|A_0)$. We define the vectors \mathbf{y}_i such that $\mathbf{y}_i \cdot h = I_h(B_i; C_i|A_i)$, and notice that the conditions $\mathbf{y}_i \cdot h \leq 0$ and $\mathbf{y}_i \cdot h = 0$ are equivalent, because $I_h(B_i; C_i|A_i) \geq 0$ always

holds. If the Exact Implication $\Sigma \Rightarrow \tau$ holds, then condition (5.10) holds. This implies condition (5.11), and inequality (5.1) in Theorem 5.2 follows by setting $\lambda \stackrel{\text{def}}{=} \max_i \theta_i$, and observing that $\|h\|_\infty = h(\Omega)$ (i.e., this is because $\max_{i \in N} h_i = h_N = h(\Omega)$).

In the rest of this section we prove Theorem 5.4. While we only need the implication (5.10) \Rightarrow (5.11), it helps to observe that the reverse holds too. Indeed, assuming $\mathbf{x} \cdot \mathbf{y}_i \leq 0$ for $i = 1, m$, the condition (5.11) implies that, for every $\varepsilon > 0$: $\mathbf{x} \cdot \mathbf{y}_0 \leq \varepsilon \|\mathbf{x}\|$. By taking $\varepsilon \rightarrow 0$ we obtain $\mathbf{x} \cdot \mathbf{y}_0 \leq 0$.

To prove the implication (5.10) \Rightarrow (5.11) we need to review some properties of cones. For any set $C \subseteq \mathbb{R}^N$, its *dual*, $C^* \subseteq \mathbb{R}^N$ is the following set:

$$C^* \stackrel{\text{def}}{=} \{\mathbf{y} \mid \forall \mathbf{x} \in C, \mathbf{x} \cdot \mathbf{y} \geq 0\} \quad (5.12)$$

It is immediate to check that C^* is a topologically closed, convex cone (because C^* is the intersection of, possibly infinitely many, linear half-spaces, each of which is a topologically closed, convex cone), formally: $C^* = \mathbf{cl}(C^*)$. We warn that the $*$ in Γ_n^* does *not* represent the dual; the notation Γ_n^* for entropic functions is by now well established, and we adopt it here too, despite its clash with the standard notation for the dual cone.

We need the following basic properties of cones:

- (A) If L is a finite set, then $\mathbf{conhull}(L)$ is topologically closed.
- (B) For any set K , $\mathbf{cl}(\mathbf{conhull}(K)) = K^{**}$.

- (C) If K_1 and K_2 are closed, convex cones then: $(K_1 \cap K_2)^* = \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^*$.

Item (A) follows immediately from the definition: $\mathbf{conhull}(\{\mathbf{x}_1, \dots, \mathbf{x}_r\}) = \{\sum_i \theta_i \mathbf{x}_i \mid \theta_1, \dots, \theta_i \in [0, 1]\}$. Item (B) is well-known and non-trivial, see for example [Stu93, Fact 6].

Item (C) is perhaps less well-known, and we include the proof here. We will use repeatedly the anti-monotonicity of $(-)^*$: $C_1 \subseteq C_2$ implies $C_1^* \supseteq C_2^*$. We first prove the inclusion:

$$\left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^* \subseteq (K_1 \cap K_2)^* \quad (5.13)$$

as follows: for $i \in \{1, 2\}$, $K_i \supseteq K_1 \cap K_2$, which implies $K_i^* \subseteq (K_1 \cap K_2)^*$, thus $K_1^* \cup K_2^* \subseteq (K_1 \cap K_2)^*$. Eq. (5.13) follows from the fact that $(K_1 \cap K_2)^*$ is both a convex cone and topologically closed. Next, we prove:

$$(K_1 \cap K_2)^* \subseteq \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^{**} \quad (5.14)$$

For $i \in \{1, 2\}$, we have that $K_i^* \subseteq \mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right)$, and therefore it holds that $K_i^{**} \supseteq \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^*$. Since K_1 and K_2 are closed convex cones then by item (B) it holds that $K_1^{**} = K_1$ and $K_2^{**} = K_2$. Therefore, for $i \in \{1, 2\}$ we have that $K_i \supseteq \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^*$. From the above we get that $K_1 \cap K_2 \supseteq \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^*$, which implies $(K_1 \cap K_2)^* \subseteq \left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^{**}$. By item (B) we have that $\left(\mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right) \right)^{**} = \mathbf{cl} \left(\mathbf{conhull} (K_1^* \cup K_2^*) \right)$, proving Eq. (5.14).

We now have the tools needed to prove Theorem 5.4. Assume condition (5.10) holds, and let $\varepsilon > 0$ be some number. Denote by $L \stackrel{\text{def}}{=} \{-\mathbf{y}_1, -\mathbf{y}_2, \dots, -\mathbf{y}_m\}$. The LHS of the implication (5.10) defines the set $\{\mathbf{x} \mid \mathbf{x} \in K, \mathbf{x} \cdot \mathbf{y}_1 \leq 0, \dots, \mathbf{x} \cdot \mathbf{y}_m \leq 0\} = \{\mathbf{x} \mid \mathbf{x} \in K \cap L^*\}$. Thus, implication (5.10) states $x \in K \cap L^* \Rightarrow \mathbf{x} \cdot \mathbf{y}_0 \leq 0$, which is equivalent to $-\mathbf{y}_0 \in (K \cap L^*)^*$. We rewrite this set as:

$$\begin{aligned} (K \cap L^*)^* &= \mathbf{cl} \left(\mathbf{conhull} (K^* \cup L^{**}) \right) && \text{Item (C)} \\ &= \mathbf{cl} \left(\mathbf{conhull} \left(K^* \cup \mathbf{cl} \left(\mathbf{conhull} (L) \right) \right) \right) && \text{Item (B)} \\ &= \mathbf{cl} \left(\mathbf{conhull} (K^* \cup \mathbf{conhull} (L)) \right) && \text{Item (A)} \\ &= \mathbf{cl} \left(\mathbf{conhull} (K^* \cup L) \right) && \text{Definition of } \mathbf{conhull} (-) \end{aligned}$$

Thus, $-\mathbf{y}_0 \in \mathbf{cl}(\mathbf{conhull}(K^* \cup L))$. By the definition of the topological closure, for all $\delta > 0$, there exists a vector in $\mathbf{conhull}(K^* \cup L)$ that is δ -close to $-\mathbf{y}_0$; in particular, we choose $\delta = \varepsilon/N$. Equivalently, there exists a vector $\mathbf{e} \in \mathbb{R}^N$ such that $\|\mathbf{e}\|_\infty < \varepsilon/N$ and:

$$-\mathbf{y}_0 + \mathbf{e} \in \mathbf{conhull}(K^* \cup L)$$

By definition of the cone hull, there exists $\mathbf{u} \in K^*$ and $\theta_1, \dots, \theta_m \geq 0$ such that:

$$-\mathbf{y}_0 + \mathbf{e} = \mathbf{u} - \theta_1 \mathbf{y}_1 - \dots - \theta_m \mathbf{y}_m$$

or, equivalently:

$$\theta_1 \mathbf{y}_1 + \dots + \theta_m \mathbf{y}_m - \mathbf{y}_0 + \mathbf{e} = \mathbf{u} \in K^*$$

By the definition of K^* , it follows that, for every $\mathbf{x} \in K$:

$$\theta_1 \mathbf{x} \cdot \mathbf{y}_1 + \dots + \theta_m \mathbf{x} \cdot \mathbf{y}_m - \mathbf{x} \cdot \mathbf{y}_0 + \mathbf{x} \cdot \mathbf{e} \geq 0$$

Statement (5.11) follows by observing that $|\mathbf{x} \cdot \mathbf{e}| \leq N \|\mathbf{x}\|_\infty \cdot \|\mathbf{e}\|_\infty < \varepsilon \|\mathbf{x}\|_\infty$.

6. RESTRICTED AXIOMS: SHANNON INEQUALITIES

The results in the previous section are mostly negative: for general constraints, relaxation fails, with the only exception of MVDs and FDs, where relaxation holds. In this section we prove that relaxation holds in general for constraints that can be inferred using only Shannon inequalities (monotonicity (2.2), and submodularity (2.3)). Equivalently, this means interpreting the constraints over the set of all polymatroids, Γ_n .

While the golden standard is the implication problem in $\mathbf{cl}(\Gamma_n^*)$, a study of the implication problem in Γ_n is important for several reasons. First, by restricting to Shannon inequalities we obtain a sound, but in general incomplete (w.r.t. $\mathbf{cl}(\Gamma_n^*)$) method for deciding implications. The incompleteness stems from the fact that the set of entropic functions (and their limit points) $\mathbf{cl}(\Gamma_n^*)$ obey additional inequalities (and hence, implications) beyond those that result from the Shannon inequalities. These are called *non-Shannon-Type inequalities* [ZY97, MMRV02], and as their name suggests, they are not implied by the Shannon inequalities. For example, Matúš's inequality (5.9) is a non-Shannon inequality. It holds only in $\mathbf{cl}(\Gamma_n^*)$ but fails in Γ_n . Second, while generally incomplete, Shannon's inequalities are complete for characterizing the inequalities that hold under certain syntactic restrictions. In particular, it follows from our results in Section 4 that they are complete for FDs and MVDs, and

it follows from results in [Ken21] that they are complete for marginal CIs⁵: in both cases, relaxation holds, with a factor $\lambda = n^2/4$.

We have already seen in Theorem 5.3 that exact implication of CIs relax over Γ_n . We start by proving an upper bound on the coefficient of the relaxation.

