Edinburgh Research Explorer

Incomplete information and certain answers in general data models

Incomplete Information and Certain Answers in General Data Models

Leonid Libkin
University of Edinburgh

ABSTRACT

While incomplete information is ubiquitous in all data models – especially in applications involving data translation or integration – our understanding of it is still not completely satisfactory. For example, even such a basic notion as certain answers for XML queries was only introduced recently, and in a way seemingly rather different from relational certain answers.

The goal of this paper is to introduce a general approach to handling incompleteness, and to test its applicability in known data models such as relations and documents. The approach is based on representing degrees of incompleteness via semantics-based orderings on database objects. We use it to both obtain new results on incompleteness and to explain some previously observed phenomena. Specifically we show that certain answers for relational and XML queries are two instances of the same general concept; we describe structural properties behind the naïve evaluation of queries; answer open questions on the existence of certain answers in the XML setting; and show that previously studied ordering-based approaches were only adequate for SQL's primitive view of nulls. We define a general setting that subsumes relations and documents to help us explain in a uniform way how to compute certain answers, and when good solutions can be found in data exchange. We also look at the complexity of common problems related to incompleteness, and generalize several results from relational and XML contexts.

1. Introduction

In most database applications, one has to deal with incomplete data – this fact has been recognized almost immediately after the birth of the relational model. And

yet for a long time the treatment of incomplete information did not receive the attention that it deserved. The design of SQL, for example, is notorious for its nulls-related features, which, as [15] nicely put it, are "in some ways fundamentally at odds with the way the world behaves" (a well-known example of such an oddity is that the statements $|X| > |Y|$ and $X - Y = \emptyset$ are logically consistent in SQL!). On the theoretical side, foundational research from the 1980s, first by Imielinski and Lipski [35, 26] and then by Abiteboul, Kanellakis, and Grahne [3] provided models of incompleteness appropriate for handling queries in different relational languages, and established the computational costs of the key tasks associated with these models. However, until recently, the topic has seen only sporadic activity.

The situation changed rather dramatically over the past several years though, and handling of incompleteness in databases was brought to the fore of database research. This happened mainly due to a number of new applications such as those dealing with data on the web, integrated data, or data that was moved between different applications. In them, the issue of incompleteness is essential. For example, in data integration and data exchange, nulls values and associated concepts such as representation systems and certain answers play a central role [1, 6, 19, 32]. Representation and querying of uncertain data has also been an active field of study [5, 37]; another active direction that borders incompleteness is handling of probabilistic data [27, 28, 40].

Most of the classical theory of incompleteness has been developed in the relational context, where the key concepts (such as, e.g., certain answers) are well understood. A number of attempts were made in the 1990s to develop a general theory of incompleteness in databases, applicable to multiple data models [9, 10, 34, 36, 39]. This was done by using techniques from semantics of programming languages [22]. The connection is rather natural: in programming semantics, a program is assigned a meaning in an ordered set, where the order describes how partial, or incomplete, a function is. In databases, we deal with domains of objects, rather than functions, but these are still ordered based on the degree of incompleteness.

Those general theories, while achieving some success

1

with nested relations [33, 34], fell well short of handling incomplete information in more complex data models, notably XML. We shall elaborate on the reasons for it later in the paper. Incompleteness in XML was addressed recently [4, 7, 16], and one issue that required a complete reworking was the concept of certain answers for queries that return XML documents (i.e., the notion of certain information in a family of trees).

Such certain answers were defined in [16] using an approach that, on the surface looked very different from the relational approach. But does it mean that incompleteness in XML is really different from relational incompleteness? And if it is not – as we shall argue – then what is the general theory that encompasses them all and can be applied to other models?

Our main goal is to develop such a general theory of incompleteness in databases, that subsumes in a uniform way existing relational and XML results, is easily extendable to other data models, an is useful in the main applications of incompleteness.

Our theory will again be based on the notion of an information ordering, to represent the degree of an incompleteness of a database object. Unlike investigations from the 1990s, however, we shall use orderings that are closely tied with the semantics of instances that most commonly arise in applications (more precisely, we deal mainly with naïve tables and their analogs). In view of this, we start by examining two different approaches to handling incompleteness: in relational databases, and in XML documents. We reformulate (and in some cases, strengthen) known results to demonstrate the usefulness of information ordering for query answering.

We then present a very general and datamodel-independent framework for handling incompleteness, that allows us to make two conclusions. First, the notions of certainty in relational and XML contexts, while very different on the surface, are really identical, and correspond to computing greatest lower bounds in ordered sets. Second, we find an easily verified criterion for queries that ensures that certain answers can be computed by *naïve evaluation*, i.e., treating nulls as if they were usual attribute values.

After that, we define the most natural ordering for both relational and XML databases: namely, an object x is less informative than an object y if x denotes more complete objects (indeed, no information means that every instance is possible, so the more objects are possible as the denotation of x, the less informative x is). These orderings can be characterized in terms of the existence of homomorphisms between instances. This immediately allows us to adapt techniques from graph theory, where lattices of graphs and their cores with respect to homomorphism-based orderings have been studied [24] (note, however, that homomorphisms of database instances are not identical to graph homomorphisms, so some work is required in adapting graph-theoretic techniques). We show how to compute greatest lower bounds for finding certain answers, and use this new view of the problem to answer several open problems about the existence of certain answers for XML queries.

To demonstrate that this approach is not limited to relational or XML data, we consider a general setting, reminiscent of data models in [14, 23, 31]. It separates structural and data aspects of a model, which is, essentially, a colored relational structure, where the color of a node determines the length of a tuple of data values attached to it. In relational databases, the underlying structure is just a set; in XML, it is, as expected, a tree.

It turns out that it is often easier to reason about such general models, deriving particular results (say, for XML) as corollaries. For example, the model immediately makes it clear how to compute lower bounds for certain answers – it is the only solution that "type-checks". We also look at upper bounds and explain that they are naturally related to building universal solutions in data exchange.

We conclude by studying computational problems typical in the context of incomplete information. As the underlying structure in our general model conveys some schema information, it gives rise to the consistency problem for incomplete objects. We also look at the membership problem: whether a complete instance represents a possible world for an incomplete one. For this problem, we prove a strong generalization of polynomial-time algorithms for both relational [3] and XML [7] cases under the Codd-interpretation of nulls (when each null occurs at most once). Finally, we look at query answering, and provide broad classes of queries for which upper bounds coincide with those for XML, even for much more general data models.

Organization In Section 2, we review incompleteness in relational and XML models (stating some new results along the way). In Section 3 we present a general model based on orderings. In Section 4 we study homomorphism-based orderings for relations and XML. In Section 5 we present a general data-model that subsumes relations and XML and study incompleteness in that model. In Section 6 we study most common computational problems associated with incompleteness in this general model. Due to space limitations, complete proofs of all the results are in the appendix.

2. Incompleteness in relations and XML

2.1 Incompleteness in relations

We start with a few standard definitions. We assume two disjoint sets of values: $\mathcal{C}$ of constants, and $\mathcal{N}$ of nulls (which will be denoted by $\bot$, with sub/superscripts). A relational schema is a set of relation names with associated arities. An incomplete relational instance associates with each k-ary relation symbol S a k-ary relation over $\mathcal{C} \cup \mathcal{N}$, i.e., a finite subset of $(\mathcal{C} \cup \mathcal{N})^k$. When the

instance is clear from the context we shall write S for the relation itself as well.

Such incomplete relational instances are referred to as *naïve* databases [2, 26]; note that a null $\perp \in \mathcal{N}$ can appear multiple times in such an instance. If each null $\perp \in \mathcal{N}$ appears at most once, we speak of *Codd* databases. If we talk about single relations, it is common to refer to them as naïve tables and Codd tables.

The semantics $[\![D]\!]$ of an incomplete database D is the set of complete databases it can represent. These are defined via *homomorphisms*. We write $\mathcal{C}(D)$ and $\mathcal{N}(D)$ for the sets of constants and nulls, resp., that occur in D. A homomorphism $h : D \to D'$ between two databases of the same schema is a map $h : \mathcal{N}(D) \to \mathcal{C}(D') \cup \mathcal{N}(D')$ such that, for every relation symbol S, if a tuple $\bar{u}$ is in relation S in D, then the tuple $h(\bar{u})$ is in the relation S in D'. As usual, we extend h to constants by letting $h(c) = c$ for every $c \in \mathcal{C}$. The semantics of an incomplete database D, denoted by $[\![D]\!]$ is then defined as the set of complete databases R such that there is a homomorphism $h : D \to R$.

For example, given a naïve table D below, the relation shown next to it is in $[\![D]\!]$, as witnesses by homomorphism $h(\perp_1) = 4$, $h(\perp_2) = 3$, and $h(\perp_3) = 5$:

$$D: \quad \begin{array}{|c|c|c|} \hline 1 & 2 & \perp_1 \\ \hline \perp_2 & \perp_1 & 3 \\ \hline \perp_3 & 5 & 1 \\ \hline \end{array} \qquad R: \quad \begin{array}{|c|c|c|} \hline 1 & 2 & 4 \\ \hline 3 & 4 & 3 \\ \hline 5 & 5 & 1 \\ \hline 3 & 7 & 8 \\ \hline \end{array}$$

Certain answers and naïve evaluation Given an incomplete database D and a query Q, one normally tries to compute *certain answers*:

$$\mathsf{certain}(Q, D) \ = \bigcap \{Q(R) \mid R \in [\![D]\!]\},$$

i.e., answers that are true regardless of the interpretation of nulls. The problem of finding certain answers is undecidable for FO queries (as it becomes finite validity), but for positive relational algebra queries there is a simple polynomial-time algorithm, described below.

Define $Q_{\mathrm{naïve}}(D)$ as follows: first, evaluate Q as if nulls were values (e.g., $\perp_1 = \perp_1$, $\perp_1 \neq \perp_2$, $\perp_1 \neq c$ for every $c \in \mathcal{C}$), and then eliminate tuples with nulls from the result. A classical result from [26] states that if Q is a union of conjunctive queries, then $\mathsf{certain}(Q, D) = Q_{\mathrm{naïve}}(D)$. We can actually show that this is essentially optimal: one cannot find a larger subclass of FO for which naïve evaluation would compute certain answers.

