Consistency, Acyclicity, and Positive Semirings

Albert Atserias
UPC Barcelona

Phokion G. Kolaitis
UC Santa Cruz & IBM Research

Local Consistency vs. Global Consistency

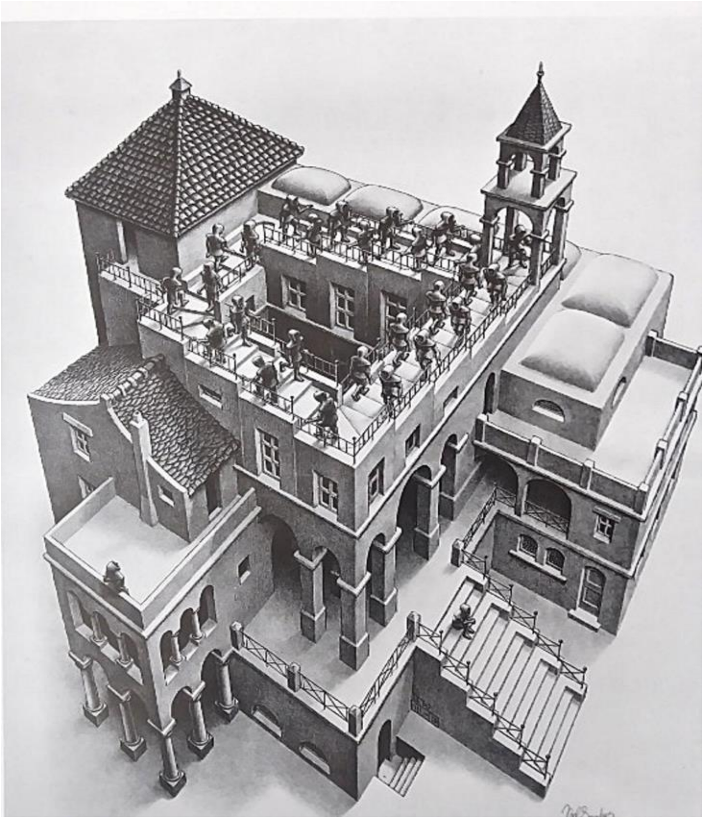

Local Consistency vs. Global Consistency

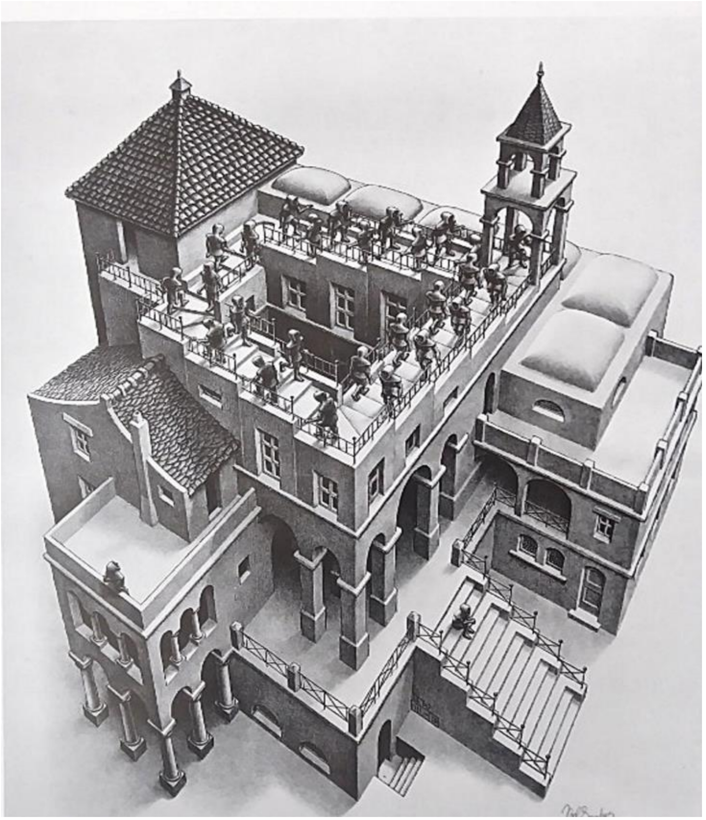

Locally Consistent
Globally Inconsistent
MC Escher

Local Consistency vs. Global Consistency

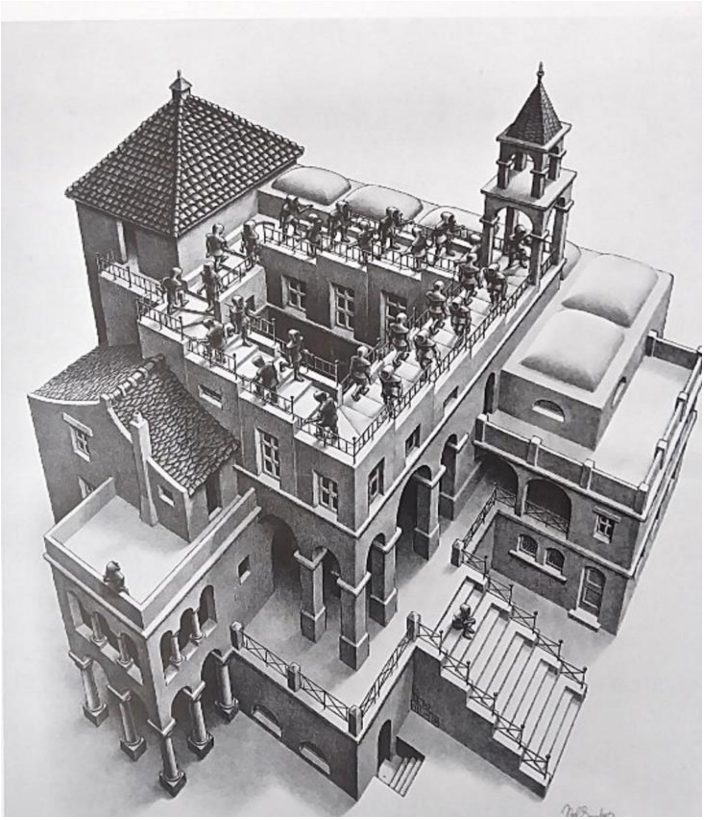

Locally Consistent
Globally Inconsistent
MC Escher

Locally Consistent
Globally Consistent
Santorini, Greece

Local Consistency vs. Global Consistency

Fact:

- In several different settings, the objects of study are "locally consistent" but they may or may not be "globally consistent". Such settings include:
  - Quantum Mechanics, Probability Theory,
  - Constraint Satisfaction, Database Theory, …

Local Consistency vs. Global Consistency

Fact:

- In several different settings, the objects of study are "locally consistent" but they may or may not be "globally consistent". Such settings include:
  - Quantum Mechanics, Probability Theory,
  - Constraint Satisfaction, Database Theory, …

Research Program:

Study the structural aspects of global consistency

- Can we unveil the "intelligible structure" of global consistency?
- When is local consistency equivalent to global consistency?

Local Consistency vs. Global Consistency

Earlier Work:

- Vorob'ev – 1962

  Characterized when a family of probability distributions defined on overlapping sets of variables has a joint distribution.
- Beeri, Fagin, Maier, Yannakakis – 1983

  Characterized when a family of database relations with overlapping sets of attributes has a universal relation.

Local Consistency vs. Global Consistency

Earlier Work:

- Vorob'ev – 1962

  Characterized when a family of probability distributions defined on overlapping sets of variables has a joint distribution.

- Beeri, Fagin, Maier, Yannakakis – 1983

  Characterized when a family of database relations with overlapping sets of attributes has a universal relation.

Goal of this work:

- A common generalization of the results of Vorob'ev and of Beeri et al.
- A unifying framework for studying local vs. global consistency that uses **K**-relations, where **K** is a positive semiring.