Theorem 6.1. *Let Σ, τ be arbitrary CIs. If $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$, then $\Gamma_n \models h(\tau) \leq (2^n)! \cdot h(\Sigma)$. In other words, CIs admit relaxation over Γ_n with coefficient $\lambda \leq (2^n)!$.*

Proof. We start by proving a lemma:

Lemma 6.2. *Consider a system of linear equations, $Ax = b$, where A is an $M \times N$ matrix with $\text{rank}(A) = M$. Suppose all entries in A and b are $-1, 0$, or 1 . Then there exists a solution x such that, for all $i = 1, N$, $|x_i| \leq M!$*

Proof. Let A_0 be a sub-matrix of A consisting of M linearly independent columns. W.l.o.g. we can assume that these are the first M columns, and write A as $A = [A_0|A_1]$. The equation $Ax = b$ becomes $A_0x_0 + A_1x_1 = b$, where x_0 represent the first M coordinates of x , and x_1 are the remaining $N - M$ coordinates. We set $x_1 = 0$, and solve $A_0x_0 = b$ using Cramer's rule: for all $i = 1, M$, $(x_0)_i = \Delta_i/\Delta$, where $\Delta = \det(A_0) \neq 0$ and Δ_i is the determinant of the matrix A_0 where the i 'th column is replaced with b . Each determinant is the sum of $M!$ terms, and each term is $-1, 0$, or 1 . It follows that $|\Delta_i| \leq M!$, hence $|(x_0)_i| \leq M!/1 = M!$. \square

Next, we briefly review two facts from linear programming, see e.g. [Sch03]. Let $A \in \mathbb{R}^{M \times N}$, $b \in \mathbb{R}^M$, $c \in \mathbb{R}^N$. Then:

- The strong duality theorem states that the primal LP and the dual LP have the same optimal values:

$$\max\{c^T x \mid x \in \mathbb{R}^N, x \geq 0, Ax \leq b\} = \min\{y^T b \mid y \in \mathbb{R}^M, y \geq 0, y^T A \geq c^T\}$$

- If the primal linear program has a finite optimal solution, then it has an optimal solution x that is a vertex of the polytope $\{x \in \mathbb{R}^N \mid x \geq 0, Ax \leq b\}$. In particular, if $r \stackrel{\text{def}}{=} \text{rank}(A)$, then there exists an optimal solution x of the primal LP that satisfies $A_0x = b$, where A_0 is some $r \times N$ sub-matrix consisting of r independent rows. As a consequence, if all entries in A, b are -1 or 0 or $+1$, then, by Lemma 6.2, there exists an optimal solution x satisfying $|x_i| \leq r!$, for every coordinate $i = 1, N$.

We prove now Theorem 6.1. Fix a set of n variables, Ω , and let Γ_n be the set of all polymatroids over variables Ω . Γ_n is defined by Shannon's inequalities, monotonicity (2.2), and submodularity (2.3), and it is known that it suffices to take only the *elemental* inequalities, i.e. inequalities of the following form (see [Yeu08, (14.12)]):

$$\forall X \in \Omega : \quad h(\Omega - X) - h(\Omega) \leq 0 \quad \text{monotonicity}$$

$$\forall X, Y \in \Omega, W \subseteq \Omega - \{X, Y\} : \quad h(W) - h(WX) - h(WY) + h(WXY) \leq 0 \quad \text{submodularity}$$

There are n elemental monotonicity constraints, and $\frac{n(n-1)}{2}2^{n-2} = n(n-1)2^{n-3}$ submodularity constraints. The total number of Shannon inequalities is $n + n(n-1)2^{n-3}$. Equivalently, we can write:

$$\Gamma_n = \{h \mid h \in \mathbb{R}^{2^n}, h \geq 0, A_S h \leq 0\} \quad (6.1)$$

⁵Marginal CIs have the form $(X; Y)$ (i.e., no conditioning)

where A_S is the $(n + n(n-1)2^{n-3}) \times 2^n$ matrix corresponding to all Shannon inequalities. Similarly, the constraints Σ on h can be defined by some $m \times 2^n$ matrix A_Σ , in other words:

$$\{h \mid h \models \Sigma\} = \{h \mid h \in \mathbb{R}^{2^n}, A_\Sigma h \leq 0\}$$

A_Σ has one row for each constraint in Σ . For example, if Σ contains the CI $Z_1; Z_2|V$, then one row in the matrix A_Σ corresponds to the assertion $I_h(Z_1; Z_2|V) \leq 0$, or, equivalently, $-h(V) + h(Z_1V) + h(Z_2V) - h(Z_1Z_2V) \leq 0$. Let $M \stackrel{\text{def}}{=} (n + n(n-1)2^{n-3}) + m$ and $N \stackrel{\text{def}}{=} 2^n$, and let A be the following $M \times N$ matrix:

$$A \stackrel{\text{def}}{=} \begin{bmatrix} A_S \\ A_\Sigma \end{bmatrix}$$

Finally, define $b \in \mathbb{R}^M$ the 0-vector, $b^T = (0, 0, \dots, 0)$, and let $c_\tau \in \mathbb{R}^N$ be the vector corresponding to the constraint τ . More precisely, assume $\tau = (X; Y|W)$. Then $c_\tau^T h$ is the linear expression $-h(W) + h(WX) + h(WY) - h(WXY)$.

With these notations, we claim that the Exact Implication $\Gamma_n \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ holds iff the optimal solution of the linear program below has value zero:

$$\max\{c_\tau^T h \mid h \in \mathbb{R}^{2^n}, h \geq 0, Ah \leq b\}$$

Recall that $b = 0$. Indeed, the constraints $h \in \mathbb{R}^{2^n}, h \geq 0, Ah \leq b$ define all polymatroids $h \in \Gamma_n$ that satisfy Σ , while $c_\tau^T h = I_h(X; Y|W)$ is the value of consequent τ . We always have $\max\{c_\tau^T h \mid h \in \mathbb{R}^{2^n}, h \geq 0, Ah \leq b\} \geq 0$, because the zero polymatroid $h \stackrel{\text{def}}{=} 0$ is a feasible solution to the LP. Moreover, if the Exact Implication holds, then every feasible solution is a polymatroid that satisfies Σ , hence it satisfies $I_h(X; Y|W) \leq 0$, in other words $c_\tau^T h \leq 0$, proving that the optimal solution to the LP is ≤ 0 . We conclude that, if the Exact Implication holds, then the optimal is $= 0$, which proves the claim.

Next, assume that the exact implication holds, thus the primal LP has optimal value 0. By the strong duality theorem, the dual LP also has optimal value 0:

$$\max\{c_\tau^T h \mid h \in \mathbb{R}^{2^n}, h \geq 0, Ah \leq b\} = \min\{y^T b \mid y \in \mathbb{R}^M, y \geq 0, y^T A \geq c_\tau^T\} = 0$$

Since $b = 0$, we have $y^T b = 0$, hence, when the exact implication holds, then $\min\{0 \mid y \in \mathbb{R}^M, y \geq 0, y^T A \geq c_\tau^T\} = 0$. Equivalently, this asserts that the constraint $y \geq 0, y^T A \geq c_\tau^T$ has a feasible solution y : otherwise the optimal of the dual solution is $\min \emptyset = \infty$.

We observe now that every feasible solution y to the dual LP represents a relaxation of the implication problem. Indeed, let's write y as $y = \begin{bmatrix} y_S \\ y_\Sigma \end{bmatrix}$, where y_S consists of the first $(n + n(n-1)2^{n-3})$ coordinates, and y_Σ of the last m coordinates of y . Since y is feasible, we have $y^T A = y_S^T A_S + y_\Sigma^T A_\Sigma \geq c_\tau^T$. Let $h \in \Gamma_n$ be any polymatroid; since $h \geq 0$, we have:

$$c_\tau^T h \leq (y_S^T A_S + y_\Sigma^T A_\Sigma)h = y_S^T (A_S h) + y_\Sigma^T (A_\Sigma h) \leq y_\Sigma^T (A_\Sigma h)$$

In the last inequality we used the fact that $A_S h \leq 0$, by definition of Γ_n in Eq. (6.1). The m coordinates of the vector $A_\Sigma h \in \mathbb{R}^m$ are the values $I_h(Z_1; Z_2|V)$, for all $(Z_1; Z_2|V) \in \Sigma$. Thus, $y_\Sigma^T (A_\Sigma h)$ can be written as a positive linear combination of the constraints in Σ , and the inequality $c_\tau^T h \leq y_\Sigma^T (A_\Sigma h)$ is precisely an Approximate Implication. If $r \stackrel{\text{def}}{=} \text{rank}(A)$, then by our discussion above, we can find a feasible solution y whose coordinates are

$\leq r! \leq (2^n)!$, since $r \leq \min(M, N) = N = 2^n$. In other words, $y \leq (2^n)!(1 \ 1 \ \dots \ 1)^T$, where $(1 \ 1 \ \dots \ 1)^T$ is the all-1 vector. Thus, we have:

$$I_h(X; Y|W) = c_\tau^T h \leq y_\Sigma^T (A_\Sigma h) \leq (2^n)!(1 \ 1 \ \dots \ 1)^T (A_\Sigma h) \leq (2^n)! \sum_{(Z_1; Z_2|V) \in \Sigma} I_h(Z_1; Z_2|V)$$

This completes the proof of the theorem. \square

Theorem 6.1 gives us a very crude upper bound on the relaxation factor for Γ_n . Next, we show a lower bound on the factor λ . We prove a lower bound of 3:

Theorem 6.3 [DFZ09]. *The following inequality holds for all polymatroids $h \in \Gamma_n$:*

$$h(Z) \leq I_h(A; B|C) + I_h(A; B|D) + I_h(C; D|E) + I_h(A; E) + 3h(Z|A) + 2h(Z|B) \quad (6.2)$$

but the inequality fails if any of the coefficients 3, 2 are replaced by smaller values. In particular, denoting τ, Σ the terms on the two sides of Eq.(6.2), the exact implication $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ holds⁶, and does not have a 1-relaxation.