Proposition 1. *If Q is a Boolean FO query and $\mathsf{certain}(Q, D) = Q_{naïve}(D)$ for all naïve databases D, then Q is equivalent to a union of conjunctive queries.*

Certain answers via preorders Note that each naïve database D is naturally viewed as a *Boolean conjunctive query (CQ) Q_D*. For example, the naïve database shown earlier is viewed as a CQ

$$\exists x_1, x_2, x_3 \ D(1, 2, x_1) \wedge D(x_2, x_1, 3) \wedge D(x_3, 5, 1),$$

that is, each null $\perp_i$ is replaced by an existentially quantified variable x_i. Likewise, each Boolean CQ Q can be viewed as a naïve database D_Q (its tableau).

By viewing incomplete databases as logical formulae, we can restate the definition of the semantics in terms of satisfaction of formulae: $R \in [\![D]\!]$ iff $R \models Q_D$.

We now relate certain answers to two standard *preorders* (recall that a preorder is a relation that is reflexive and transitive). The first is what we referred to as the *information ordering* in the introduction: a naïve database D_1 is *less informative* than D_2 if it represents more databases, i.e. $D_1 \preceq D_2$ iff $[\![D_2]\!] \subseteq [\![D_1]\!]$. Notice that this is a preorder rather than a partial order: if $[\![D]\!] = [\![D']\!]$, then both $D \preceq D'$ and $D' \preceq D$ hold.

For queries, we have a well-known preorder: containment (denoted, as usual, by $Q_1 \subseteq Q_2$). The following is essentially a restatement of known results:

Proposition 2. *For a Boolean conjunctive query Q and a naïve database D, the following are equivalent:*

1. $\mathsf{certain}(Q, D) = true$;
2. $D_Q \preceq D$;
3. $Q_D \subseteq Q$.

This also immediately implies that $\mathsf{certain}(Q, D) = \bigwedge \{Q(D') \mid D \preceq D'\}$, where by $Q(D')$ we mean running a query over an incomplete database as if it were a complete database.

These observations indicate that information ordering and lower bounds have a tight connection to finding certain answers; this will be made much more precise later in the paper.

2.2 Incompleteness in XML

Although rather elaborate models of incompleteness for XML exist [4, 7], for our purposes we shall present only a fairly simple model, with nulls occurring as attribute values. This will suffice to explain the main concept used for defining certain answers for XML queries.

The data model is unranked trees, with nodes labeled by letters from a finite alphabet Σ. For each letter $a \in \Sigma$, we have an associated arity $ar(a)$ telling us the number of attributes such nodes carry. More precisely, a tree over Σ is defined as $T = \langle V, E, \lambda, \rho \rangle$, where

- $\langle V, E \rangle$ is a rooted directed unranked tree, with nodes V and the child relation E;
- $\lambda : V \to \Sigma$ is a labeling function; if $\lambda(v) = a$ then we refer to v as an a-labeled node;
- the function ρ assigns to each a-labeled node an $ar(a)$-tuple of data values from $\mathcal{C} \cup \mathcal{N}$.

An example is shown below. Here the alphabet is $\{a, b, c\}$, the arity of a is 2, both b and c have arity 1, and data values are shown next to the labels.

3
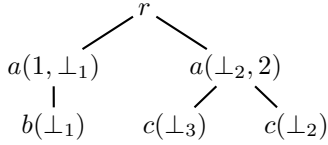
If all the data values in $\rho(v)$ come from $\mathcal{C}$ and the root is labeled by a designated root symbol $r \in \Sigma$, then we refer to a *complete* tree T; otherwise we speak of *incomplete trees.* The semantics of incomplete trees is again defined by means of homomorphisms which act on both tree nodes and data values.

Given a tree T, we use the notation $\mathcal{C}(T)$ for the set of constants in $\mathcal{C}$ that appear as data values in T, and $\mathcal{N}(T)$ for the set of nulls that appear in it. A *homomorphism* $h : T \to T'$ from a tree $T = \langle V, E, \lambda, \rho \rangle$ into a tree $T' = \langle V', E', \lambda', \rho' \rangle$ is a pair of mappings $h_1 : V \to V'$ and $h_2 : \mathcal{N}(T) \to \mathcal{C}(T') \cup \mathcal{N}(T')$ such that:

1. if $(x, y) \in E$ then $(h_1(x), h_1(y)) \in E'$;
2. if x is labeled a in T, then $h_1(x)$ is labeled a in T', i.e., $\lambda'(h_1(x)) = \lambda(x)$; and
3. the data values of $h_1(x)$ are the result of h_2 on the data values of x, i.e., $\rho'(h_1(x)) = h_2(\rho(x))$.

As usual, we extend h_2 to be the identity on constants.

Now the semantics $[\![T]\!]$ of an incomplete tree T is defined as the set of all complete trees T' such that there is a homomorphism $h : T \to T'$.

Certain answers via max-descriptions To define relational certain answers, we used intersection. But how do we define them for trees – what is an analog of intersection then? Most papers addressing incompleteness in XML only looked at queries returning relations, to find a way around this problem.

For proper XML-to-XML queries, the problem was addressed in [16]. If T is an incomplete XML tree, and Q is a query returning XML trees, then certain answers $\mathsf{certain}(Q, T)$ should be the certain information in the set $Q([\![T]\!]) = \{Q(T') \mid T' \in [\![T]\!]\}$. Thus, we need to know how to extract certain information from a collection of XML documents.

For this, it is convenient to use the analogy with naïve tables and CQs, and view trees as both objects and formulae: we say $T \models T'$ iff there is a homomorphism from T' to T (i.e., T satisfies what is described by an incomplete tree T'; this is the analog of $D \models Q_{D'}$ we used above to describe when D is more informative than D'). In XML, this corresponds to describing incompleteness via tree patterns, as done in, e.g., [4, 7, 8].

Our goal now is to extract the certain information from a set $\mathcal{T}$ of incomplete trees. The proposal of [16] was to view incomplete trees as *both* models and formulas, and reconstruct a classical model/formulas Galois connection (cf. [43]). More precisely, for $\mathcal{T}$ we define:

- its theory $\mathrm{Th}(\mathcal{T}) = \{T' \mid \forall T \in \mathcal{T} : T \models T'\}$;
- its models $\mathrm{Mod}(\mathcal{T}) = \{T' \mid \forall T \in \mathcal{T} : T' \models T\}$.

Under this view, $\mathrm{Th}(\mathcal{T})$ describes the certain knowledge conveyed by $\mathcal{T}$. If we want to capture it by an object, we look for a tree T whose models are the models of the certain knowledge, i.e., $\mathrm{Mod}(T) = \mathrm{Mod}(\mathrm{Th}(\mathcal{T}))$. Such trees T were called *max-description* in [16].

Note that all max-descriptions T, T' are equivalent, i.e., $T \models T'$ and $T' \models T$. The definition of certain answers $\mathsf{certain}(Q, T)$ in [16] was simply to take a max-description of $Q([\![T]\!])$. It was shown that this agrees perfectly with the relational definition when XML documents model relations. It was also shown that every set of complete XML documents has a max-description, and every finite set of complete or incomplete XML documents has a max-description. In the finite case, these are computable in time polynomial in the total size of $\mathcal{T}$, and exponential in the number of trees in $\mathcal{T}$.

It was left open however whether arbitrary sets (in fact, even computable sets) of XML documents have max-descriptions. We shall answer this soon. To do so, we first build a general theory (which will completely demystify max-descriptions, among other things), and then apply it to relations, trees, and beyond.

3. Ordered sets and incompleteness

The notions of certain answers for relational and XML queries, on the surface, appear to be very different. But in reality, they are not; in fact, they are the same. To see this, we now describe a very general framework for handling incompleteness and certain answers. In this framework, all that we are going to assume is the presence of an information ordering on a set of database objects. Even with this, we can reason about certain answers, max-descriptions, query answering, and naïve evaluation.

First, some very basic definitions. A *preorder* $\preceq$ is a binary relation that is transitive and reflexive. Associated with the preorder we have an equivalence relation $x \sim y$ defined by $(x \preceq y) \wedge (y \preceq x)$, and the quotient of $\preceq$ by $\sim$ is a partial order. We let $\uparrow x = \{y \mid x \preceq y\}$ and $\downarrow x = \{y \mid y \preceq x\}$. Also $\uparrow X = \bigcup_{x \in X} \uparrow x$ and likewise for $\downarrow X$. A lower bound of a set X is an element y such that $y \preceq x$ for all $x \in X$. A *greatest lower bound* *(glb)* of X is a lower bound y such that $y' \preceq y$ whenever y' is another lower bound of X. It is denoted by $\bigwedge X$. Note that in a preorder, $\bigwedge X$ is an equivalence class wrt $\sim$, but whenever we do not distinguish between equivalent elements we shall write, slightly abusing notation, $y = \bigwedge X$. Glb's need not exist in general.

By a *database domain* we mean a set $\mathcal{D}$ with a preorder $\preceq$ on it. The interpretation is as follows: $\mathcal{D}$ is a set of database objects (say, incomplete relational databases or incomplete trees over the same schema), and $\preceq$ is the

information ordering. That is, we assume that incomplete database objects come with some notion of their semantics $[\![\cdot]\!]$ (what exactly it is, is immaterial to us in this section; this will become important when we look at concrete models), and the interpretation of the pre-order is, as before $x \preceq y$ iff $[\![y]\!] \subseteq [\![x]\!]$. The associated equivalence relation $x \sim y$ is simply $[\![x]\!] = [\![y]\!]$.

Next, we need to define *certain information* in a set $X \subseteq \mathcal{D}$ of objects. If this certain information is represented by an object c, then:

1. c is less informative than every object in X, since it defines information common to all objects in X; that is, c is a lower bound of X;

2. if there are two such objects c, c' and $c' \preceq c$, then we prefer c as it has more information.