Local vs. Global Consistency in Databases

Basic Concepts:

- Attribute: a symbol A with an associated set dom(A) of values.
- R(X): relation R with X as its set of attributes (names of columns)
- R[Z] with $Z \subseteq X$: the projection of R on Z

Example:

R(A,B,C)

A	B	C
1	2	3
1	2	5
2	4	6

R[A,B]

A	B
1	2
2	4

Local vs. Global Consistency in Databases

Definition:

- Two relations $R(X)$ and $S(Y)$ are consistent if there is a relation T over $X \cup Y$ such that $T[X] = R$ and $T[Y] = S$.
- $R_1(X_1), \ldots, R_n(X_n)$ are globally consistent if there is a relation T over $X_1 \cup \cdots \cup X_n$ such that $T[X_1] = R_1, \ldots, T[X_n] = R_n$.

Local vs. Global Consistency in Databases

Definition:

- Two relations R(X) and S(Y) are consistent if there is a relation T over $X \cup Y$ such that $T[X] = R$ and $T[Y] = S$.
- $R_1(X_1), \ldots, R_n(X_n)$ are globally consistent if there is a relation T over $X_1 \cup \cdots \cup X_n$ such that $T[X_1] = R_1, \ldots, T[X_n] = R_n$.

Basic Facts:

- If $R_1(X_1), \ldots, R_n(X_n)$ are globally consistent, then they are pairwise consistent, i.e., R_i and R_j are consistent for all i and j.
- The converse is not always true, i.e., there are relations $R_1(X_1), \ldots, R_n(X_n)$ that are pairwise consistent but are not globally consistent.

Hardy’s Paradox

$R(A_1,B_1)$

A_1	B_1
0	0
0	1
1	0
1	1

$S(A_1, B_2)$

A_1	B_2
0	1
1	0
1	1

$T(A_2, B_1)$

A_2	B_1
0	1
1	0
1	1

$U(A_2, B_2)$

A_2	B_2
0	0
0	1
1	0

Local-to-Global Consistency for Relations

Definition: Schema $H = (X_1, \ldots, X_n)$ of sets of attributes.

H has the local-to-global consistency property for relations if every pairwise consistent collection $R_1(X_1), \ldots, R_n(X_n)$ of relations is globally consistent.

Local-to-Global Consistency for Relations

Definition: Schema $H = (X_1, \ldots, X_n)$ of sets of attributes.

H has the local-to-global consistency property for relations if every pairwise consistent collection $R_1(X_1), \ldots, R_n(X_n)$ of relations is globally consistent.

Example 1: The schema

$$H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\})$$

has the local-to-global consistency property for relations.

Local-to-Global Consistency for Relations

Definition: Schema $H = (X_1, \ldots, X_n)$ of sets of attributes.

H has the local-to-global consistency property for relations if every pairwise consistent collection $R_1(X_1), \ldots, R_n(X_n)$ of relations is globally consistent.

Example 1: The schema

$$H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\})$$

has the local-to-global consistency property for relations.

Example 2: The schema

$$H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\}, \{A_4,A_1\})$$

does not have the local-to-global consistency property for relations.

Local-to-Global Consistency for Relations

Definition: Schema $H = (X_1, \ldots, X_n)$ of sets of attributes.

H has the local-to-global consistency property for relations if every pairwise consistent collection $R_1(X_1), \ldots, R_n(X_n)$ of relations is globally consistent.

Theorem (Beeri, Fagin, Maier, Yannakakis – 1983):

The following are equivalent for a schema $H = (X_1, \ldots, X_n)$

1. H is an acyclic hypergraph.
2. H is a conformal and chordal hypergraph.
3. H has the running intersection property.
4. H has a join tree.
5. H has the local-to-global consistency property for relations.

Conformal and Chordal Hypergraphs

Definition: Let H be a hypergraph

- The primal graph of H is the undirected graph whose edges are pairs of nodes that appear together in at least one hyperedge of H.
- H is conformal if every clique of the primal graph of H is contained in some hyperedge of H.
- H is chordal if its primal graph is chordal (i.e., every cycle of length at least four of has a chord).

Conformal and Chordal Hypergraphs

Definition: Let H be a hypergraph

- The primal graph of H is the undirected graph whose edges are pairs of nodes that appear together in at least one hyperedge of H.
- H is conformal if every clique of the primal graph of H is contained in some hyperedge of H.
- H is chordal if its primal graph is chordal (i.e., every cycle of length at least four of has a chord).