Proof. We make use the following inequality that was proved in Lemma 1 in [DFZ09].

$$h(Z|R) + I_h(R; S|T) \geq I_h(Z; S|T) \quad (6.3)$$

We apply (6.3) three times:

- (1) $h(Z|A) + I_h(A; B|C) \geq I_h(Z; B|C)$
- (2) $h(Z|A) + I_h(A; B|D) \geq I_h(Z; B|D)$
- (3) $h(Z|A) + I_h(A; E) \geq I_h(Z; E)$

Plugging back into the formula we get that:

$$\begin{aligned} & I_h(A; B|C) + I_h(A; B|D) + I_h(C; D|E) + I_h(A; E) + 3h(Z|A) + 2h(Z|B) \\ & \geq I_h(Z; B|C) + I_h(Z; B|D) + I_h(Z; E) + I_h(C; D|E) + 2h(Z; B) \end{aligned} \quad (6.4)$$

We now apply this identity twice more:

- (1) $h(Z|B) + I_h(Z; B|C) \geq I_h(Z; Z|C) = h(Z|C)$
- (2) $h(Z|B) + I_h(Z; B|D) \geq I_h(Z; Z|D) = h(Z|D)$

Plugging back into (6.4) we get that:

$$\begin{aligned} & I_h(Z; B|C) + I_h(Z; B|D) + I_h(Z; E) + I_h(C; D|E) + 2h(Z|B) \\ & \geq h(Z|C) + h(Z|D) + I_h(Z; E) + I_h(C; D|E) \\ & = h(ZC) - h(C) + h(ZD) - h(D) + h(Z) + h(E) - h(ZE) + h(CE) + h(DE) - h(E) - h(CDE) \\ & = I_h(Z; E|C) + h(ZEC) + I_h(Z; E|D) + h(ZED) + h(Z) - h(ZE) - h(CDE) \\ & = I_h(Z; E|C) + I_h(Z; E|D) + I_h(C; D|ZE) + h(CDZE) - h(CDE) + h(Z) \\ & = I_h(Z; E|C) + I_h(Z; E|D) + I_h(C; D|ZE) + h(Z|CDE) + h(Z) \\ & \geq h(Z) \end{aligned} \quad \square$$

We remark that inequality (6.2) can be verified using known tools for testing whether an inequality holds for all polymatroids (e.g., ITIP⁷, and XITIP⁸). It is still open whether

⁶This means that if $h(\Sigma) = 0$ where $\Sigma = \{(A; B|C), (A; B|D), \dots, h(Z|B)\}$, then $h(\tau) = h(Z) = 0$. That is, Z is deterministic.

⁷<https://user-www.ie.cuhk.edu.hk/~ITIP/>

⁸<http://xitip.epfl.ch/>

the coefficients 3 and 2 (for $h(Z|A)$ and $h(Z|B)$, respectively) are tight. To show that, we would need to present a polymatroid for which (6.2) holds with equality. Using XITIP, we were able to test that the coefficients 3 and 2 could not be reduced even by 0.0001. While this does not rule out the possibility of having the inequality hold for some coefficient $3 - \epsilon$ for a small-enough ϵ , it does allow us to conclude that the exact implication corresponding to (6.2) does not have a 1-relaxation.

7. RESTRICTED MODELS: POSITIVE I-MEASURE

While relaxation fails in its most general setting, we have seen that relaxation holds if we either restrict the type of constraints to FDs and MVDs, or if we restrict the implications to those that can be inferred from Shannon’s inequalities. In this section we consider a different restriction: we will restrict the types of models, or databases, over which the constraints are interpreted: more precisely we restrict the set of entropic functions to the step functions, \mathcal{S}_n , or, equivalently, to their conic hull, which we denote by $\mathcal{P}_n \stackrel{\text{def}}{=} \mathbf{conhull}(\mathcal{S}_n)$. In Section 8 we show that these entropic functions are precisely those with a *positive I-measure*, a notion introduced by Yeung [Yeu91, Yeu08]. In this section, we prove that all EIs admit a 1-relaxation over entropic functions with positive I-measure (i.e. over \mathcal{P}_n):

Theorem 7.1. *Every implication $\Sigma \Rightarrow \tau$ admits a 1-relaxation over \mathcal{P}_n , where Σ, τ are arbitrary CIs. In other words, if $\mathcal{P}_n \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ then $\forall h \in \mathcal{P}_n, h(\tau) \leq h(\Sigma)$.*

The golden standard for the semantics of constraints is to interpret them over Γ_n^* , hence the reader may wonder what we gain by restricting them to (the conic hull of) the step functions, or, equivalently, what we can learn by checking implications only on the uniform 2-tuple distributions. We have two motivations. First, checking an implication only on the uniform 2-tuple distributions leads to a complete, but unsound procedure for checking implication over Γ_n^* . In other words, by testing an implication on all uniform 2-tuple distributions we can detect if an implication *fails*, thus, the procedure is complete. Of course, the procedure is not sound, because an implication may hold on all step functions but fail in general. For a simple example, the inequality $I_h(X; Y|Z) \leq I_h(X; Y)$ holds for all step functions, but fails on the “parity function” in Fig. 1 (b). The second motivation is more interesting. It turns out that restricting the models to uniform 2-tuple distributions leads to a sound and complete procedure in some important special cases. We saw one such case in Section 4: when the constraints are restricted to FDs and MVDs, then in order to check an implication in Γ_n^* , it suffices to check that it holds on all uniform 2-tuple distributions. We will present in this section a second case: checking *differential constraints* in market basket analysis [SVG05].

In Section 8 we use the I-measure theory to characterize the conic hull of step functions \mathcal{P}_n , and provide an alternative proof to the main results of this section: Theorems 7.1 and 7.4.

7.1. I-Measure Constraints. Recall the following definition in information theory:

Definition 7.2. Fix a set Ω of n variables, and let $h : 2^\Omega \rightarrow \mathbb{R}$ be any function. Let $W, Y \subseteq \Omega$ be two sets of variables. Then *conditional multi-variate mutual information* is

defined as:

$$I_h(Y|W) \stackrel{\text{def}}{=} - \sum_{Z:W \subseteq Z \subseteq W \cup Y} (-1)^{|Z-W|} h(Z) \quad (7.1)$$

The quantity $I_h(Y|W)$ is usually written as $I_h(y_1; \dots; y_m|W)$, where $Y = \{y_1, \dots, y_m\}$. In the literature, $I_h(Y|W)$ is defined only for entropic functions h , but we remove this restriction here. When $m = 2$ then this is precisely the mutual information $I_h(y_1; y_2|W) = -h(W) + h(y_1W) + h(y_2W) - h(y_1y_2W)$, and when $m = 1$ then it becomes the conditional $I_h(y|W) = -h(W) + h(yW) = h(y|W)$. Notice that $I_h(Y|W) = I_h(Y - W|W)$. When $Y \cup W = \Omega$, then we say that $I_h(Y|W)$ is *saturated*.

Definition 7.3. Fix a set Ω of n variables. An *I-measure constraint statement* is a formula of the form $Y|W$, where $W, Y \subseteq \Omega$. We call the constraint *saturated* if $Y \cup W = \Omega$. We say that a function $h : 2^\Omega \rightarrow \mathbb{R}$ *satisfies* the constraint statement, if $I_h(Y|W) = 0$. An *implication* is a formula $\Sigma \Rightarrow (Y|W)$, where Σ is a set of I-measure constraints and $Y|W$ is an I-measure constraint. Fix a set K s.t. $\mathcal{S}_n \subseteq K \subseteq \Gamma_n$. The *exact implication* $\Sigma \Rightarrow (Y|W)$ holds in K , in notation $K \models_{\text{EI}} \Sigma \Rightarrow (Y|W)$, if $\forall h \in K, \bigwedge_{X|V \in \Sigma} I_h(X|V) = 0$ implies $I_h(Y|W) = 0$. The λ -*approximate implication* holds in K , if $\forall h \in K, I_h(Y|W) \leq \lambda \sum_{X|V \in \Sigma} I_h(X|V)$. We say that the exact implication problem admits a λ -*relaxation* in K , for some $\lambda > 0$, if every exact implication $K \models_{\text{EI}} \Sigma \Rightarrow (Y|W)$ implies the λ -approximate implication $\forall h \in K, I_h(Y|W) \leq \lambda \sum_{X|V \in \Sigma} I_h(X|V)$.

We prove our main result in this section:

Theorem 7.4. *Exact implications of I-measure constraints admit a 1-relaxation in \mathcal{P}_n . More precisely, if $\Sigma \models (Y|W)$ is an implication of I-measure constraints and $\mathcal{P}_n \models_{\text{EI}} \Sigma \Rightarrow (Y|W)$, then $\forall h \in \mathcal{P}_n, I_h(Y|W) \leq \sum_{X|V \in \Sigma} I_h(X|V)$.*

Notice that Theorem 7.1 follows immediately as a special case, since every CI is, in particular, an I-measure constraint.

To prove theorem 7.4, we need to prove two lemmas. In both lemmas below we will use a simple property. For any two sets A, B :

$$\sum_{C:A \subseteq C \subseteq B} (-1)^{|C-A|} = \begin{cases} 1 & \text{when } A = B \\ 0 & \text{otherwise} \end{cases} \quad (7.2)$$

To prove the identity of (7.2), it suffices to show that $\sum_{C:A \subseteq C \subseteq B} (-1)^{|C-A|} = 0$ when $A \subsetneq B$: fix any element $b \in B - A$, and notice that the sets C that don't contain b are in 1-1 correspondence with those that do contain b , via the mapping $C \mapsto C \cup \{b\}$. Since the two sets C and $C \cup \{b\}$ have different parities, their terms cancel out, $(-1)^{|C-A|} + (-1)^{|C \cup \{b\} - A|} = 0$, and the entire sum is zero.

Recall from Section 2 that h_Z denotes the ‘‘step function at Z ’’, where $Z \subsetneq \Omega$ (see the definition in Eq. (2.9)). By definition, their conic hull $\mathcal{P}_n \stackrel{\text{def}}{=} \mathbf{conhull}(\mathcal{S}_n)$ consists of all functions of the form $h = \sum_{Z \subsetneq \Omega} d_Z h_Z$, where $d_Z \geq 0$, for $Z \subsetneq \Omega$ are arbitrary coefficients. These coefficients are precisely the saturated conditional multi-variate mutual information. This follows from a more general lemma:

Lemma 7.5. *Consider the vector space $\mathbf{V} = \{h \mid h : 2^\Omega \rightarrow \mathbb{R}, h(\emptyset) = 0\}$. Then, the $2^n - 1$ step functions $\{h_Z \mid Z \subsetneq \Omega\}$ are basis for \mathbf{V} . Moreover, for every function $h \in \mathbf{V}$, denoting by $d_Z \stackrel{\text{def}}{=} I_h(\Omega - Z|Z)$, for all $Z \subsetneq \Omega$ (where I_h is defined in Eq. (7.1)), we have: $h =$*

$\sum_{Z \subsetneq \Omega} d_Z h_Z$. In other words, the projections of h on the basis consisting of the step functions are precisely the saturated conditional multi-variate mutual informations.