Thus, we want the most informative object that defines information common to all objects in X, i.e., a maximal element among all lower bounds of X. That is, *the certain information in the set X is $\bigwedge X$*, the glb of X.

If we look at relational databases without nulls (which arise as elements of $Q([\![D]\!])$), then the semantic ordering for them is $\subseteq$. Thus, for any collection $\mathcal{X}$ of complete databases, we have $\bigwedge \mathcal{X} = \bigcap_{D \in \mathcal{X}} D$ (i.e., in the case of relations certain answers are obtained as intersections).

Certain information defined by max-descriptions does not at first seem to be related to greatest lower bounds; nonetheless, we show that this is precisely what it is.

Max-descriptions deconstructed As we did with naïve databases and XML trees, we can view objects as partial descriptions. Then $x \models y$ is just $y \preceq x$: an object x satisfies every description that is less informative than itself. Therefore, $\mathrm{Mod}(x) = \uparrow x$ and $\mathrm{Th}(x) = \downarrow x$; these are extended to sets, as usual, by $\mathrm{Mod}(X) = \bigcap_{x \in X} \mathrm{Mod}(x)$ and similarly for Th.

The definition of certain information in a set X used in [16] (in the XML context) was a max-description of X: an object x such that $\mathrm{Mod}(x) = \mathrm{Mod}(\mathrm{Th}(X))$.

Theorem 1. *Given a subset X of a database domain $\mathcal{D}$, an element x is a max-description of X iff $x = \bigwedge X$. In particular, a max-description of a set exists iff its greatest lower bound exists.*

So max-descriptions are precisely glb's. The XML case was not different after all: for both relations and XML, the certain information in a set of objects is its glb.

Certain answers for queries We can now use the notion of certain information in a set of objects to define certain answers to queries. An abstract view of a query is as follows. For each schema σ, we have a domain $\mathcal{D}_\sigma$ of database objects of that schema, as well as a preorder $\preceq_\sigma$ on it (if clear from the context which relation $\preceq_\sigma$ we refer to, we write just $\preceq$). A *query* is then a mapping $Q : \mathcal{D}_\sigma \to \mathcal{D}_\tau$ between two domains. If we have a set

$X \subseteq \mathcal{D}_\sigma$ of objects, then $Q(X) = \{Q(x) \mid x \in X\}$, and certain answers to Q over X are defined as

$$\mathsf{certain}(Q, X) = \bigwedge Q(X).$$

There is a simple recipe for finding certain answers when a query Q is *monotone*, i.e., when $x \preceq_\sigma y$ implies $Q(x) \preceq_\tau Q(y)$.

A *basis* of $X \subseteq \mathcal{D}$ is a set $B \subseteq X$ such that $\uparrow B = \uparrow X$. The following simple observation was already made in some concrete cases (like incompleteness in XML [16]):

Lemma 1. *If Q is monotone and B is a basis of X, then $\mathsf{certain}(Q, X) = \bigwedge Q(B)$.*

Hence, if we can find a finite basis $B = \{b_1, \dots, b_n\}$ of X, then we can compute certain answers $\mathsf{certain}(Q, X)$ as $Q(b_1) \wedge \dots \wedge Q(b_n)$. Thus, in such a case it is essential to be able to compute the glb of two (and hence finitely many) elements. The meaning of monotonicity depends on the exact interpretation of the information ordering; this will be addressed in detail in Section 4.

While cases of finite but non-singleton bases do arise in applications [16], most often one computes certain answers over sets X of the form $[\![x]\!]$. In the very general ordered setting, one way of defining the semantics is by $[\![x]\!] = \uparrow x$, i.e., an object represents all objects which are more informative than itself. Then Lemma 1 implies:

Corollary 1. $\mathsf{certain}(Q, \uparrow x) = Q(x)$ *for a monotone Q.*

It says that certain answers can be computed by simply evaluating Q on x. However, in many models the semantics is more refined than just $\bigwedge Q(\uparrow x)$, and is based on *complete* objects that do not have null values. We now study how to incorporate those into our setting.

Complete objects

We extend our setting with the notion of complete objects, i.e., objects without nulls, and explain the properties that lead to naïve evaluation of queries. *Database domains with complete objects* are structures of the form $\langle \mathcal{D}, \preceq, \mathcal{C} \rangle$ where $\preceq$ is a preorder and $\mathcal{C} \subseteq \mathcal{D}$ is a set of objects we view as those having no nulls. To state some basic requirements for these sets, we look at their behavior in the standard models, i.e., in naïve tables.

1. For each object x, the set $\uparrow_{\mathrm{cpl}} x = \uparrow x \cap \mathcal{C}$ of more informative complete objects is not empty (guaranteeing well-defined semantics).

2. For each object x, there is a unique maximal complete object $\pi_{\mathrm{cpl}}(x)$ under x (think of a relation obtained by removing all the rows with nulls from a naïve table). The function π_{cpl} is monotone, and the identity on $\mathcal{C}$. In the standard terminology, this means that $\pi_{\mathrm{cpl}} : \mathcal{D} \to \mathcal{C}$ is a retraction.

3. There are complete objects in $\uparrow_{\mathrm{cpl}} y - \uparrow_{\mathrm{cpl}} x$, unless x is less informative than y. In essence, it says that there are sufficiently many complete objects.

These conditions are satisfied in the standard domains, such as naïve databases or XML documents with nulls. In those domains we define the semantics of x as $\uparrow_{\mathrm{cpl}}x$, the set of complete objects above a given object. It follows immediately from the definition that this notion of the semantics agrees with the ordering.

Lemma 2. *If $\langle \mathcal{D}, \preceq, \mathcal{C} \rangle$ is a database domain with complete objects, then $x \preceq y \Leftrightarrow \uparrow_{\mathrm{cpl}}y \subseteq \uparrow_{\mathrm{cpl}}x$.*

We denote the glb in the restriction $\langle \mathcal{C}, \preceq \rangle$ (when it exists) by $\bigwedge_{\mathrm{cpl}}$. This gives us the notion of certain answers based on complete objects, for a query $Q : \mathcal{D} \to \mathcal{D}'$ that send complete objects to complete objects. It is the glb of the answers to Q over complete objects dominating the input:

$$\mathsf{certain}_{\mathrm{cpl}}(Q, x) \;=\; \bigwedge\nolimits_{\mathrm{cpl}} Q\big(\uparrow_{\mathrm{cpl}}x\big).$$

We say that *certain answers are computed by naïve evaluation* if $\mathsf{certain}_{\mathrm{cpl}}(Q, x) = \pi_{\mathrm{cpl}}(Q(x))$: in other words, they are obtained by running Q and then finding the complete approximation of the result.

For relations, these are of course the familiar concepts, as $\mathsf{certain}_{\mathrm{cpl}}(Q, x)$ is exactly $\bigcap\{Q(R) \mid R \in [\![D]\!]\}$, and $\pi_{\mathrm{cpl}}(Q(D))$ is $Q_{\mathrm{naïve}}(D)$.

To provide a criterion for checking whether certain answers can be computed by naïve evaluation, we say that a function $f : \mathcal{D} \to \mathcal{D}'$ between two database domains with complete objects $\mathcal{C}$ and $\mathcal{C}'$ has the *complete saturation property* if it maps complete objects to complete objects and the following conditions hold:

- if $f(x) \in \mathcal{C}'$ then $f(c) = f(x)$ for some $c \in \uparrow_{\mathrm{cpl}}x$;
- if $f(x) \notin \mathcal{C}'$ and $f(x) \not\preceq c' \in \mathcal{C}'$, then $f(c)$ and c' are incompatible for some $c \in \uparrow_{\mathrm{cpl}}x$.

Intuitively, $\uparrow_{\mathrm{cpl}}x$ has enough complete objects to witness different relationships that involve $f(x)$.

Theorem 2. *For every query that is monotone and has the complete saturation property, certain answers are computed by naïve evaluation.*

Over relational databases, complete saturation is very easy to check for unions of CQs, providing an alternative proof that the naïve evaluation works for them (see Section 4). Note, however, that Theorem 2 is not limited to any language, as it identifies a structural property behind naïve evaluation.

4. Homomorphism-based ordering

We now apply the general theory to relational databases with nulls and to incomplete XML trees. In both cases, the ordering was based on the semantics: $D_1 \preceq D_2$ iff $[\![D_2]\!] \subseteq [\![D_1]\!]$ and likewise for trees. There is a simple way to characterize this ordering:

Proposition 3.   • *If D, D' are naïve databases over the same schema, then $D \preceq D'$ iff there is a homomorphism from D to D'.*

- *If T, T' are XML documents over the same alphabet, then $T \preceq T'$ iff there is a homomorphism from T to T'.*

Hence, we deal with a well-established notion, and the ordering corresponds to a concept studied in fields such as graph theory [24] and constraint satisfaction [30]. This raises two questions: how does it relate to orderings previously studied in connection with incompleteness, and what can we get from known results, in particular in graph theory? We now address those.

Information ordering and other orderings As mentioned in the introduction, orderings have been used to provide semantics of incompleteness in the past [9, 10, 34, 36, 39]. A typical ordering (say, for relational model) would be defined as follows. For two tuples over constants and nulls we let $(a_1, \ldots, a_m) \sqsubseteq (b_1, \ldots, b_m)$ if, for each each i, either a_i is a null, or both a_i and b_i are the same constant (i.e., each null is less informative than each constant). This would be lifted to sets by $X \sqsubseteq^\flat Y \Leftrightarrow \forall x \in X \exists y \in Y : x \sqsubseteq y$.

In general, $\sqsubseteq^\flat$ cannot coincide with $\preceq$, as testing the existence of a homomorphism is NP-complete, and $\sqsubseteq^\flat$ is easily testable in quadratic time. But they do coincide for Codd databases (where nulls cannot be reused).