Example 1: H = ({A_1,A_2}, {A_2,A_3}, {A_3,A_4}) is conformal and chordal.

Conformal and Chordal Hypergraphs

Definition: Let H be a hypergraph

- The primal graph of H is the undirected graph whose edges are pairs of nodes that appear together in at least one hyperedge of H.
- H is conformal if every clique of the primal graph of H is contained in some hyperedge of H.
- H is chordal if its primal graph is chordal (i.e., every cycle of length at least four of has a chord).

Example 1: H = $(\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\})$ is conformal and chordal.

Example 2: H = $(\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\}, \{A_4,A_1\})$ is conformal but not chordal.

Conformal and Chordal Hypergraphs

Definition: Let H be a hypergraph

- The primal graph of H is the undirected graph whose edges are pairs of nodes that appear together in at least one hyperedge of H.
- H is conformal if every clique of the primal graph of H is contained in some hyperedge of H.
- H is chordal if its primal graph is chordal (i.e., every cycle of length at least four of has a chord).

Example 1: $H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\})$ is conformal and chordal.

Example 2: $H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\}, \{A_4,A_1\})$ is conformal but not chordal.

Example 3: $H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_1\})$ is chordal but not conformal.

The Running Intersection Property

Definition: A hypergraph H has the running intersection property if there is an ordering $X_1, \ldots, X_n$ of its hyperedges such that for every $i \leq n$, there is a $j < i$ such that

$$X_i \cap (X_1 \cup \cdots \cup X_{i-1}) \subseteq X_j.$$

The Running Intersection Property

Definition: A hypergraph H has the running intersection property if there is an ordering $X_1, \ldots, X_n$ of its hyperedges such that for every $i \leq n$, there is a $j < i$ such that

$$X_i \cap (X_1 \cup \cdots \cup X_{i-1}) \subseteq X_j.$$

Example 1: $H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\})$ has the running intersection property.

Example 2: $H = (\{A_1,A_2\}, \{A_2,A_3\}, \{A_3,A_4\}, \{A_4,A_1\})$ does not have the running intersection property.

Local-to-Global Consistency for Relations

Definition: Schema $H = (X_1, \ldots, X_n)$ of sets of attributes.

H has the local-to-global consistency property for relations if every pairwise consistent collection $R_1(X_1), \ldots, R_n(X_n)$ of relations is globally consistent.

Theorem (Beeri, Fagin, Maier, Yannakakis – 1983):

The following are equivalent for a schema $H = (X_1, \ldots, X_n)$

1. H is an acyclic hypergraph.
2. H is a conformal and chordal hypergraph.
3. H has the running intersection property.
4. H has a join tree.
5. H has the local-to-global consistency property for relations.

Structural Notions (1–4)

Semantic Notion (5)

Positive Semirings

Definition: A positive semiring is a structure **K** = (K, +, ×, 0, 1) such that

- $+$ and $\times$ are binary operations that are commutative and associative and have 0 and 1 as their identity elements;
- $0 \neq 1$;
- $\times$ distributes over $+$, i.e., $a \times (b + c) = (a \times b) + (a \times c)$, for all a, b, c;
- 0 annihilates K, i.e., $0 \times a = 0$, for all a;
- $a + b = 0$ implies $a = 0$ and $b = 0$, for all a, b (i.e., **K** is plus-positive);
- $a \times b = 0$ implies $a = 0$ or $b = 0$, for all a, b (i.e., **K** has no zero divisors).

Note:

- Plus-positivity ensures that the sum of non-zero elements is non-zero.
- No zero divisors ensures that if a product is 0, then at least one factor is 0.