Proof. We prove the following claim: for every $W \subseteq \Omega$, it holds that

$$h(W) = \sum_{Z: Z \subsetneq \Omega} d_Z h_Z(W),$$

where $d_Z \stackrel{\text{def}}{=} I_h(\Omega - Z|Z)$. The claim implies that the step functions span the entire vector space $\{h \mid h: 2^\Omega \rightarrow \mathbb{R}, h(\emptyset) = 0\}$. This immediately implies that they are linearly independent, hence are a basis for \mathbf{V} . Indeed, the dimensionality of \mathbf{V} is $2^n - 1$, and this coincides with the number of step functions: since they span the entire vector space, they must be linearly independent. It remains to prove the claim.

To prove the claim, for any $W \subseteq \Omega$, define $f(W) \stackrel{\text{def}}{=} h(\Omega) - h(W)$. We want to express $f(W)$ as $f(W) = \sum_{Z: W \subsetneq Z \subsetneq \Omega} g(Z)$, for some function $g: 2^\Omega \rightarrow \mathbb{R}$. Such a function g is the unique Möbius inverse of f . Recall that the Möbius inversion formula states that the two expressions below are equivalent:

$$\forall W: f(W) = \sum_{Z: W \subsetneq Z \subsetneq \Omega} g(Z) \quad \forall W: g(W) = \sum_{Z: W \subsetneq Z \subsetneq \Omega} (-1)^{|Z-W|} f(Z) \quad (7.3)$$

Since $f(\emptyset) = h(\Omega)$, the first equality implies:

$$h(\Omega) = f(\emptyset) = \sum_{Z: Z \subsetneq \Omega} g(Z)$$

Therefore, the left equation in (7.3) can be rewritten as:

$$h(W) = h(\Omega) - f(W) = \sum_{Z: Z \subsetneq \Omega} g(Z) - \sum_{Z: W \subsetneq Z \subsetneq \Omega} g(Z) = \sum_{Z: W \not\subsetneq Z} g(Z) \quad (7.4)$$

On the other hand, we can express $g(W)$ in terms of the conditional multi-variate information, by expanding $f(Z) = h(\Omega) - h(Z)$ in the right equation in (7.3), then using Eq. (7.2):

$$\begin{aligned} g(W) &= \left(\sum_{Z: W \subsetneq Z} (-1)^{|Z-W|} \right) h(\Omega) - \left(\sum_{Z: W \subsetneq Z} (-1)^{|Z-W|} h(Z) \right) \\ &= \begin{cases} h(\Omega) - h(\Omega) = 0 & \text{if } W = \Omega \\ 0 + I_h(\Omega - W|W) & \text{if } W \subsetneq \Omega \end{cases} \end{aligned}$$

We compute now $h(W)$ from (7.4):

$$h(W) = \sum_{Z: W \not\subsetneq Z} g(Z) = \sum_{Z: W \not\subsetneq Z} I_h(\Omega - Z|Z) = \sum_{Z \subsetneq \Omega} I_h(\Omega - Z|Z) h_Z(W)$$

where the last equality holds because $h_Z(W) = 1$ when $W \not\subsetneq Z$ and $h_Z(W) = 0$ otherwise. This proves the claim, as required. \square

The lemma implies that, if $h \in \mathcal{P}_n$, then $I_h(\Omega - Z|Z) \geq 0$ for all $Z \subsetneq \Omega$. Indeed, by definition of the conic hull $\mathcal{P}_n = \mathbf{conhull}(\mathcal{S}_n)$, the function h can be written as $h = \sum_{Z: Z \subsetneq \Omega} d_Z h_Z$, where $d_Z \geq 0$. By the lemma, we have $d_Z = I_h(\Omega - Z|Z)$, proving the claim.

The second lemma is:

Lemma 7.6. *For any two sets $W, Y \subseteq \Omega$ s.t. $Y \neq \emptyset$, the following identity holds:*

$$I_h(Y|W) = \sum_{V:W \subseteq V \subseteq \Omega - Y} I_h(\Omega - V|V)$$

Proof. We expand I_h in the RHS according to its definition:

$$\begin{aligned} \text{RHS} &= \sum_{V:W \subseteq V \subseteq \Omega - Y} I_h(\Omega - V|V) = - \sum_{V:W \subseteq V \subseteq \Omega - Y} \left(\sum_{T:V \subseteq T \subseteq \Omega} (-1)^{|T-V|} h(T) \right) \\ &= - \sum_{T:W \subseteq T} h(T) \left(\sum_{V:(W \subseteq V \subseteq \Omega - Y) \wedge (V \subseteq T)} (-1)^{|T-V|} \right) \\ &= - \sum_{T:W \subseteq T} h(T) \left(\sum_{V:(W \subseteq V \subseteq (T \cap (\Omega - Y)))} (-1)^{|T-V|} \right) \end{aligned}$$

By Eq. (7.2), the inner sum is $= 1$ when $W = T \cap (\Omega - Y)$ and $= 0$ otherwise. The condition $W = T \cap (\Omega - Y)$ is equivalent to $W \subseteq T \subseteq W \cup Y$ and thus we obtain:

$$\text{RHS} = - \sum_{T:W \subseteq T \subseteq W \cup Y} (-1)^{|T-W|} h(T)$$

The latter expression is equal to $I_h(Y|W)$ by Def. 7.2. \square

We will now prove Theorem 7.4. Consider an exact implication $\mathcal{P}_n \models_{\text{EI}} \Sigma \Rightarrow (Y|W)$. By Lemma 7.6, $I_h(Y|W) = \sum_{V:W \subseteq V \subseteq \Omega - Y} I_h(\Omega - V|V)$. We claim that, for every set V s.t. $W \subseteq V \subseteq \Omega - Y$, there exists a constraint $X|U \in \Sigma$ such that $U \subseteq V \subseteq \Omega - X$. In other words, if we expand $I_h(X|U)$ according to Lemma 7.6, then one of the terms will be $I_h(\Omega - V|V)$. To prove the claim, we consider the step function at V , h_V , and use the fact that the exact implication must hold for h_V . We notice that $I_{h_V}(\Omega - V|V) = 1$ and $I_{h_V}(\Omega - U|U) = 0$ for $U \neq V$: this follows by considering the expansion of h_V given by Lemma 7.5, $h_V = \sum_{U:U \subseteq \Omega} d_U h_U$, where $d_U = I_{h_V}(\Omega - U|U)$ and noting that, since h_V is part of the basis, we have $I_{h_V}(\Omega - V|V) = 1$ and $I_{h_V}(\Omega - U|U) = 0$ for $U \neq V$. In particular, $I_{h_V}(Y|W) = 1$, in other words, the constraint $Y|W$ does not hold in h_V . Since the exact implication holds, it must be the case that Σ does not hold for h_V either, hence there exists an I-constraint $X|U \in \Sigma$ such that $I_{h_V}(X|U) > 0$. When expanding it according to Lemma 7.6, $I_{h_V}(X|U) = \sum_{T:U \subseteq T \subseteq \Omega - X} I_{h_V}(\Omega - T|T)$. Since $I_{h_V}(\Omega - T|T) = 0$ for all $T \neq V$, we conclude that one of the terms must be $I_{h_V}(\Omega - V|V)$, proving the claim.

We use claim to prove the theorem. Let h be any function in \mathcal{P}_n . To prove the inequality $I_h(Y|W) \leq \sum_{X|U \in \Sigma} I_h(X|U)$, we expand both sides:

$$\begin{aligned} I_h(Y|W) &= \sum_{V:W \subseteq V \subseteq \Omega - Y} I_h(\Omega - V|V) \\ \sum_{(X|U) \in \Sigma} I_h(X|U) &= \sum_{T:U \subseteq T \subseteq \Omega - X, (X|U) \in \Sigma} I_h(\Omega - T|T) \end{aligned}$$

We have shown that every term $I_h(\Omega - V|V)$ in the first sum occurs at least once as a term $I_h(\Omega - T|T)$ in the second sum. Since all terms are ≥ 0 , it follows that $I_h(Y|W) \leq \sum_{(X|U) \in \Sigma} I_h(X|U)$, as required.

7.2. Differential Constraints in Market Basket Analysis. We end this section by describing the tight connection between I-measure constraints and differential constraints for Market Basket Analysis, introduced by Sayrafi and Van Gucht [SVG05].

Consider a set of items $\Omega = \{X_1, \dots, X_n\}$, and a set of baskets $\mathcal{B} = \{b_1, \dots, b_N\}$ where every basket is a subset $b_i \subseteq \Omega$. The *support function* $f^{\mathcal{B}} : 2^\Omega \rightarrow \mathbb{N}$ assigns to every subset $W \subseteq \Omega$ the number of baskets in \mathcal{B} that contain the set W :

$$f^{\mathcal{B}}(W) = |\{i \mid i \in [n], W \subseteq b_i\}| \quad (7.5)$$

The function $f^{\mathcal{B}}$ is anti-monotone: $W_1 \subseteq W_2$ implies $f^{\mathcal{B}}(W_1) \geq f^{\mathcal{B}}(W_2)$. Similarly define $d^{\mathcal{B}}(W)$ the number of baskets in \mathcal{B} that are equal to the set W :

$$d^{\mathcal{B}}(W) = |\{i \mid i \in [n], W = b_i\}|$$

Then the following two identities are easily verified:

$$\forall W \subseteq \Omega : \quad f^{\mathcal{B}}(W) = \sum_{Z:W \subseteq Z} d^{\mathcal{B}}(Z) \quad d^{\mathcal{B}}(W) \stackrel{\text{def}}{=} \sum_{Z:W \subseteq Z} (-1)^{|Z-W|} f^{\mathcal{B}}(Z) \quad (7.6)$$

The first identity follows immediately from the definitions of $f^{\mathcal{B}}$ and $d^{\mathcal{B}}$. The second is Möbius' inversion function. Building on the identity (7.6), Sayrafi and Van Gucht define the *density* of any function f as follows:

Definition 7.7. Let $f : 2^\Omega \rightarrow \mathbb{R}$ be any function. Its *density* is the unique function $d_f : 2^\Omega \rightarrow \mathbb{R}$ is defined as:

$$\forall W \subseteq \Omega : \quad d_f(W) \stackrel{\text{def}}{=} \sum_{Z:W \subseteq Z} (-1)^{|Z-W|} f(Z)$$

When f is the support function $f^{\mathcal{B}}$ associated to a set of baskets \mathcal{B} , then its density d_f is precisely the function $d^{\mathcal{B}}$, because it satisfies the left equation in (7.6).