Proposition 4. *If D and D' are Codd databases, then $D \preceq D'$ iff $D \sqsubseteq^\flat D'$.*

Thus, the orderings used back in the 1990s were well suited for the Codd interpretation of nulls, but not the most commonly used naïve interpretation, which arises in applications such as integration and exchange of data.

The lattice of cores Orderings induced by homomorphisms are well known in graph theory. We look at graphs $G = \langle V, E \rangle$, where V is a set of nodes and E is a set of (directed) edges. We write $G \preceq G'$ if there is a homomorphism $h : G \to G'$, i.e., a map $h : V \to V'$ such that $(x, y) \in E$ implies $(h(x), h(y)) \in E'$. We use the same notation $\preceq$ as before, but it will always be clear from the context which homomorphisms we talk about. Clearly $\preceq$ is a preorder, and the associated equivalence relation $G \sim G'$ is given by $\mathrm{core}(G) = \mathrm{core}(G')$. Recall that the core of a graph G is the smallest subgraph G_0 of G such that there is a homomorphism from G to G_0; the core is unique up to isomorphism [24].

When restricted to cores, $\preceq$ defines a lattice with the glb $\wedge$ and the least upper bound $\vee$ given by [24]

- $G \wedge G' = \mathrm{core}(G \times G')$, and
- $G \vee G' = \mathrm{core}(G \bigsqcup G')$ ($\bigsqcup$ is the disjoint union).

In general, $G \times G'$ is an element of the equivalence class defining $G \wedge G'$, and $G \bigsqcup G'$ is an element of the equivalence class defining $G \vee G'$.

There are many facts known about this lattice; many of them stem from the classical result of Erdős [18] that there are graphs of arbitrarily large chromatic number and odd girth (the minimum length of an odd cycle). Since the former is monotone with respect to $\preceq$ and the latter is antimonotone, this lets one construct arbitrary antichains and dense chains inside $\preceq$. In fact, a very deep result [25] says that every countable partial order can be embedded into $\preceq$.

However, "database homomorphisms" and "graph homomorphisms" are not the same. In naïve databases, homomorphisms are only defined on nulls, but their range is unrestricted. In XML, they are pairs of mappings, one on tree nodes, and the other on nulls. Thus, we cannot use standard graph-theoretic results, but of course we can adapt them.

Lower bounds for naïve databases We have seen that lower bounds are essential for computing certain answers. We now show how to compute the glb of two naïve tables (the construction extends to databases, simply by doing it relation-by-relation). The construction, as suggested by graph-theoretic results, is essentially a product that accounts for the different roles of nulls and constants.

Let R, R' be two naïve tables over the same set of attributes. Take any 1-1 mapping that assigns every pair $(x, y) \in \mathcal{C} \cup \mathcal{N}$ a null $\perp_{xy}$ that does not belong to $\mathcal{N}(R) \cup \mathcal{N}(R')$. As instances R, R' will always be clear from the context, we simply write $\perp_{xy}$ instead of the more formal but cumbersome $\perp_{x,y,R,R'}$. Then, for two tuples $t = (a_1, \ldots, a_m)$ and $t' = (b_1, \ldots, b_m)$, we define

$$t \otimes t' = (c_1, \ldots, c_m): \quad c_i = \begin{cases} a_i & \text{if } a_i = b_i \in \mathcal{C} \\ \perp_{a_i b_i} & \text{otherwise} \end{cases} \quad (1)$$

Proposition 5. *The set $\{t \otimes t' \mid t \in R, \ t' \in R'\}$ is $R \wedge R'$, i.e., a glb of R and R' in the preorder $\preceq$.*

Thus, for every finite collection $\mathcal{X}$ of naïve tables, $\bigwedge \mathcal{X}$ exists. In fact, measuring $\|\mathcal{X}\|$ as the total number of tuples in all relations $R \in \mathcal{X}$, one can easily use the inequality between the arithmetic and geometric means to derive that for $\mathcal{X}$ with n tables, $|\bigwedge \mathcal{X}| \leq (\|\mathcal{X}\|/n)^n$, for $\bigwedge$ constructed in Proposition 5. Results of [16] can also be adapted to show that even $|\text{core}(\bigwedge(\mathcal{X}))|$ is necessarily exponential in the number of relations in $\mathcal{X}$.

What about infinite collections? Recall that in the case of XML, the existence of glb's (called at that time max-descriptions) for arbitrary collections was an open problem. We now show that in general, for arbitrary collections (even recursive ones), glb's need not exist.

Theorem 3. *There is an infinite family $\mathcal{X}$ of naïve tables such that $\bigwedge \mathcal{X}$ does not exist (in fact, there are uncountably many such families). Furthermore, there exist (countably many) recursive families $\mathcal{X}$ of relational databases such that $\bigwedge \mathcal{X}$ does not exist.*

Proof sketch. For the first statement, we use the fact that $\langle \mathbb{Q}, < \rangle$ can be embedded into the lattice of cores. If we look at naïve tables that only contain nulls, we have an embedding of $\langle \mathbb{Q}, < \rangle$ into the ordering $\preceq$. If each set had a glb, the Dedekind-MacNeille completion of $\langle \mathbb{Q}, < \rangle$ would be embeddable into the (countable) set of naïve tables, which is impossible by the cardinality argument (this also shows that the number of sets without a glb is uncountable).

For the second statement, we show that the family of directed cycles whose length is a power of two does not have a glb. □

Lower bounds for XML First, observe that the notion $T \models T'$ used in [16] to define certain answers is precisely $T \preceq T'$ (the existence of homomorphisms) and thus, by Theorem 1, max-descriptions for XML as defined in [16] are the glb's in the ordering $\preceq$. Now by coding naïve tables as XML documents (for each tuple there is a child of the root, whose attribute values are the values in the tuple), and using Theorem 3, we answer the open question from [16].

Corollary 2. *There are recursive collections $\mathcal{X}$ of XML documents (of depth 2) for which $\bigwedge \mathcal{X}$ does not exist.*

We shall have more to say about glb's for XML documents in the next section. For now, we offer one comment on the interaction between glb's and sibling-ordering in XML documents. Note that all the work on computing certain answers in XML was done for unordered documents. It turns out that if we add sibling-ordering, then glb's do not exist even for finite collections, which justifies restricting to unordered documents for handling certain answers to queries.

Proposition 6. *There are ordered XML trees T, T' such that $T \wedge T'$ does not exist.*

Proof sketch. Both T and T' have a root and two children, labeled a and b, but ordered differently. It is easy to show that $T \wedge T'$ does not exist. □

Naïve evaluation under $\preceq$ Theorem 2 stated two sufficient conditions for queries to admit naïve evaluation: monotonicity and complete saturation property. We now see what they mean for relational databases wrt to the semantic ordering $D \preceq D' \Leftrightarrow \llbracket D' \rrbracket \subseteq \llbracket D \rrbracket$. Note that monotonicity wrt this ordering is *not* the usual database notion of monotonicity which uses the partial order $\subseteq$: instead, it corresponds to preservation under homomorphisms. The complete saturation property, it turns out, is very easy to check for unions of CQs.

Proposition 7. *Every union of conjunctive queries is monotone and has the complete saturation property with respect to $\preceq$.*

Proof. Monotonicity wrt $\preceq$ (i.e., preservation under homomorphisms) is well known for unions of CQs. To show complete saturation, let Q be a union of UCQs. If $Q(D)$ does not have nulls, then, for h that maps all

nulls in D into distinct constants different from those in D, by genericity we get $Q(h(D)) = h(Q(D)) = Q(D)$. For the second property, take a complete $R \not\preceq Q(D)$, assume that $Q(D)$ has nulls, and let h be a 1-1 map that sends nulls to constants not present in D and R. Take the complete database $h(D) \succeq D$. By genericity, $Q(h(D)) = h(Q(D)) \not\subseteq R$. If $R \subseteq Q(h(D))$ then again by genericity $R = h^{-1}(R) \subseteq h^{-1}(Q(h(D))) = Q(D)$, a contradiction which proves complete saturation. $\square$

Together with Theorem 2 this gives an alternative proof of the classical result that certain answers for unions conjunctive queries can be obtained by naïve evaluation. Note that Theorem 2 is a general result that is not restricted to FO queries, as it states a structural condition behind naïve evaluation.

A remark about CWA We worked with the Open World Assumption (OWA) here: the ordering $\preceq$ is defined as the existence of an *into* homomorphism. The Closed World Assumption corresponds to the existence of an *onto* homomorphism $h : D \to D'$; in this case we write $D \preceq_{\text{CWA}} D'$. We now offer one comment on the relationship between this ordering, and orderings corresponding to CWA that were considered in the past. Typically (see [9, 34]), to model CWA, one used the Plotkin ordering [22] for sets, defined as $X \sqsubseteq^\natural Y$ iff $\forall x \in X \exists y \in Y : x \sqsubseteq y$ and $\forall y \in Y \exists x \in X : x \sqsubseteq y$, to lift the ordering on tuples to the ordering on sets. We saw that $\sqsubseteq^\flat$ coincides with $\preceq$ over Codd databases. Will $\preceq_{\text{CWA}}$ coincide with $\sqsubseteq^\natural$ under the same restriction?

It turns out that they do not coincide, but they are very close. Recall that a relation $S \subseteq A \times B$ satisfies Hall's condition iff $|S(U)| \geq |U|$ for every $U \subseteq A$ (this is the requirement for the existence of a perfect matching).

Proposition 8. *Over Codd databases, $D \preceq_{\text{CWA}} D'$ iff $D \sqsubseteq^\natural D'$ and $\sqsubseteq^{-1}$ satisfies Hall's condition.*

5. Incompleteness in general data models

To further understand incompleteness in relational, XML, and potentially other models, we need to provide algorithms for computing glb's for certain answers, and to understand the basic computational properties associated with incompleteness (i.e., membership, query answering, consistency). The most general setting provided in Section 3 is too general for reasoning about such concrete tasks, and especially their complexity, and in Section 4 we handled relational and XML cases separately, as they had usually been handled in the literature. But we show now that this artificial separation is not necessary and we can derive many results simultaneously for a variety of data models.