Positive Semirings Are Everywhere

- Boolean semiring: **B** = ({0 ,1}, $\wedge$, $\vee$, 0, 1)
- Bag semiring: **N** = (N, +, $\times$, 0, 1), where N = {0,1,2, …}
  (SQL semantics of database queries via multisets)
- Non-negative reals: $\mathbf{R}^+$ = ([0, ∞), +, $\times$, 0, 1)
- Tropical semiring **T** = ([0, ∞], min, +, ∞, 0)
  (shortest paths in graphs)
- Viterbi semiring **V** = ([0,1], max, $\times$, 0, 1)
  (confidence scores – isomorphic to **T** via $h(x) = e^{-x}$)
- Fuzzy semiring: **F** = ([0,1], max, min, 0, 1)
  (fuzzy logic semantics)
- Polynomial semiring: **N**[X] = (N[X], +, $\times$, 0, 1) with X a set of variables, N[X] all polynomials with variables from X and coefficients from N
  (database provenance – where the answers come from and how)

K-Relations

- Attribute A with dom(A) as its set of values
- $X = \{A_1, \ldots, A_k\}$ set of attributes
- $Tup(X) = dom(A_1) \times \ldots \times dom(A_k)$

Definition: **K** = (K, +, ×, 0, 1) positive semiring, X be a set of attributes.

A **K**-relation over X is a function R: Tup(X) → K having finite support R', i.e.,

$$R' = \{ t \in Tup(X): R(t) \neq 0 \} \text{ is finite.}$$

K-Relations

- Attribute A with dom(A) as its set of values
- $X = \{A_1, \ldots, A_k\}$ set of attributes
- $Tup(X) = dom(A_1) \times \ldots \times dom(A_k)$

Definition: $\mathbf{K} = (K, +, \times, 0, 1)$ positive semiring, X be a set of attributes.

A **K**-relation over X is a function $R: Tup(X) \rightarrow K$ having finite support R', i.e.,

$$R' = \{ t \in Tup(X): R(t) \neq 0 \} \text{ is finite.}$$

Examples:

- Relations are **B**-relations, where $\mathbf{B} = (\{0, 1\}, \wedge, \vee, 0, 1)$.
- Bags are **N**-relations, where $\mathbf{N} = (N, +, \times, 0, 1)$
  (each tuple has a non-negative integer as multiplicity).
- Probability distributions of finite support are $\mathbf{R^+}$-relations P such that $\sum_{t \in Tup(X)} P(t) = 1$, where $\mathbf{R^+} = ([0, \infty), +, \times, 0, 1)$.

Equivalence of **K**-relations

Definition: Let R(X) and S(Y) be two K-relations.

$R \equiv S$ if there are non-zero elements a,b in K such that $aR = bS$,

where $aR : Tup(X) \rightarrow K$ with $(aR)(t) = a \times R(t)$, and similarly for bS.

Fact: $\equiv$ is an equivalence relation on the collection of all K-relations.

Equivalence of **K**-relations

Definition: Let R(X) and S(Y) be two K-relations.
$R \equiv S$ if there are non-zero elements a,b in K such that $aR = bS$,
where $aR : Tup(X) \rightarrow K$ with $(aR)(t) = a \times R(t)$, and similarly for bS.

Fact: $\equiv$ is an equivalence relation on the collection of all K-relations.

Note:

- If R and S are **B**-relations, then $R \equiv S$ if and only if $R = S$.
- There are **N**-relations (bags) R(X) and S(Y) such that $R \equiv S$ but $R \neq S$.
- If R and S are probability distributions of finite support, then $R \equiv S$ if and only if $R = S$.
- For every $\mathbf{R}^+$-relation R, there is a probability distribution P of finite support such that $R \equiv P$ (normalize R to get P).

Local vs. Global Consistency for **K**-Relations

Definition: Let $R_1(X_1), \ldots, R_n(X_n)$ be **K**-relations

- $R_1(X_1), \ldots, R_n(X_n)$ are globally consistent if there is a **K**-relation T over $X_1 \cup \cdots \cup X_n$ such that $T[X_1] \equiv R_1, \ldots, T[X_n] \equiv R_n$.

Basic Facts:

- If $R_1(X_1), \ldots, R_n(X_n)$ are globally consistent **K**-relations, then they are pairwise consistent, i.e., R_i and R_j are consistent for all i and j.
- The converse is not always true, i.e., there are relations $R_1(X_1), \ldots, R_n(X_n)$ that are pairwise consistent but are not globally consistent.

Local-to-Global Consistency for **K**-relations

Definition: **K** positive semiring, Schema H = (X_1, …, X_n) of sets of attributes.
H has the local-to-global consistency property for **K**-relations if every pairwise consistent collection $R_1(X_1)$, …, $R_n(X_n)$ of **K**-relations is globally consistent.