For any function $d : 2^\Omega \rightarrow \mathbb{R}$ and any two sets of items $W, Y \subseteq \Omega$, we denote by $d(Y|W) \stackrel{\text{def}}{=} \sum_{V:W \subseteq V \subseteq W \cup Y} d(V)$.

A *differential constraint* is an expression $Y|W$. A set of baskets \mathcal{B} *satisfies* the differential constraint⁹, if $d^{\mathcal{B}}(Y|W) = 0$. Sayrafi and Van Gucht [SVG05] define an implication, to be an assertion $\Sigma \Rightarrow (Y|W)$, where Σ is a set of differential constraints. In addition to that, we also define here an approximate implication:

Definition 7.8. Fix a set of items Ω . An *implication* is a formula $\Sigma \Rightarrow (Y|W)$ where Σ is a set of differential constraints, and $Y|W$ is one differential constraint. The *exact implication* $\Sigma \Rightarrow (Y|W)$ holds, in notation $\models_{\text{EI}} (\Sigma \Rightarrow (Y|W))$ if, for every set of baskets \mathcal{B} , $\bigwedge_{(X|V) \in \Sigma} d^{\mathcal{B}}(X|V) = 0$ implies $d^{\mathcal{B}}(Y|W) = 0$. The *λ -approximate implication* holds, if, for every set of baskets \mathcal{B} , $d^{\mathcal{B}}(Y|W) \leq \lambda \sum_{(X|V) \in \Sigma} d^{\mathcal{B}}(X|V)$. We write $\models_{\text{APX}} (\Sigma \Rightarrow (Y|W))$ if there exists a finite $\lambda \geq 0$ such that the λ -approximate implication holds.

In order to check any implication, it suffices to check whether it holds for *singleton* sets of baskets, $\mathcal{B} = \{b\}$. This follows from the linearity of the functions $f^{\mathcal{B}}$ and $d^{\mathcal{B}}$: if $\mathcal{B} = \{b_1, \dots, b_N\}$, then $f^{\mathcal{B}}(W) = \sum_{i=1, N} f^{\{b_i\}}(W)$ and $d^{\mathcal{B}}(W) = \sum_i d^{\{b_i\}}(W)$. An Exact Implication $\Sigma \Rightarrow (Y|W)$ holds, iff for every singleton sets of baskets $\mathcal{B} = \{b\}$,

⁹The notation used by Sayrafi and Van Gucht [SVG05] is $W \rightarrow Y$. They actually define slightly more general differential constraints, where Y is allowed to be a set of sets of items instead of a set of items. To simplify the presentation, in this paper we only consider Y as a set of items.

$\bigwedge_{(X|V) \in \Sigma} d^{\{b\}}(X|V) = 0$ implies $d^{\{b\}}(Y|W) = 0$, and similarly for a λ -approximate implication.

With this observation, we can now explain the connection to information theory. For any $b \subseteq \Omega$, we denote as usual h_b the step function at b when $b \neq \Omega$, and denote $h_b = 0$ (the constant function 0) when $b = \Omega$. Then:

Theorem 7.9. *Fix a set Ω of n items, and let $\mathcal{B} = \{b\}$ be a set consisting of a single basket, $b \subseteq \Omega$. The following hold:*

- For all $W \subseteq \Omega$: $f^{\{b\}}(W) = 1 - h_b(W)$.
- For all $W \subseteq \Omega$: $d^{\{b\}}(W) = I_{h_b}(\Omega - W|W)$.
- Every exact implication $\models_{\text{EI}} \Sigma \Rightarrow W$ relaxes to a 1-approximate implication: $d^{\{b\}}(W) \leq \sum_{V \in \Sigma} d^{\{b\}}(V)$.

Proof. The first item follows immediately from the definitions: $f^{\{b\}}(W) = 1$ iff $W \subseteq b$, iff $h_b(W) = 0$. The second item follows also from the definitions:

$$d^{\{b\}}(W) = \sum_{Z:W \subseteq Z} f^{\{b\}}(Z) = \sum_{Z:W \subseteq Z} (1 - h_b(Z)) = \begin{cases} 1 - h_b(\Omega) = 0 = I_{h_b}(0|\Omega) & \text{if } W = \Omega \\ 0 + I_{h_b}(\Omega - W|W) & \text{otherwise} \end{cases}$$

(We used Eq. (7.2).) Finally, the last item follows immediately from the second item and from Theorem 7.4. \square

Example 7.10. Consider five items $\Omega = \{A, B, C, D, E\}$. The differential constraint $C|AB$ asserts that “every customer who bought items A, B also bought item C ”. Indeed, if a set of baskets \mathcal{B} satisfies the constraint $C|AB$ then $d^{\mathcal{B}}(C|AB) = d^{\mathcal{B}}(AB) + d^{\mathcal{B}}(ABD) + d^{\mathcal{B}}(ABE) + d^{\mathcal{B}}(ABDE) = 0$; this simply says that no basket exists that contains both A and B but does not contain C . For a second example, consider the differential constraint $CD|AB$. It asserts “every customer who bought items A, B also bought either C or D , or both”, because $d^{\mathcal{B}}(CD|AB) = d^{\mathcal{B}}(AB) + d^{\mathcal{B}}(ABE) = 0$ implies that no basket exists that contains AB and none of C or D .

Consider now the following implication: $(CD|\emptyset) \Rightarrow (CD|AB)$. Obviously, the exact implication holds, and the 1-approximate implication holds too, because $d^{\mathcal{B}}(CD|AB) = d^{\mathcal{B}}(AB) + d^{\mathcal{B}}(ABE) \leq d^{\mathcal{B}}(CD|\emptyset) = d^{\mathcal{B}}(\emptyset) + d^{\mathcal{B}}(A) + d^{\mathcal{B}}(B) + d^{\mathcal{B}}(E) + d^{\mathcal{B}}(AB) + d^{\mathcal{B}}(AE) + d^{\mathcal{B}}(BE) + d^{\mathcal{B}}(ABE)$.

8. ADDITIONAL RESULTS AND PROOFS THAT RELY ON THE I-MEASURE

The I-measure, developed by Yeung [Yeu91, Yeu08], is a theory that establishes a one-to-one correspondence between Shannon’s information measures (i.e., entropy, conditional entropy, mutual information, and conditional mutual information) and set theory. In Section 8.1, we use the I-measure to prove Theorem 4.11. In Section 8.2 we provide an alternative proof to Theorems 7.1 and 7.4. We begin by briefly describing the main ideas and theorems of the I-measure theory that we use in the remainder of this section; for further details and examples, see chapter 3 of [Yeu08].

Let $h \in \Gamma_n$ denote a polymatroid defined over the variables $\{X_1, \dots, X_n\}$. As mentioned, the I-Measure theory formulates a one-to-one correspondence between Shannon’s information measures and set theory. Hence, every variable X_i of the polymatroid, is associated with a set $m(X_i)$. The universal set is $\Lambda \stackrel{\text{def}}{=} \bigcup_{i=1}^n m(X_i)$, the union of all sets associated with the

Information Measures	μ^*
$H(X)$	$\mu^*(\mathfrak{m}(X))$
$H(XY)$	$\mu^*(\mathfrak{m}(X) \cup \mathfrak{m}(Y))$
$H(X Y)$	$\mu^*(\mathfrak{m}(X) \cap \mathfrak{m}^c(Y))$
$I_H(X; Y)$	$\mu^*(\mathfrak{m}(X) \cap \mathfrak{m}(Y))$
$I_H(X; Y Z)$	$\mu^*(\mathfrak{m}(X) \cap \mathfrak{m}(Y) \cap \mathfrak{m}^c(Z))$

TABLE 3. Information measures and associated I-measure

variables. The complement of the set $\mathfrak{m}(X_i)$ is $\mathfrak{m}^c(X_i) \stackrel{\text{def}}{=} \Lambda \setminus \mathfrak{m}(X_i)$. Let $\mathcal{I} \subseteq [n]$. We recall that $X_{\mathcal{I}}$ is the joint random variable $X_{\mathcal{I}} \stackrel{\text{def}}{=} (X_i : i \in \mathcal{I})$ (see Notation 2.1), and we denote $\mathfrak{m}(X_{\mathcal{I}}) \stackrel{\text{def}}{=} \bigcup_{i \in \mathcal{I}} \mathfrak{m}(X_i)$.

Definition 8.1. The field \mathcal{F}_n generated by sets $\mathfrak{m}(X_1), \dots, \mathfrak{m}(X_n)$ is the collection of sets which can be obtained by any sequence of usual set operations (union, intersection, complement, and difference) on $\mathfrak{m}(X_1), \dots, \mathfrak{m}(X_n)$.

The *atoms* of \mathcal{F}_n are sets of the form $\bigcap_{i=1}^n Y_i$, where Y_i is either $\mathfrak{m}(X_i)$ or $\mathfrak{m}^c(X_i)$. We denote by \mathcal{A} the atoms of \mathcal{F}_n . Note that by our choice of the universal set, $\Lambda \stackrel{\text{def}}{=} \bigcup_{i=1}^n \mathfrak{m}(X_i)$, the atom $\bigcap_{i=1}^n \mathfrak{m}^c(X_i)$ degenerates to the empty set. This is because:

$$\bigcap_{i=1}^n \mathfrak{m}^c(X_i) = (\mathfrak{m}(X_1) \cup \dots \cup \mathfrak{m}(X_n))^c = \Lambda^c = \emptyset \quad (8.1)$$

Hence, there are $2^n - 1$ non-empty atoms in \mathcal{F}_n .

We call a function $\mu : \mathcal{F}_n \rightarrow \mathbb{R}$ *set additive* if, for every pair of disjoint sets A and B , it holds that $\mu(A \cup B) = \mu(A) + \mu(B)$. A real function μ defined on \mathcal{F}_n is called a *signed measure* if it is set additive, and $\mu(\emptyset) = 0$.

The I-measure μ^* on \mathcal{F}_n is defined by

$$\mu^*(\mathfrak{m}(X_{\mathcal{I}})) = H(X_{\mathcal{I}}) \quad (8.2)$$

where $\mathcal{I} \subseteq \{1, \dots, n\}$, and H is the entropy (see (2.8)). Theorem 8.2 [Yeu91, Yeu08] establishes the one-to-one correspondence between Shannon's information measures and μ^* .