Of course looking for general data models subsuming many others is nothing new in database theory, and several have been proposed (e.g., [14, 23, 31]). What we do here is similar in spirit (and in fact closest to [14]), and the model is well-suited to talk about relations in XML in one common framework.

The basic idea is simple: databases and documents have a structural part (e.g., trees for XML) and data part (tuples of attributes attached to tree nodes). The number of attributes attached to a node is determined by the label of that node. This model also subsumes relations, as the simplest case: the structural part is just a set, as we are about to see.

5.1 Generalized databases and information ordering

We now formalize this as follows. A *generalized schema* $\mathbb{S} = \langle \Sigma, \sigma, ar \rangle$ consists of a finite alphabet Σ, a relational vocabulary σ, and a function $ar : \Sigma \to \mathbb{N}$. To define databases, we assume sets $\mathcal{C}$ of constants and $\mathcal{N}$ of nulls, as before, as well as a set $\mathcal{V}$ of nodes to describe the structural part. A *generalized database* over $\mathbb{S}$ is then $\mathfrak{D} = \langle \mathbf{M}, \lambda, \rho \rangle$ where

- $\mathbf{M}$ is a finite σ-structure over $\mathcal{V}$;
- λ is a labeling of elements of $\mathbf{M}$ in Σ, and
- ρ assigns to each node ν in $\mathbf{M}$ a k-tuple over $\mathcal{C} \cup \mathcal{N}$ where k is the arity of the label of ν, i.e., $ar(\lambda(\nu))$.

Both relational and XML databases are special cases of generalized databases. To code a relational database, let $\sigma = \emptyset$ (i.e., $\mathbf{M}$ is just a set), and let Σ contain relation names, with the same arity as the relations themselves. For example, a relational instance $\{R(1, \perp_1), S(\perp_1, \perp_2, 2)\}$ is represented by a set $\{\nu_1, \nu_2\}$ with $\lambda(\nu_1) = R$ and $\lambda(\nu_2) = S$; the arities of R and S are 2 and 3, and $\rho(\nu_1) = (1, \perp_1)$ and $\rho(\nu_2) = (\perp_1, \perp_2, 2)$. For XML, σ is a vocabulary for unranked trees (there could be several; in the example in Sec. 2 we used only the edge relation, but one can use other axes such as next-sibling), and generalized databases themselves are trees with attribute tuples attached to nodes (again, as shown in the example in Sec. 2).

We shall write $\mathbf{M}_\lambda$ for the colored structure $\langle \mathbf{M}, \lambda \rangle$; technically, it is a structure in the vocabulary σ expanded with unary relations P_a for each $a \in \Sigma$, whose interpretation is the set of nodes labeled a. We also write $\langle \mathbf{M}_\lambda, \rho \rangle$ for a generalized database.

To define the semantics in terms of complete generalized databases (i.e., those not using nulls), we use, as before, homomorphisms between generalized databases. We let $\mathcal{V}(\mathfrak{D})$, $\mathcal{C}(\mathfrak{D})$, and $\mathcal{N}(\mathfrak{D})$ stand for the sets of vertices, constants, and nulls in $\mathfrak{D}$. A *homomorphism* $h : \mathfrak{D} \to \mathfrak{D}'$ between $\mathfrak{D} = \langle \mathbf{M}_\lambda, \rho \rangle$ and $\mathfrak{D}' = \langle \mathbf{M}'_{\lambda'}, \rho' \rangle$ is a pair $h = (h_1, h_2)$ of mappings $h_1 : \mathcal{V}(\mathfrak{D}) \to \mathcal{V}(\mathfrak{D}')$ and $h_2 : \mathcal{N}(\mathfrak{D}) \to \mathcal{C}(\mathfrak{D}') \cup \mathcal{N}(\mathfrak{D}')$ such that:

1. h_1 is a usual homomorphism $\mathbf{M}_\lambda \to \mathbf{M}'_{\lambda'}$;
2. if $\nu \in \mathcal{V}(\mathfrak{D})$ and $\nu' = h_1(\nu)$ then $\rho'(\nu') = h_2(\rho(\nu))$.

As always, we extend h_2 to be the identity on constants. Observe also that in condition 2, ν and ν' have the same

label (since h_1 is a homomorphism) and hence tuples $\rho(\nu)$ and $\rho(\nu')$ have the same length. This notion of homomorphism becomes one of the standard notions when we consider relational databases or XML documents.

We now define $[\![\mathfrak{D}]\!]$, the semantics of $\mathfrak{D}$, as the set of all complete generalized databases $\mathfrak{D}'$ such that there is a homomorphism $h : \mathfrak{D} \to \mathfrak{D}'$. As usual, the information ordering is

$$\mathfrak{D} \preceq \mathfrak{D}' \;\Leftrightarrow\; [\![\mathfrak{D}']\!] \subseteq [\![\mathfrak{D}]\!].$$

We have a standard characterization of it:

Proposition 9. $\mathfrak{D} \preceq \mathfrak{D}'$ *iff there is a homomorphism* $h : \mathfrak{D} \to \mathfrak{D}'$.

We now look at constructions of upper and lower bounds wrt the information ordering. The greatest lower bound can be used for computing certain answers, while upper bounds naturally correspond to constructing target instances in data exchange.

5.2 Greatest lower bounds for certain answers

We have already argued that we need greatest lower bounds, or glb's, to compute certain answers to queries. As glb's of arbitrary sets need not exist in general, we want to compute finite bases of databases and calculate glb's over the results of queries on such finite bases. Lemma 1 tells us that this will give us certain answers for queries which are monotone wrt the information ordering.

However, simply calculating glb's as products is not always going to work. For example, the product of two trees is not a tree, so this will not work in the case of XML. In such a case we need glb's in the subclass of generalized databases for which the structural part is a tree. Generally, if we have a class $\mathcal{K}$ of structures $\mathbf{M}_\lambda$, we need a glb $\wedge_{\mathcal{K}}$ of $\mathcal{K}$-*generalized databases*, i.e., those generalized databases in which the structural part is restricted to be in $\mathcal{K}$.

We now show how to construct such a glb using the minimal possible set of assumptions. In fact, in a way it is easier to do the construction in this general setting, without knowing what the concrete structures in $\mathcal{K}$ are, as the construction is fully guided by the setup: it is the only one that "typechecks". It consists of two steps: first, compute the lower bound for the structural part, and then add data as we did for relations.

Our minimal assumption is that we have a glb $\mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'}$ of structures in class $\mathcal{K}$ (that is, a glb that itself is a member of $\mathcal{K}$). Without knowing anything at all about its structure, we can only conclude that $\mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \preceq \mathbf{M}_\lambda \times \mathbf{M}'_{\lambda'}$, since the latter is a glb without the restriction to $\mathcal{K}$. That is, there is a homomorphism $h : \mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \to \mathbf{M}_\lambda \times \mathbf{M}'_{\lambda'}$. Composing h with first and second projections from $\mathbf{M}_\lambda \times \mathbf{M}'_{\lambda'}$ to $\mathbf{M}_\lambda$ and $\mathbf{M}'_{\lambda'}$ gives us homomorphisms

$$\iota : \mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \to \mathbf{M}_\lambda \;\; \text{and} \;\; \iota' : \mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \to \mathbf{M}'_{\lambda'}.$$

This provides the construction of $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}'$ of a glb of two $\mathcal{K}$-generalized databases. We let

$$\langle \mathbf{M}_\lambda, \rho \rangle \;\wedge_{\mathcal{K}}\; \langle \mathbf{M}'_{\lambda'}, \rho' \rangle \;\stackrel{\text{def}}{=}\; \langle \mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle \quad (2)$$

where $\rho \otimes \rho'(\nu) = \rho(\iota(\nu)) \otimes \rho'(\iota'(\nu))$.

Here $\otimes$ is the merge operation on tuples (1) we used to define relational glb in Section 4 (technically speaking, it depends on $\mathfrak{D}, \mathfrak{D}'$ as the nulls it generates must be outside $\mathcal{N}(\mathfrak{D}) \cup \mathcal{N}(\mathfrak{D}')$, but the instances are always clear from the context). Since ι and ι' are homomorphisms, the colors of $\iota(\nu)$ and $\iota'(\nu)$ are the same, and hence $\rho(\iota(\nu)) \otimes \rho'(\iota'(\nu))$ is well-defined.

Theorem 4. *For every class $\mathcal{K}$ of labeled σ-structures that admits glbs, $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}'$ given by (2) is a glb in the class of $\mathcal{K}$-generalized databases ordered by $\preceq$.*

When $\sigma = \emptyset$, i.e., when we deal with relational databases, this yields precisely the construction for relations in Proposition 5.

When $\mathcal{K}$ is the class of unranked trees (in the vocabulary of the child relations and labels), $\wedge_{\mathcal{K}}$ is precisely the construction shown for computing certain answers in [16]. The construction of the glb for trees themselves is standard, and is done inductively, level by level, by pairing nodes with the same labels.

5.3 Least upper bounds for data exchange

We now cast the well-studied problem of finding solutions in data exchange [6, 19] in our framework, proving a partial explanation why in some cases finding solutions is easy, and in some it is not. We shall present both relational and XML data exchange uniformly, using our notion of generalized databases.

In data exchange, we have a source schema $\mathbb{S}$, a target schema $\mathbb{S}'$, and a schema mapping $\mathbb{M}$ which consists of the rules of the form $\mathfrak{I} \to \mathfrak{I}'$, where $\mathfrak{I}$ and $\mathfrak{I}'$ are instances over $\mathbb{S}$ and $\mathbb{S}'$, respectively. Given a complete instance $\mathfrak{D}$ over $\mathbb{S}$, an instance $\mathfrak{D}'$ over $\mathbb{S}'$ is called a *solution* for $\mathfrak{D}$ if for every rule $\mathfrak{I} \to \mathfrak{I}'$ in $\mathbb{M}$, and every homomorphism $(h_1, h_2) : \mathfrak{I} \to \mathfrak{D}$, there exists a homomorphism $(g_1, g_2) : \mathfrak{I}' \to \mathfrak{D}'$ such that g_2 coincides with h_2 on nulls common to $\mathfrak{I}$ and $\mathfrak{I}'$. A solution $\mathfrak{D}'$ is called *universal* (or $\mathcal{K}$-universal) if there is exists a homomorphism from $\mathfrak{D}'$ to every other solution (or every other solution whose structural part belongs to the class $\mathcal{K}$).