Main Theorem: Let **K** be a positive semiring.
The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H is a conformal and chordal hypergraph.
3. H has the running intersection property.
4. H has a join tree.
5. H has the local-to-global consistency property for **K**-relations.

Structural Notions (1–4)

Semantic Notion (5)

Local-to-Global Consistency for **K**-relations

Main Theorem: Let **K** be a positive semiring.

The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for **K**-relations.

Local-to-Global Consistency for **K**-relations

Main Theorem: Let **K** be a positive semiring.

The following are equivalent for a schema $H = (X_1, \ldots, X_n)$

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for **K**-relations.

Proof Hint: Different proof architecture than the BFMY Theorem.

Step 1: If H has the running intersection property, then H has the local-to-global consistency property for **K**-relations.

Step 2: If H is not conformal or H is not chordal, then H does not have the local-to-global consistency property for **K**-relations.

Local-to-Global Consistency for **K**-relations

Step 1: If H has the running intersection property, then H has the local-to-global consistency property for **K**-relations.

Proof Outline: Let $X_1, \ldots, X_n$ be an ordering of the hyperedges of H such that for every $i \leq n$, there is a $j < i$ such that $X_i \cap (X_1 \cup \cdots \cup X_{i-1}) \subseteq X_j$.

- Assume that $R_1(X_1), \ldots, R_n(X_n)$ are pairwise consistent **K**-relations. By induction on $i \leq n$, show that $R_1(X_1), \ldots, R_i(X_i)$ are globally consistent.

- Assume that $R_1(X_1), \ldots, R_{i-1}(X_{i-1})$ are globally consistent and let $W(X_1 \cup \ldots \cup X_{i-1})$ be a **K**-relation witnessing their global consistency.

- Define the notion of the join $R \bowtie S$ of two K-relations and show that $W \bowtie R_i$ witnesses the consistency of W and R_i.

Note: The definition of $R \bowtie S$ is rather delicate (and is not the obvious one).

Local-to-Global Consistency for **K**-relations

Step 2: If H is not conformal or H is not chordal, then H does not have the local-to-global consistency property for **K**-relations.

Local-to-Global Consistency for **K**-relations

Step 2: If H is not conformal or H is not chordal, then H does not have the local-to-global consistency property for **K**-relations.

Proof Outline: Show the following intermediate results:

- If H is not conformal or H is not chordal, then H contains a "simple" induced hypergraph H* with hyperedges of one of the forms:
  - $\{ V \setminus A : A \in V \}$, for some set V with $|V| \geq 3$.
  - $\{ \{ A_1, A_2 \}, \ldots, \{ A_{n-1}, A_n \}, \{ A_n, A_1 \} \}$ with $n \geq 4$.
- If H has the local-to-global consistency property for **K**-relations, then so do the above "simple" induced hypergraphs.
- The "simple" induced hypergraphs H* do not have the local-to-global consistency property for **K**-relations.

  Explicit construction of **K**-relations that are pairwise consistent but not globally consistent; inspired by Tseitin's hard-to-prove tautologies.

Local-to-Global Consistency

Main Theorem: Let **K** be a positive semiring.

The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for **K**-relations.

Corollary: The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for relations.
3. H has the local-to-global consistency property for probability distributions of finite support.

Local-to-Global Consistency

Main Theorem: Let **K** be a positive semiring.

The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for **K**-relations.

Corollary: The following are equivalent for a schema H = (X_1, …, X_n)

1. H is an acyclic hypergraph.
2. H has the local-to-global consistency property for relations.
3. H has the local-to-global consistency property for probability distributions of finite support.

Note:

- The equivalence between 1. and 2. is the BFMY result.
- How is the equivalence between 1. and 3. related to Vorob'ev's work?

Vorob'ev's Theorem and Related Work

Vorob'ev's Theorem - 1962:

The following are equivalent for a schema H = (X_1, …, X_n)

- The hypergraph H = (X_1, …, X_n) is regular.
- H has the local-to-global consistency property for probability distributions of finite support.

Note:

- In the paper, we give a direct proof that H is regular iff H is acyclic.