Theorem 8.2 [Yeu91, Yeu08]. μ^* is the unique signed measure on \mathcal{F}_n which is consistent with all Shannon's information measures (i.e., entropies, conditional entropies, mutual information, and conditional mutual information).

Let $X, Y, Z \subseteq \Omega$ denote three subsets of variables; Z may be empty. We refer to objects of the form (X) and $(X|Z)$ as entropy and conditional entropy terms, respectively. Likewise, we refer to triples of the form $(X; Y|Z)$ as conditional mutual information terms. Collectively, we refer to these as information-theoretic terms. Let $\sigma = (X; Y|Z)$ be a mutual information term. We denote by $\mathfrak{m}(\sigma) = \mathfrak{m}(X) \cap \mathfrak{m}(Y) \cap \mathfrak{m}^c(Z)$ the (I-Measure) set associated with σ (see [Yeu08] for intuition and examples). Table 3 summarizes the extension of the I-measure μ^* to the rest of the Shannon measures (i.e., conditional entropy, mutual information, and conditional mutual information). For a set of mutual information and entropy terms Σ , we

let:

$$m(\Sigma) \stackrel{\text{def}}{=} \bigcup_{\sigma \in \Sigma} m(\sigma).$$

Theorem 8.3 [Yeu08]. *Let $\Omega = \{X_1, \dots, X_n\}$, and let \mathcal{F}_n be the field generated by sets $m(X_1), \dots, m(X_n)$. Let $\mu_+ : \mathcal{F}_n \rightarrow \mathbb{R}_+$ be any set-additive function that assigns non-negative values to the atoms of \mathcal{F}_n . Then the function $H : 2^\Omega \rightarrow \mathbb{R}_+$, defined as $H(X_\alpha) = \mu_+(m(X_\alpha))$ is a polymatroid (i.e., meets the polymatroid inequalities (2.2) and (2.3)).*

In the following Lemma, we apply Theorem 8.3 to characterize exact implication in the polymatroid cone Γ_n .

Lemma 8.4. *Let Σ denote a set of CIs. If $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$ then $m(\tau) \subseteq m(\Sigma)$.*

Proof. Assume, by contradiction, that $m(\tau) \not\subseteq m(\Sigma)$, and let $b \in m(\tau) \setminus m(\Sigma)$. By Theorem 8.3 the I -measure μ^* can take the following non-negative values:

$$\mu^*(a) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}$$

It is evident that $\mu^*(\Sigma) = 0$ while $\mu^*(\tau) = 1$, which contradicts the implication. \square

8.1. Proof from Section 4. In this section we formally prove Theorem 4.11 from Section 4 that relies on the I -measure. For brevity, we denote the intersection of the sets corresponding to variables A and B (e.g., $m(A) \cap m(B)$) by $m(A)m(B)$.

THEOREM 4.11. *Let Σ be a set of saturated, pairwise disjoint CI terms (Def. 4.6), and τ be a saturated CI. Then, $\mathcal{S}_n \models_{\text{EI}} (\Sigma \Rightarrow \tau)$ implies $\Gamma_n \models h(\tau) \leq h(\Sigma)$.*

Proof. From Corollary 4.2 we have that $\Gamma_n \models_{\text{EI}} \Sigma \Rightarrow \tau$. Let $\tau \stackrel{\text{def}}{=} (A; B|C)$ be the saturated implied CI, and let $\sigma = (X; Y|Z) \in \Sigma$ be a saturated CI. Recalling that $m(\tau) = m^c(C)m(A)m(B)$ (see Table 3)¹⁰ we get that:

$$\begin{aligned} m^c(\tau) &= (m^c(C)m(A)m(B))^c \\ &= (m^c(C))^c \cup m^c(A) \cup m^c(B) \\ &= m(C) \cup m^c(C)m^c(A) \cup m^c(C)m(A)m^c(B) \end{aligned} \quad (8.3)$$

Furthermore, since $\sigma = (X; Y|Z)$ is saturated, we get that $C = C_X C_Y C_Z$ where $C_X = C \cap X$, $C_Y = C \cap Y$, and $C_Z = C \cap Z$. From this, we get that:

$$m(C) = m(C_X) \cup m^c(C_X)m(C_Y) \cup m^c(C_X)m^c(C_Y)m(C_Z) \quad (8.4)$$

Then, from (8.3), and the set-additivity of μ^* , we get that:

$$\begin{aligned} &\mu^*(m(\sigma) \cap m^c(\tau)) \\ &= \mu^*(m(\sigma) \cap m(C)) + \mu^*(m(\sigma) \cap m^c(C)m^c(A)) + \mu^*(m(\sigma) \cap m^c(C)m(A)m^c(B)) \end{aligned} \quad (8.5)$$

¹⁰And recall that $m(A) \cap m(B)$ is denoted $m(A)m(B)$

We consider each one of the terms in (8.5) separately. For the first term, by (8.4), we have that:

$$\begin{aligned} & \mu^*(\mathfrak{m}(\sigma) \cap \mathfrak{m}(C)) \\ &= \mu^*(\mathfrak{m}(\sigma)\mathfrak{m}(C_X)) + \mu^*(\mathfrak{m}(\sigma)\mathfrak{m}^c(C_X)\mathfrak{m}(C_Y)) + \mu^*(\mathfrak{m}(\sigma)\mathfrak{m}^c(C_X)\mathfrak{m}^c(C_Y)\mathfrak{m}(C_Z)) \\ &= \mu^*(\mathfrak{m}(C_X)\mathfrak{m}(Y)\mathfrak{m}^c(Z)) + \mu^*(\mathfrak{m}(X \setminus C_X)\mathfrak{m}(C_Y)\mathfrak{m}^c(ZC_X)) \\ & \quad + \mu^*(\mathfrak{m}(X \setminus C_X)\mathfrak{m}(Y \setminus C_Y)\mathfrak{m}(C_Z)\mathfrak{m}^c(ZC_X C_Y)) \end{aligned} \quad (8.6)$$

$$= \mu^*(\mathfrak{m}(C_X)\mathfrak{m}(Y)\mathfrak{m}^c(Z)) + \mu^*(\mathfrak{m}(X \setminus C_X)\mathfrak{m}(C_Y)\mathfrak{m}^c(ZC_X)) \quad (8.7)$$

$$= I_h(C_X; Y|Z) + I_h(X \setminus C_X; C_Y|ZC_X) \quad (8.8)$$

where transition (8.6) follows from substituting $\mathfrak{m}(\sigma) = \mathfrak{m}(X)\mathfrak{m}(Y)\mathfrak{m}^c(Z)$, and transition (8.7) is because $C_Z \subseteq Z$ and thus $\mathfrak{m}(C_Z)\mathfrak{m}^c(ZC_X C_Y) = \emptyset$, and $\mu^*(\emptyset) = 0$. We now consider the second term of (8.5).

$$\begin{aligned} \mu^*(\mathfrak{m}(\sigma) \cap \mathfrak{m}^c(C)\mathfrak{m}^c(A)) &= \mu^*(\mathfrak{m}(X)\mathfrak{m}(Y)\mathfrak{m}^c(Z)\mathfrak{m}^c(C)\mathfrak{m}^c(A)) \\ &= \mu^*(\mathfrak{m}(X)\mathfrak{m}(Y)\mathfrak{m}^c(ZCA)) \\ &= I_h(X; Y|ZCA) \end{aligned} \quad (8.9)$$

Finally, we consider the third term of (8.5). Since $\sigma = (X; Y|Z)$ is saturated, we get that $A = A_X A_Y A_Z$ where $A_X = A \cap X$, $A_Y = A \cap Y$, and $A_Z = A \cap Z$. Therefore:

$$\mathfrak{m}(A) = \mathfrak{m}(A_X) \cup \mathfrak{m}^c(A_X)\mathfrak{m}(A_Y) \cup \mathfrak{m}^c(A_X)\mathfrak{m}^c(A_Y)\mathfrak{m}(A_Z) \quad (8.10)$$

From (8.10), and the set-additivity of μ^* , we get that:

$$\begin{aligned} & \mu^*(\mathfrak{m}(\sigma) \cap \mathfrak{m}^c(C)\mathfrak{m}(A)\mathfrak{m}^c(B)) \\ &= \mu^*(\mathfrak{m}(X)\mathfrak{m}(Y)\mathfrak{m}^c(Z)\mathfrak{m}^c(C)\mathfrak{m}(A)\mathfrak{m}^c(B)) \\ &= \mu^*(\mathfrak{m}(X)\mathfrak{m}(Y)\mathfrak{m}(A)\mathfrak{m}^c(ZCB)) \\ &= \mu^*(\mathfrak{m}(A_X)\mathfrak{m}(Y)\mathfrak{m}^c(ZCB)) + \mu^*(\mathfrak{m}(X \setminus A_X)\mathfrak{m}(A_Y)\mathfrak{m}^c(ZCBA_X)) \\ & \quad + \mu^*(\mathfrak{m}(X \setminus A_X)\mathfrak{m}(Y \setminus A_Y)\mathfrak{m}(A_Z)\mathfrak{m}^c(ZCBA_X A_Y)) \end{aligned} \quad (8.11)$$

$$= \mu^*(\mathfrak{m}(A_X)\mathfrak{m}(Y)\mathfrak{m}^c(ZCB)) + \mu^*(\mathfrak{m}(X \setminus A_X)\mathfrak{m}(A_Y)\mathfrak{m}^c(ZCBA_X)) \quad (8.12)$$

$$= I_h(A_X; Y|ZCB) + I_h(X \setminus A_X; A_Y|ZCBA_X) \quad (8.13)$$

where transition (8.11) follows from substituting the formula for $\mathfrak{m}(A)$ (see (8.10)), and where transition (8.12) is because $A_Z \subseteq Z$, and thus $\mathfrak{m}(A_Z) \cap \mathfrak{m}^c(ZCBA_X) = \emptyset$, and $\mu^*(\emptyset) = 0$. By (8.8), (8.9), and (8.13) we get that:

$$\begin{aligned} \mu^*(\mathfrak{m}(\sigma) \cap \mathfrak{m}^c(\tau)) &= I_h(C_X; Y|Z) + I_h(X \setminus C_X; C_Y|ZC_X) \\ & \quad + I_h(X; Y|ZCA) + I_h(A_X; Y|ZCB) + I_h(X \setminus A_X; A_Y|ZCBA_X) \end{aligned} \quad (8.14)$$