It is very easy to check that in the cases of relations and XML, this abstract view coincides precisely with the standard definitions of [6, 19] (in the case of XML, we deal of course with $\mathcal{K}$-universal solutions when $\mathcal{K}$ is the class of trees). To give a concrete example, consider a relational rule (these are known as st-tgds)

$$S(x, y, u) \to T(x, z), T(z, y).$$

For a source relational database $\mathfrak{D}_S$ with a ternary relation S over $\mathcal{C}$, a target database $\mathfrak{D}_T$ with a binary

relation T over $\mathcal{C}$ and $\mathcal{N}$ is a solution if $\mathfrak{D}_S$ and $\mathfrak{D}_T$ satisfy the sentence $\forall x, y, u \; (S(x, y, u) \rightarrow \exists z \; (T(x, z) \wedge T(z, y)))$. We can view the left and right-hand sides of the rule as databases $\mathfrak{I}_S = \{S(x, y, u)\}$ and $\mathfrak{I}_T = \{T(x, z), T(z, y)\}$ over nulls. Then the rule is satisfied iff for each homomorphism $h : \mathfrak{I}_S \rightarrow \mathfrak{D}_S$ (i.e., a fact $S(a, b, c)$ in $\mathfrak{D}_S$, where $h(x) = a, h(y) = b, h(u) = c$) we have tuples $T(a, v), T(v, b)$ for some value v: in other words, a homomorphism $g : \mathfrak{I}_T \rightarrow \mathfrak{D}_T$ such that $g(z) = v$, and g coincides with h on x and y. Thus, we indeed deal with the standard concepts from data exchange.

We now show that universal solutions are *least upper bounds* (lub's) in the preorder $\preceq$; these will be denoted by $\bigvee$ or $\bigvee_{\mathcal{K}}$ when restricted to $\mathcal{K}$-generalized databases.

We let $\mathrm{Hom}(\cdot, \cdot)$ denote the set of homomorphisms between two instances. For a source instance $\mathfrak{D}$ and a mapping $\mathbb{M}$, define

$$\mathbb{M}(\mathfrak{D}) = \{h_2(\mathfrak{I}') \mid \mathfrak{I} \rightarrow \mathfrak{I}' \in \mathbb{M}, \; (h_1, h_2) \in \mathrm{Hom}(\mathfrak{I}, \mathfrak{D})\}.$$

Intuitively, $(h_1, h_2) \in \mathrm{Hom}(\mathfrak{I}, \mathfrak{D})$ provides an instantiation of a rule in the mapping, and then $\mathbb{M}(\mathfrak{D})$ is the set of single-rule applications of rules in $\mathbb{M}$ to the source $\mathfrak{D}$. Computing those is often the first – and easy – step of building a solution in data exchange. Now we can relate $\mathbb{M}(\mathfrak{D})$ to universal solutions.

Theorem 5. *For a mapping $\mathbb{M}$ and a source $\mathfrak{D}$, the $\mathcal{K}$-universal solutions are precisely the elements of the $\sim$-equivalence class $\bigvee_{\mathcal{K}} \mathbb{M}(\mathfrak{D})$.*

This result partly explains why in some cases computing solutions in data exchange is easy, and in some it is not. Consider, for example, relational databases without putting any restrictions on them. Then lub's exist – as we saw before, they are essentially disjoint unions (technically, disjoint unions after renaming of nulls). Indeed, then $\bigsqcup \mathbb{M}(\mathfrak{D})$ is what is called canonical universal solution in data exchange (without constraints on the target). Furthermore, the canonical representative of the equivalence class $\bigvee \mathbb{M}(\mathfrak{D})$ in this case is the core solution in data exchange, that is, precisely $\mathrm{core}(\bigsqcup \mathbb{M}(\mathfrak{D}))$.

However, if we move to XML, and let $\mathcal{K}$ be the class of all unranked trees, we immediately encounter the following problem:

Proposition 10. *Least upper bounds do not always exist in the restriction of $\preceq$ to labeled unranked trees.*

Hence, in the case of XML (especially with additional schema information) and relations with extra target constraints, one needs to find not an lub (as it may not exist) but rather *some* upper bound. This loss of canonicity in the choice of a solution leads to the necessarily ad-hoc – and varying – choices for solutions, which we see in data exchange literature outside of restricted cases of nicely behaved mappings.

6. Computational problems

We now address the complexity of the key computational problems associated with incompleteness. We do it in our general model from Section 5, and show that this form of reasoning allows us to get some results for both relational and XML cases in a uniform way.

The standard problems to consider are [2, 3, 4, 7, 26]:

- Consistency: does an incomplete database have a completion satisfying some conditions? This problem is commonly considered in the XML context, where schemas are usually more complex. This problem tends to be NP-complete, and in PTIME with suitable restrictions [7].

- Membership: does an incomplete database represent a complete one? It is NP-complete for naïve databases and XML documents, and in PTIME in both cases under the Codd interpretation (although the proofs of these facts in [3, 7] use very different techniques).

- Query answering. Typically one asks whether a given tuple of values is a certain answer. Complexity tends to range, depending on the language, from undecidable (for sufficiently expressive languages to encode the validity problem) to coNP-complete to PTIME [3, 7].

We now look at these problems in our general setting.

Consistency

Let φ be a condition, given by a logical formula, over the structures $\mathbf{M}_\lambda$. The consistency problem is:

PROBLEM:	$\mathrm{CONS}(\varphi)$
INPUT:	a generalized database $\langle \mathbf{M}_\lambda, \rho \rangle$
QUESTION:	is there $\langle \mathbf{M}'_{\lambda'}, \rho' \rangle \in [\![\langle \mathbf{M}_\lambda, \rho \rangle]\!]$ such that $\mathbf{M}'_{\lambda'} \models \varphi$?

In general we should avoid undecidable classes of formulae defining structural conditions; indeed, since every generalized database belongs to $[\![\emptyset]\!]$, the consistency problem checks, in particular, satisfiability of sentences. Even though we only deal with data complexity (φ is fixed), we still prefer to avoid problems whose combined complexity is undecidable.

As a sample result on the complexity of the consistency problem, we consider the well-known decidable case of $\exists^* \forall^*$-formulas (i.e., the Bernays-Schönfinkel class, consisting of formulae of the form $\exists x_1, \ldots, x_k \forall y_1, \ldots, y_m \alpha(\bar{x}, \bar{y})$, where α is quantifier-free), and classify the complexity of $\mathrm{CONS}(\varphi)$ based on the exact shape of the quantifier prefix.

Proposition 11.   • *If φ is an $\exists^* \forall^*$ sentence, then $\mathrm{CONS}(\varphi)$ is in NP.*

- *There is a $\exists^*\forall$ sentence φ_0 such that $\mathrm{CONS}(\varphi_0)$ is NP-complete.*
- *If φ is an $\exists^*$ sentence, then $\mathrm{CONS}(\varphi)$ is in PTIME.*

Notice that we consider *data* complexity, i.e., the sentence φ is fixed, so there is no contradiction with the known higher complexity for the satisfiability problem for the Bernays-Schönfinkel class.

Membership

The membership question is whether $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$ for a complete database $\mathfrak{D}'$ and an incomplete database $\mathfrak{D}$. More generally, we can ask whether $\mathfrak{D} \preceq \mathfrak{D}'$ (which is the membership problem when $\mathfrak{D}'$ has no nulls).

Since checking whether $\mathfrak{D} \preceq \mathfrak{D}'$ amounts to checking the existence of a homomorphism from $\mathfrak{D}$ to $\mathfrak{D}'$, we deal with the general constraint-satisfaction problem [30]. Such a problem is in NP, and often NP-complete, even for a fixed $\mathfrak{D}'$. One case that is solvable in PTIME in both relational and XML contexts is the case of the Codd interpretation of nulls. The proofs of these results are very different however: for relational Codd tables, [3] reduced the problem to finding matchings in bipartite graphs, and for XML, [7] used an analog of CTL-model-checking algorithms on finite Kripke structures.

Now we provide a uniform explanation. We say that in a generalized database $\langle \mathbf{M}_\lambda, \rho \rangle$, the function ρ has *Codd interpretation* if each null occurs as its value at most once.

Theorem 6. *For each fixed $k > 0$, checking whether $\langle \mathbf{M}_\lambda, \rho \rangle \preceq \langle \mathbf{M}'_{\lambda'}, \rho' \rangle$ can be done in polynomial time if ρ has Codd interpretation, and the treewidth of $\mathbf{M}_\lambda$ is at most k.*

Both relational and XML polynomial-time algorithms for the Codd interpretation of nulls are special cases of Theorem 6 when $k = 1$. This result is not a corollary of the standard results on the tractability of constraint satisfaction problems with bounded treewidth (cf. [20, 30]) due to the presence of data values and special conditions on homomorphisms.

Query answering

To answer questions about the complexity of query answering, we need, of course, a query language for generalized databases. Here we look at a natural analog of FO. For relational databases, we know that query answering is in PTIME for unions of conjunctive queries (i.e., existential positive FO sentences) but undecidable for all of FO [3]. For XML, even for analogs of conjunctive queries on trees [21, 8] the complexity of finding certain answers can be coNP-complete, but this is typically caused by missing structural information or the presence of a schema [7].

Since generalized databases are two-sorted structures, it appears to be natural to consider a two-sorted version of FO. We can however rather easily avoid the cumbersome multi-sorted presentation by considering, for a generalized schema $\mathbb{S} = \langle \Sigma, \sigma, ar \rangle$, the logic $\mathrm{FO}(\mathbb{S}, \sim)$, which is first-order over σ, the labeling predicates, and predicates $=_{ij} (x, y)$ meaning that the ith attribute of x equals the jth attribute of y.