- Thus, Vorob'ev's Theorem and the BFMY Theorem are instances of a single unifying result.

Consistency over Positive Monoids

Observations: Let **K** = (K, +, ×, 0, 1) be a positive semiring.

- The definition of the projection R[Z] uses only addition +
- Multiplication × was used to define
  - the equivalence relation R ≡ S (there are a, b such that aR=bS) and
  - the join operation R ⋈ S.

Consistency over Positive Monoids

Observations: Let **K** = (K, +, $\times$, 0, 1) be a positive semiring.

- The definition of the projection R[Z] uses only addition +
- Multiplication $\times$ is used to define
  - the equivalence relation R $\equiv$ S (there are a, b such that aR=bS) and
  - the join operation R $\bowtie$ S.

Definitions:

- A positive monoid is a commutative monoid **K** = (K, +, 0) such that a + b = 0 implies a = 0 and b = 0, for all a, b $\in$ K.
- Two **K**-relations R(X) and S(Y) are strictly consistent if there is a **K**-relation T(X $\cup$ Y) such that T[X] = R and T[Y] = S.
- Define analogously the notions of strict global consistency property and strict local-to-global consistency property for a hypergraph H.

Consistency over Positive Monoids

Results (work in progress):

- Let **K** = (K, +, 0) be a positive monoid and let H be a hypergraph. If H has the strict local-to-global consistency property for **K**-relations, then H is acyclic.

- There are positive monoids **K** and acyclic hypergraphs H such that H does not have the strict local-to-global consistency property for **K**-relations.

- We characterize the positive monoids **K** for which every acyclic hypergraph H has the strict local-to-global consistency property for **K**-relations.

A new expanded framework for local vs. global consistency

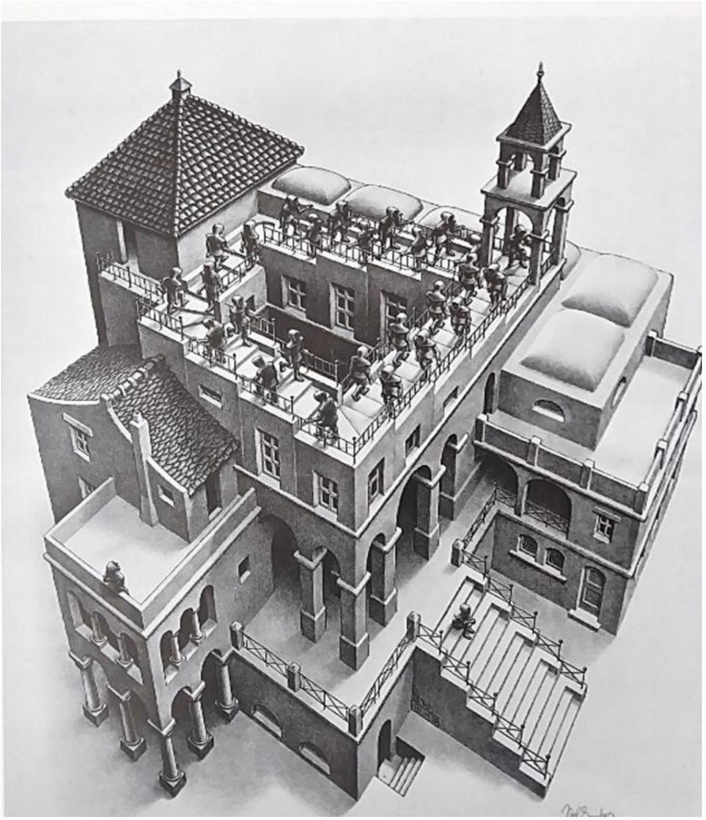

Thank you for your attention!

Backup Slides

The Join of two **K**-relations

If R(X) and S(Y) are **K**-relations, then the join of R and S is the **K**-relation $R \bowtie S$ over $X \cup Y$ such that for every $(X \cup Y)$-tuple t, we have

$$(R \bowtie S)(t) = R(t[X]) \times S(t[Y]) \times \prod_{u \neq t[X \cap Y]} S[X \cap Y](u)$$