Therefore, by (8.14), we get that: $\mu^*(\mathfrak{m}(\sigma) \cap \mathfrak{m}^c(\tau)) \geq 0$ for every $\sigma \in \Sigma$. By Lemma 8.4 we have that $\mathfrak{m}(\tau) \subseteq \mathfrak{m}(\Sigma) = \bigcup_{\sigma \in \Sigma} \mathfrak{m}(\sigma)$. And since Σ is pairwise disjoint then:

$$\mu^*(\mathfrak{m}(\tau)) = \mu^*\left(\bigcup_{\sigma \in \Sigma} \mathfrak{m}(\sigma) \cap \mathfrak{m}(\tau)\right) \underset{\Sigma \text{ is pairwise disjoint}}{=} = \sum_{\sigma \in \Sigma} \mu^*(\mathfrak{m}(\tau) \cap \mathfrak{m}(\sigma))$$

The result then follows from noting that due to the set-additivity of μ^* we have:

$$h(\Sigma) = \sum_{\sigma \in \Sigma} \mu^*(m(\sigma)) = \underbrace{\sum_{\sigma \in \Sigma} \mu^*(m(\sigma) \cap m(\tau))}_{=\mu^*(m(\tau))=h(\tau) \text{ because } \Sigma \text{ is disjoint}} + \underbrace{\sum_{\sigma \in \Sigma} \mu^*(m(\sigma) \cap m^c(\tau))}_{\geq 0 \text{ (by (8.14))}}. \quad \square$$

8.2. Exact implications in the cone \mathcal{P}_n admit a 1-relaxation. We recall the signed measure μ^* (Theorem 8.2) defined over the field \mathcal{F}_n (Definition 8.1) whose atoms we denote by \mathcal{A} . In particular, μ^* can be negative.

Example 8.5. Consider the parity function on three binary random variables. That is, $Z = X \oplus Y$ where X, Y are uniformly distributed binary variables. We first observe that the variables are pairwise independent. Clearly, X and Y are pairwise independent. It is also easy to see that $P(Z = z) = \frac{1}{2}$ for $z \in \{0, 1\}$. Observe that X and Z are also independent:

$$\begin{aligned} P(X = x, Z = z) &= P(x, z, Y = z \oplus x) + P(x, z, Y \neq z \oplus x) \\ &= P(z|x, Y = z \oplus x)P(x, Y = z \oplus x) + P(z|x, Y \neq z \oplus x)P(x, Y \neq z \oplus x) \\ &= P(z|x, Y = z \oplus x)P(x, Y = z \oplus x) \\ &= 1 \cdot P(x, Y = z \oplus x) \\ &= \frac{1}{4} \\ &= P(x) \cdot P(z) \end{aligned}$$

Further, every variable is functionally determined by the other two (e.g., $X = 0, Z = 0 \Rightarrow Y = 0$). Let h_P denote the entropic function associated with P . Since X and Y are independent, then $I_{h_P}(X, Y) = 0$. On the other hand, given $Z = z$ variables X, Y are not independent. We can see this formally:

$$\begin{aligned} I_{h_P}(X; Y|Z) &= h_P(XZ) + h_P(YZ) - h_P(Z) - h_P(XYZ) \\ &= h_P(X) + h_P(Y) + h_P(Z) - (h_P(Z|XY) + h(XY)) \\ &= h_P(X) + h_P(Y) + h_P(Z) - h_P(XY) \\ &= h_P(Z) \\ &= 2 \cdot \frac{1}{2} \log 2 = \log 2 > 0 \end{aligned} \tag{8.15}$$

We now see how this is related to the i-measure μ^* . By definition of μ^* we have that (see Table 3):

$$\begin{aligned} 0 &= I_{h_P}(X, Y) = \mu^*(m(X) \cap m(Y)) \\ &= \mu^*(m(X) \cap m(Y) \cap m(Z)) + \mu^*(m(X) \cap m(Y) \cap m^c(Z)) \end{aligned}$$

Therefore,

$$\begin{aligned} \mu^*(m(X) \cap m(Y) \cap m(Z)) &= -\mu^*(m(X) \cap m(Y) \cap m^c(Z)) \\ &= -I_{h_P}(X; Y|Z) && \text{Table 3} \\ &= -\log 2 && (8.15) \\ &< 0 \end{aligned}$$

If $\mu^*(a) \geq 0$ for all atoms $a \in \mathcal{A}$ then it is called a *positive measure*. A polymatroid is said to be *positive* if its I-measure is positive, and we denote by Δ_n the cone of positive polymatroids.

Theorem 8.3 states that any I-measure assigning non-negative values to its atoms is a polymatroid. As a corollary, we get, in Lemma 8.4, a set-theoretic characterization of exact implication in the polymatroid cone. The result of Lemma 8.4, combined with the assumption of a non-negative i-measure, gives us a 1-relaxation in the cone of positive polymatroids Δ_n (Theorem 8.6).

Theorem 8.6. *Exact implication in the cone of positive polymatroids admits a 1-relaxation.*

$$\text{If } \Delta_n \models_{\text{EI}} \Sigma \Rightarrow \tau \text{ then } \Delta_n \models h(\Sigma) \geq h(\tau).$$

Proof. By Lemma 8.4 we have that $m(\tau) \subseteq m(\Sigma)$. Hence, $m(\tau) = \cup_{\sigma \in \Sigma} (m(\tau) \cap m(\sigma))$. By the assumption that the I-measure μ^* is non-negative, we get that:

$$\mu^*(m(\tau)) = \mu^*\left(\cup_{\sigma \in \Sigma} (m(\tau) \cap m(\sigma))\right) \leq \sum_{\sigma \in \Sigma} \mu^*(m(\tau) \cap m(\sigma)) \quad (8.16)$$

By the set-additivity of μ^* we get that:

$$h(\Sigma) \stackrel{\text{def}}{=} \sum_{\sigma \in \Sigma} \mu^*(m(\sigma)) = \sum_{\sigma \in \Sigma} \mu^*(m(\sigma) \cap m(\tau)) + \sum_{\sigma \in \Sigma} \mu^*(m(\sigma) \cap m^c(\tau)) \quad (8.17)$$

From (8.17), (8.16), and non-negativity of μ^* we get that:

$$\begin{aligned} \mu^*(m(\tau)) &\leq \sum_{\sigma \in \Sigma} \mu^*(m(\sigma)) - \sum_{\sigma \in \Sigma} \mu^*(m(\sigma) \cap m^c(\tau)) \\ &\leq \sum_{\sigma \in \Sigma} \mu^*(m(\sigma)). \end{aligned}$$

Since, by Theorem 8.2, μ^* is the unique signed measure consistent with all Shannon information measures, the result follows. \square

\mathcal{P}_n coincides with the cone of step functions. We describe the structure of the I-measure of a step function, and prove that the conic hull of step functions $\mathcal{P}_n = \text{conhull}(\mathcal{S}_n)$ and positive polymatroids $\Delta_n \subseteq \Gamma_n$ coincide. That is, $\Delta_n = \mathcal{P}_n$. We let $U \subseteq [n]$, and let s_U denote the step function at U . In the rest of this section we prove Theorem 8.7.

Theorem 8.7. *It holds that $\Delta_n = \mathcal{P}_n$.*

For proving this Theorem, we require Lemma 8.8 that characterizes the i-measure function for step functions. In particular, this lemma shows that step functions are positive polymatroids (i.e., have positive i-measures).

Lemma 8.8. *Let s_U be the step function at $U \subseteq [n]$, and $R = [n] \setminus U$. The unique I-measure for any step function s_U is:*

$$\mu^*(a) = \begin{cases} 1 & \text{if } a = (\cap_{X \in R} m(X)) \cap (\cap_{X \in U} m^c(X)) \\ 0 & \text{otherwise} \end{cases} \quad (8.18)$$

Proof. We define the atom a^* by:

$$a^* \stackrel{\text{def}}{=} (\cap_{X \in R} m(X)) \cap (\cap_{X \in U} m^c(X))$$

Let $S \subseteq U$. We show that $h(S) = \mu^*(m(S)) = 0$. By definition, $m(S) = \cup_{Y \in S} m(Y)$. Therefore, every atom $a \in m(S)$ has at least one set $m(Y)$, where $Y \in S$, that appears in positive form. In particular, $a^* \notin m(S)$. Therefore, by (8.18), we have that $h(S) = \mu^*(m(S)) = 0$ as required.

Now, let $T \not\subseteq U$, and let $X \in T \setminus U$. Since $m(X)$ appears in positive form in a^* , and since $m(X) \subseteq m(T)$, then we have that $h(T) = 1$. Therefore, μ^* as defined in (8.18) is the I-measure for s_U . By Theorem 8.2, the I-measure μ^* is *unique*, and is therefore the only I-measure for s_U . \square

By Lemma 8.8, every step function has a non-negative I-measure. As a consequence, any conic combination of step functions has a positive I-measure, proving that $\mathcal{P}_n \subseteq \Delta_n$.

Now, let μ^* denote any non-negative I-measure. For every atom a , we let $\mathbf{neg}(a)$ denote the set of variables whose sets appear in negated form (e.g., if $a = (\cap_{X \in R} m(X)) \cap (\cap_{X \in U} m^c(X))$, then $\mathbf{neg}(a) = U$). Let \mathcal{A} denote the atoms of \mathcal{F}_n . We define the mapping $f : \mathcal{A} \mapsto \mathcal{P}_n$ as follows.

$$f(a) \stackrel{\text{def}}{=} \mu^*(a) \cdot s_{\mathbf{neg}(a)}$$

where $s_{\mathbf{neg}(a)}$ is the step function at $\mathbf{neg}(a)$. Since, by Lemma 8.8, the I-measure of $s_{\mathbf{neg}(a)}$ is 1 at atom a and 0 everywhere else we get that:

$$\mu^* = \sum_{a \in \mathcal{A}} f(a) = \sum_{a \in \mathcal{A}} \mu^*(a) \cdot s_{\mathbf{neg}(a)}$$

proving that $\mu^* \in \mathcal{P}_n$, and that $\Delta_n \subseteq \mathcal{P}_n$. This concludes the proof of Theorem 8.7.