Note that this covers both relational and XML Boolean conjunctive queries (for XML, σ defines the set of axes). So the natural problem we have is the following, where φ is a sentence of $\mathrm{FO}(\mathbb{S}, \sim)$:

PROBLEM:	$\mathrm{QA}(\varphi)$
INPUT:	a generalized database $\mathfrak{D}$
QUESTION:	is $\mathsf{certain}(\varphi, \mathfrak{D}) = \mathrm{true}$?

We now show that all three cases witnessed for relations and XML – tractability, intractability, and undecidability – are possible.

Theorem 7.
- *If φ is an existential positive sentence of $\mathrm{FO}(\mathbb{S}, \sim)$, then $\mathrm{QA}(\varphi)$ is in DLOGSPACE.*
- *If φ is an existential sentence of $\mathrm{FO}(\mathbb{S}, \sim)$, then $\mathrm{QA}(\varphi)$ is in coNP. Moreover, there exists a generalized schema $\mathbb{S}$ and an existential sentence φ_0 of $\mathrm{FO}(\mathbb{S}, \sim)$ so that $\mathrm{QA}(\varphi_0)$ is coNP-complete.*
- *There exists a generalized schema $\mathbb{S}$ and a sentence φ_1 of $\mathrm{FO}(\mathbb{S}, \sim)$ so that $\mathrm{QA}(\varphi_1)$ is undecidable.*

7. Future work

We briefly outline some directions for future work. Many results in this paper are of generic nature, and it would be nice to understand how they work when they are applied for particular classes of structures. For example, Theorem 7 is stated over the class of all generalized databases, and one would like to see how similar results change when we impose restrictions on the structural part (e.g., require it to be a tree).

We have not looked at constraints, but they are known to cause problems in the presence of incompleteness; in particular, they affect complexity and decidability results [12, 13]. We also would like to see if the structural study of constraints imposed on target instances in data exchange will help determine classes for which least upper bounds, and thus universal solutions, exist; to start with, one can attempt to extract such structural conditions from cases when the chase procedure is known to work (e.g. [19, 17]), or from restricted classes of tgds studied in ontological reasoning [11].

The general model we used in Section 5, while inspired by several other graph-based models, is rather close to the model of [14]. The focus of [14] was rather different from ours, but a more detailed study of the connections with that paper may be warranted.

Finally, all the results here (except Proposition 8) are based on the open world assumption. An ordering can naturally be extracted from the closed world semantics as well, and we plan to study it in the future. Going from OWA to CWA is often known to increase the complexity of the main computational tasks though [41].

8. References

[1] S. Abiteboul, O. Duschka. Complexity of answering queries using materialized views. In *PODS 1998*, pages 254–263.

[2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[3] S. Abiteboul, P Kanellakis, and G. Grahne. On the representation and querying of sets of possible worlds. *TCS*, 78(1):158–187, 1991.

[4] S. Abiteboul, L. Segoufin, and V. Vianu. Representing and querying XML with incomplete information. *ACM TODS*, 31(1):208–254, 2006.

[5] L. Antova, C. Koch, D. Olteanu. 10^{10^6} worlds and beyond: efficient representation and processing of incomplete information. *VLDB J.* 18(5): 1021-1040 (2009).

[6] M. Arenas, P. Barceló, L. Libkin, F. Murlak. *Relational and XML Data Exchange*. Morgan & Claypool, 2010.

[7] P. Barceló, L. Libkin, A. Poggi, and C. Sirangelo. XML with incomplete information. *J. ACM* 58(1): 1–62 (2010).

[8] H. Björklund, W. Martens, and T. Schwentick. Conjunctive query containment over trees. In *DBPL'07*, pages 66–80.

[9] P. Buneman, A. Jung, A. Ohori. Using powerdomains to generalize relational databases. *TCS* 91 (1991), 23–55.

[10] P. Buneman, S. Davidson, A. Watters. A semantics for complex objects and approximate answers. *JCSS* 43(1991), 170–218.

[11] A. Calì, G. Gottlob, T. Lukasiewicz. Datalog$^{\pm}$: a unified approach to ontologies and integrity constraints. In *ICDT'10*, pages 14–30.

[12] A. Calì, D. Lembo, and R. Rosati. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *PODS'03*, pages 260–271.

[13] A. Calì, D. Martinenghi. Querying incomplete data over extended ER schemata. *TPLP* 10(3): 291-329 (2010).

[14] S. Cohen and Y. Sagiv. An abstract framework for generating maximal answers to queries. In *ICDT 2005*, pages 129–143.

[15] C. Date and H. Darwin. *A Guide to the SQL Standard*. Addison-Wesley, 1996.

[16] C. David, L. Libkin, F. Murlak. Certain answers for XML queries. In *PODS 2010*, pages 191-202.

[17] A. Deutsch, A. Nash, J. Remmel. The chase revisited. In *PODS'08*, pages 149–158.

[18] P. Erdős. Graph theory and probability. *Canad. J. Math.* 11 (1959), 34–38.

[19] R. Fagin, Ph. Kolaitis, R. Miller, and L. Popa. Data exchange: semantics and query answering. *TCS*, 336(1):89–124, 2005.

[20] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[21] G. Gottlob, C. Koch, and K. Schulz. Conjunctive queries over trees. *J. ACM* 53(2):238-272, 2006.

[22] C. Gunter. *Semantics of Programming Languages*. The MIT Press, 1992.

[23] M. Gyssens, J. Paredaens, J. Van den Bussche, D. Van Gucht. A graph-oriented object database model *IEEE TKDE* 6(4):572–586, 1994.

[24] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.

[25] J. Hubička and J. Nešetřil. Finite paths are universal. *Order* 22(1):21–40, 2005.

[26] T. Imieliński and W. Lipski. Incomplete information in relational databases. *J. ACM*, 31(4):761–791, 1984.

[27] B. Kimelfeld, Y. Sagiv. Modeling and querying probabilistic XML data. *SIGMOD Record*, 37(4):69–77, 2008.

[28] C. Koch. On query algebras for probabilistic databases. *SIGMOD Record* 37(4): 78-85 (2008).

[29] P. Kolaitis and M. Vardi. Conjunctive-query containment and constraint satisfaction. *JCSS* 61(2): 302-332 (2000).

[30] P. Kolaitis and M. Vardi. A logical approach to constraint satisfaction. In *Finite Model Theory and Its Applications*, Springer 2007, pages 339–370.

[31] G. Kuper and M. Vardi. The logical data model. *ACM TODS* 18(3):379–413 (1993).

[32] M. Lenzerini. Data integration: a theoretical perspective. In *PODS'02*, pages 233–246.

[33] M. Levene and G. Loizou. Semantics of null extended nested relations. *ACM TODS* 18 (1992), 414-459.

[34] L. Libkin. A semantics-based approach to design of query languages for partial information. In *Semantics in Databases*, LNCS 1358, 1998, pages 170–208.

[35] W. Lipski. On semantic issues connected with incomplete information in databases. *ACM TODS* 4 (1979), 262–296.

[36] A. Ohori. Semantics of types for database objects. *Theoretical Computer Science* 76 (1990), 53–91.

[37] D. Olteanu, C. Koch, L. Antova. World-set decompositions: expressiveness and efficient algorithms. *TCS* 403 (2008), 265–284.

[38] B. Rossman. Homomorphism preservation theorems. *J. ACM* 55(3): (2008).

[39] B. Rounds. Situation-theoretic aspects of databases. In *Situation Theory and Appl.*, CSLI vol. 26, 1991, pages 229-256.

[40] D. Suciu. Probabilistic databases. *Encyclopedia of Database Systems*, 2009, pages 2150-2155.

[41] M. Vardi. On the integrity of databases with incomplete information. In *PODS'86*, pages 252–266.

[42] M. Vardi. On the complexity of bounded-variable queries. In *PODS'95*, pages 266–276.

[43] W. Wechler. *Universal Algebra for Computer Scientists*. Springer, 1992.

APPENDIX

Proofs

We shall use the following notations throughout the appendix. If h is a homomorphism between two database instances D and D' (i.e., a mapping that sends nulls to constants or nulls) we shall write $D \xrightarrow{h} D'$. When we deal with XML trees or generalized databases where homomorphisms are pairs of mappings (h_1, h_2), the first one on nodes and the second one on nulls, we shall write $\mathfrak{D} \xrightarrow{(h_1, h_2)} \mathfrak{D}'$. When we deal with the usual graph-theoretic homomorphisms (i.e., nodes are sent to nodes), we write $\mathbf{M} \xmapsto{g} \mathbf{M}'$ to indicate that g is such a homomorphism between structures $\mathbf{M}$ and $\mathbf{M}'$.

Proof of Proposition 1

Note that for a Boolean FO query Q (i.e., an FO sentence) $\mathsf{certain}(Q, D) = Q_{\text{naïve}}(D)$, is equivalent to the following statement being true:

$$D \models Q \;\Leftrightarrow\; \forall D': \; \left(\text{if } D \xrightarrow{h} D' \text{ for } h : \mathcal{N} \to \mathcal{C}, \text{ then } D' \models Q \right) \tag{3}$$

In the next statement we use the graph-theoretic notion of homomorphisms, rather than database homomorphism. That is, such a homomorphism between two structures R and R' of the same vocabulary over $\mathcal{C}$ is a mapping g from the universe of R to the universe of R' that satisfies the usual condition of a homomorphism (thus, the difference is that it does not need to be the identity on constants, as database homomorphism). Recall that for such homomorphisms we write $R \xmapsto{g} R'$.

Next, recall that an FO query Q is preserved under homomorphisms (in the finite) if for every finite structure R over $\mathcal{C}$, we have

$$R \models Q \text{ and } R \xmapsto{g} R' \text{ imply } R' \models Q. \tag{4}$$

We now show that (3) is true for all naïve databases D iff (4) is true for all databases R over $\mathcal{C}$ (of the same schema, of course). By Rossman's theorem [38], this will imply that Q is equivalent to a union of conjunctive queries.