9. DISCUSSION AND FUTURE WORK

Number of Repairs. A natural way to measure the degree of a constraint in a relation instance R is by the number of repairs needed to enforce the constraint on R . In the case of a key constraint, $X \rightarrow Y$, where $XY = \Omega$, our information-theoretic measure is naturally related to the number of repairs, as follows. If $h(Y|X) = c$, where h is the entropy of the empirical distribution on R , then one can check $|R|/|\Pi_X(R)| \leq 2^c$. Thus, the number of repairs $|R| - |\Pi_X(R)|$ is at most $(2^c - 1)|\Pi_X(R)|$. We leave for future work an exploration of the connections between the number of repairs and information theoretic measures.

Small Model Property. We have proven in Sec. 4 that several classes of implications (including saturated CIs, FDs, and MVDs) have a “small model” property: if the implication holds for all uniform, 2-tuple distributions, then it holds in general. In other words, it suffices to check the implication on the step functions \mathcal{S}_n . One question is whether this small model property continues to hold for other tractable classes of implications in the literature. For example, Geiger and Pearl [GP93] give an axiomatization (and, hence, a decision procedure) for *marginal* CIs. However, marginal CIs do not have the same small model property. Indeed, the implication $(X \perp Y), (X \perp Z) \Rightarrow (X \perp YZ)$ holds for all uniform 2-tuple distributions (because $I_h(X; YZ) \leq I_h(X; Y) + I_h(X; Z)$ holds for all step functions), however it fails for the “parity distribution” in Fig.1(b). We leave for future work an investigation of the small model property for other classes of constraints.

Bounds on the Factor λ . In the early stages of this work we conjectured that all CIs in Γ_n admit a 1-relaxation, until we discovered the counterexample in Th. 6.3, where $\lambda = 3$. On the other hand, the only general upper bound is $(2^n)!$. None of them is likely to be tight. We leave for future work the task of finding tighter bounds for λ .

REFERENCES

- [AD80] William Ward Armstrong and Claude Delobel. Decomposition and functional dependencies in relations. *ACM Trans. Database Syst.*, 5(4):404–430, 1980. doi:10.1145/320610.320620.
- [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. doi:10.1145/2983323.2983781.
- [Bee80] Catriel Beeri. On the membership problem for functional and multivalued dependencies in relational databases. *ACM Trans. Database Syst.*, 5(3):241–259, September 1980. URL: <http://doi.acm.org/10.1145/320613.320614>, doi:10.1145/320613.320614.
- [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. doi:10.1145/509404.509414.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. doi:10.1017/CB09780511804441.
- [Cha11] Terence Chan. Recent progresses in characterising information inequalities. *Entropy*, 13(2):379–401, 2011. URL: <http://www.mdpi.com/1099-4300/13/2/379>, doi:10.3390/e13020379.
- [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. doi:10.1109/ICDE.2014.6816746.
- [CP01] Toon Calders and Jan Paredaens. Axiomatization of frequent sets. In Jan Van den Bussche and Victor Vianu, editors, *Database Theory — ICDT 2001*, pages 204–218, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
- [DFZ09] Randall Dougherty, Christopher F. Freiling, and Kenneth Zeger. Linear rank inequalities on five or more variables. *CoRR*, abs/0910.0284, 2009. URL: <http://arxiv.org/abs/0910.0284>, arXiv:0910.0284.
- [DR00] Mehmet M. Dalkilic and Edward L. Roberston. Information dependencies. In *Proceedings of the Nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '00*, pages 245–253, New York, NY, USA, 2000. ACM. URL: <http://doi.acm.org/10.1145/335168.336059>, doi:10.1145/335168.336059.
- [Fag82] Ronald Fagin. Horn clauses and database dependencies. *J. ACM*, 29(4):952–985, 1982. doi:10.1145/322344.322347.
- [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. URL: <http://proceedings.mlr.press/v48/friesen16.html>.
- [GNG14] Marc Gyssens, Mathias Niepert, and Dirk Van Gucht. On the completeness of the semigraphoid axioms for deriving arbitrary from saturated conditional independence statements. *Inf. Process. Lett.*, 114(11):628–633, 2014. doi:10.1016/j.ipl.2014.05.010.
- [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. URL: <http://www.jstor.org/stable/2242326>.
- [GR04] Chris Giannella and Edward L. Robertson. On approximation measures for functional dependencies. *Inf. Syst.*, 29(6):483–507, 2004. doi:10.1016/j.is.2003.10.006.

- [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. doi:10.1016/j.ic.2006.09.002.
- [IC15] Ihab F. Ilyas and Xu Chu. Trends in cleaning relational data: Consistency and deduplication. *Foundations and Trends in Databases*, 5(4):281–393, 2015. doi:10.1561/19000000045.
- [Ken21] Batya Kenig. Approximate implication with d-separation. *CoRR*, abs/2105.14463, 2021. URL: <https://arxiv.org/abs/2105.14463>, arXiv:2105.14463.
- [KLV13] Juha Kontinen, Sebastian Link, and Jouko Väänänen. Independence in database relations. In Leonid Libkin, Ulrich Kohlenbach, and Ruy de Queiroz, editors, *Logic, Language, Information, and Computation*, pages 179–193, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [KN18] Sebastian Kruse and Felix Naumann. Efficient discovery of approximate dependencies. *PVLDB*, 11(7):759–772, 2018. URL: <http://www.vldb.org/pvldb/vol11/p759-kruse.pdf>, doi:10.14778/3192965.3192968.
- [KR13] Tarik Kaced and Andrei E. Romashchenko. Conditional information inequalities for entropic and almost entropic points. *IEEE Trans. Information Theory*, 59(11):7149–7167, 2013. doi:10.1109/TIT.2013.2274614.
- [KS20] Batya Kenig and Dan Suciu. Integrity constraints revisited: From exact to approximate implication. In Carsten Lutz and Jean Christoph Jung, editors, *23rd International Conference on Database Theory, ICDT 2020, March 30-April 2, 2020, Copenhagen, Denmark*, volume 155 of *LIPIcs*, pages 18:1–18:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICDT.2020.18.
- [Lee87] 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. doi:10.1109/TSE.1987.232847.
- [Li05] Chen Li, editor. *Proceedings of the Twenty-fourth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 13-15, 2005, Baltimore, Maryland, USA*. ACM, 2005. URL: <http://dl.acm.org/citation.cfm?id=1065167>.
- [Mai83] David Maier. *Theory of Relational Databases*. Computer Science Pr, 1983.
- [Mal86] Francesco M. Malvestuto. Statistical treatment of the information content of a database. *Inf. Syst.*, 11(3):211–223, 1986. doi:10.1016/0306-4379(86)90029-3.
- [Mat07] F. Matúš. Infinitely many information inequalities. In *2007 IEEE International Symposium on Information Theory*, pages 41–44, June 2007. doi:10.1109/ISIT.2007.4557201.
- [MMRV02] Konstantin Makarychev, Yury Makarychev, Andrei E. Romashchenko, and Nikolai K. Vereshchagin. A new class of non-Shannon-type inequalities for entropies. *Commun. Inf. Syst.*, 2(2):147–166, 2002. doi:10.4310/cis.2002.v2.n2.a3.
- [MVM⁺18] Alejandro Molina, Antonio Vergari, Nicola Di Mauro, Sriraam Natarajan, Floriana Esposito, and Kristian Kersting. Mixed sum-product networks: A deep architecture for hybrid domains. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 3828–3835, 2018. URL: <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16865>.
- [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. URL: https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=2194&proceeding_id=27.
- [PE08] Jean-Philippe Pellet and André Elisseeff. Using markov blankets for causal structure learning. *Journal of Machine Learning Research*, 9:1295–1342, 2008. URL: <http://doi.acm.org/10.1145/1390681.1442776>, doi:10.1145/1390681.1442776.
- [PP86] Judea Pearl and Azaria Paz. Graphoids: Graph-based logic for reasoning about relevance relations or when would x tell you more about y if you already know z? In *ECAI*, pages 357–363, 1986.
- [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. URL: <http://dl.acm.org/citation.cfm?id=1164187>.

- [Sch03] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
- [SDPF81] Yehoshua Sagiv, Claude Delobel, D. Scott Parker, Jr., and Ronald Fagin. An equivalence between relational database dependencies and a fragment of propositional logic. *J. ACM*, 28(3):435–453, July 1981. URL: <http://doi.acm.org/10.1145/322261.322263>, doi:10.1145/322261.322263.
- [SGG08] Bassem Sayrafi, Dirk Van Gucht, and Marc Gyssens. The implication problem for measure-based constraints. *Inf. Syst.*, 33(2):221–239, 2008. doi:10.1016/j.is.2007.07.005.
- [SGS00] Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT press, 2000.
- [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. URL: <http://doi.acm.org/10.1145/3183713.3196914>, doi:10.1145/3183713.3196914.
- [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.
- [Stu93] Milan Studený. Convex cones in finite-dimensional real vector spaces. *Kybernetika*, 29(2):180–200, 1993. URL: <http://www.kybernetika.cz/content/1993/2/180>.
- [SVG05] Bassem Sayrafi and Dirk Van Gucht. Differential constraints. In *Proceedings of the Twenty-fourth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '05*, pages 348–357, New York, NY, USA, 2005. ACM. URL: <http://doi.acm.org/10.1145/1065167.1065213>, doi:10.1145/1065167.1065213.
- [WBW00] S. K. Michael Wong, Cory J. Butz, and Dan Wu. On the implication problem for probabilistic conditional independency. *IEEE Trans. Systems, Man, and Cybernetics, Part A*, 30(6):785–805, 2000. doi:10.1109/3468.895901.
- [Yeu91] Raymond W. Yeung. A new outlook of Shannon’s information measures. *IEEE Trans. Information Theory*, 37(3):466–474, 1991. doi:10.1109/18.79902.
- [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. doi:10.1109/18.641561.
- [ZY98] Zhen Zhang and Raymond W. Yeung. On characterization of entropy function via information inequalities. *IEEE Trans. Information Theory*, 44(4):1440–1452, 1998. doi:10.1109/18.681320.