First assume that (3) is true for all naïve databases D. Consider an arbitrary database R over $\mathcal{C}$ such that $R \models Q$, and let $R \xmapsto{g} R'$. Assume that the universe of R is $\{c_1, \ldots, c_m\}$. Let $\bot_1, \ldots, \bot_m$ be distinct nulls from $\mathcal{N}$, and let f be the map that sends each c_i to $\bot_i$, for $i \le m$. Let $R_\bot = f(R)$. Since f is an isomorphism of first-order structures, we have $R_\bot \models Q$. Now consider a mapping $g \circ f^{-1}$ from $R_\bot$ to R'. This mapping is a (database) homomorphism (since it is only defined on nulls in $R_\bot$). Hence applying (3) to $R_\bot$ and the homomorphism $R_\bot \xrightarrow{f^{-1} \circ g} R'$ we conclude that $R' \models Q$. This proves the first implication.

Now assume (4) is true for all databases R over $\mathcal{C}$. We have to prove that (3) is true for all naïve databases D. This of course means proving both directions in (3). First, assume that for a given D, we have $D' \models Q$ whenever $D \xrightarrow{h} D'$ for some $h : \mathcal{N} \to \mathcal{C}$. Let the domain of D contain nulls $\bot_1, \ldots, \bot_n$ and constants $c_1, \ldots, c_m$. Define $h(\bot_i) = c'_i$, for $i \le n$, where $c'_1, \ldots, c'_n$ are distinct constants different from $c_1, \ldots, c_m$. Let $D' = h(D)$. By the assumption, we have $D \models Q$. But then D' is isomorphic to D as a first-order structure, so we conclude $D \models Q$.

Conversely, let $D \models Q$, where D is a naïve database, and let $D \xrightarrow{h} D'$ for some $h : \mathcal{N} \to \mathcal{C}$. We have to prove $D' \models Q$. Assume again that the domain of D contain nulls $\bot_1, \ldots, \bot_n$ and constants $c_1, \ldots, c_m$, and let $f(\bot_i) = c'_i$, for $i \le n$, where $c'_1, \ldots, c'_n$ are distinct constants different from $c_1, \ldots, c_m$. As before, $f(D) \models Q$. Now $f(D)$ is a database over constants, and $h \circ f^{-1}$ is a homomorphism between $f(D)$ and D', that is, $f(D) \xmapsto{h \circ f^{-1}} D'$. Hence, applying (4), we conclude $D' \models Q$. This concludes the proof of the Proposition. $\qquad\square$

Proof of Proposition 2

$(\mathbf{1}) \Rightarrow (\mathbf{2})$ Assume that $\mathsf{certain}(Q, D)$ evaluates to true, i.e., for every R over $\mathcal{C}$ such that $D \xrightarrow{h} R$, we have $R \models Q$.

To show that $D_Q \preceq D$, assume that we have some R over constants such that $D \xrightarrow{h} R$. Then $R \models Q$ by the

assumption, which means that there is a homomorphism h' such that $D_Q \xrightarrow{h'} R$. This shows $[\![D]\!] \subseteq [\![D_Q]\!]$ and thus $D_Q \preceq D$.

(2) $\Rightarrow$ **(3)** If $D_Q \preceq D$, then by Proposition 3 (proved independently of this one), we have $D_Q \xrightarrow{h} D$. Since these are tableaux of Q and Q_D, resp., we conclude $Q_D \subseteq Q$.

(3) $\Rightarrow$ **(1)** If $Q_D \subseteq Q$, then there is a homomorphism from the tableau of Q to the tableau of Q_D, i.e., $D_Q \xrightarrow{h} D$. Now suppose $D \xrightarrow{g} R$ for some R over constants. Then $D_Q \xrightarrow{g \circ h} R$ which implies $R \models Q$. Since this is true for an arbitrary R over constants, we conclude $\mathsf{certain}(Q, D) = \mathsf{true}$. $\qquad\square$

Proof of Theorem 1

Assume first that $\bigwedge X$ exists. Since $\bigwedge X$ is an equivalence class with respect to $\sim$, and since so is the set of max-descriptions (if it exists), it suffices to show that every element u in $\bigwedge X$ is a max-description of X; it will then follow that glb's and max-descriptions coincide in this case.

Let $LB(X)$ denote the set of all lower bounds of X. Then

$$\mathrm{Mod}(\mathrm{Th}(X)) \;=\; \bigcap_{y \in LB(X)} \uparrow y.$$

Hence, $\mathrm{Mod}(\mathrm{Th}(X)) \subseteq \uparrow u$ since $u \in LB(X)$. Conversely, if y is an arbitrary element of $LB(X)$, then $\uparrow u \subseteq \uparrow y$ and thus $\uparrow u \subseteq (\bigcap_{y \in LB(X)} \uparrow y) = \mathrm{Mod}(\mathrm{Th}(X))$. Hence, $\uparrow u = \mathrm{Mod}(\mathrm{Th}(X))$ and u is a max-description.

For the converse we need a simple observation: since by the properties of Galois connections $X \subseteq \mathrm{Mod}(\mathrm{Th}(X))$, we have that if u is a max-description, then $u \in LB(X)$. In particular, a max-description u, if it exists, must be in $(\bigcap_{y \in LB(X)} \uparrow y) \cap LB(X)$.

Now assume that $\bigwedge X$ does not exist and look at the set $LB(X)$ of all lower bounds of X. Assume that $LB(X) = \emptyset$; then $\mathrm{Th}(X) = \emptyset$ and $\mathrm{Mod}(\emptyset) = \mathcal{D}$. A max-description, if it existed, would have been the minimal element of $\mathcal{D}$, but such an element does not exist, for otherwise it would have been in $LB(X)$. So now suppose $LB(X)$ is not empty. Then either there is no maximal element in it, or there are two or more incomparable maximal elements. In both cases we get $(\bigcap_{y \in LB(X)} \uparrow y) \cap LB(X) = \emptyset$, which, by the preceding observation, implies that there is no max-description. This concludes the proof. $\qquad\square$

Proof of Lemma 1

First we show that $Q(B)$ is a basis of $Q(X)$. Since $B \subseteq X$, we have $\uparrow Q(B) \subseteq \uparrow Q(X)$. Conversely, let y be $\uparrow Q(X)$. There there is $x \in X$ such that $Q(x) \preceq y$. Since B is a basis, we have $b \preceq x$ for some $b \in B$. Hence $Q(b) \preceq Q(x) \preceq y$, showing $y \in \uparrow Q(B)$. Thus, $Q(B)$ is a basis of $Q(X)$, and now it suffices to show that if sets Y and Z are such that $\uparrow Y = \uparrow Z$ then their glb's coincide (when they exist). This follows from an even stronger fact that the sets of lower bounds of Y and Z are the same: if y is a lower bound of Y and $z \in Z$, then $z \in \uparrow Y$, i.e., for some $y' \in Y$ we have $z \succeq y' \succeq y$, proving that y is a lower bound for Z. Due to complete symmetry, the lower bounds of Y and Z coincide, proving the lemma. $\qquad\square$

Proof of Theorem 2

Let $\mathcal{D}$ and $\mathcal{D}'$ be two database domains with complete objects. By convention, we shall refer to sets of complete objects as $\mathcal{C}$ and $\mathcal{C}'$, resp. We shall use the same symbol $\preceq$ for ordering as it is always clear from the context which domain one refers to. We also use the notation $x|y$ when two elements are incomparable, i.e. neither $x \preceq y$ nor $y \preceq x$ is true.

We start with two observations. First, for each $x \in \mathcal{D}$, the glb $\bigwedge_{\mathrm{cpl}} \uparrow_{\mathrm{cpl}} x$ exists and equals $\pi_{\mathrm{cpl}}(x)$. For this, it suffices to show that every lower bound c of $\uparrow_{\mathrm{cpl}} x$ is lower than x wrt $\preceq$. This will suffice since $\pi_{\mathrm{cpl}}(x)$ is the greatest element of $\mathcal{C}$ under x, and since $\pi_{\mathrm{cpl}}(x)$ itself is a lower bound of $\uparrow_{\mathrm{cpl}} x$. So assume now that c is a lower bound of $\uparrow_{\mathrm{cpl}} x$ but $c \npreceq x$. Then there is an element $y \in \uparrow_{\mathrm{cpl}} x$ such that $y \notin \uparrow_{\mathrm{cpl}} c$, or, in other words, $c \npreceq y$. But contradicts the assumption that c is a lower bound of $\uparrow_{\mathrm{cpl}} x$, proving the equation $\bigwedge_{\mathrm{cpl}} \uparrow_{\mathrm{cpl}} x = \pi_{\mathrm{cpl}}(x)$.

We call a function $f : \mathcal{D} \to \mathcal{D}'$ between two database domains with complete objects *naïve-admissible* if it is both monotone and has the complete saturation property. We now show that for every naïve-admissible function f,

14

certain answers can be found by naïve evaluation; that is,

$$\bigwedge\nolimits_{\mathrm{cpl}} f(\uparrow_{\mathrm{cpl}} x) \;=\; \pi_{\mathrm{cpl}}(f(x)) \tag{5}$$

To prove (5), it suffices to show that

$$\bigwedge\nolimits_{\mathrm{cpl}} f(\uparrow_{\mathrm{cpl}} x) \;=\; \bigwedge\nolimits_{\mathrm{cpl}} \uparrow_{\mathrm{cpl}} f(x) \tag{6}$$

since the right-hand side of (6), by the fact proven above, exists, and is equal to $\pi_{\mathrm{cpl}}(f(x))$.

We now prove (6). For this, it suffices to show that $LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x)) = LB_{\mathrm{cpl}}(\uparrow_{\mathrm{cpl}} f(x))$, where $LB_{\mathrm{cpl}}(\cdot)$ refers to the set of lower bounds in the set of complete objects (i.e., $LB_{\mathrm{cpl}}(X) = \{c \mid c$ is complete and $c \preceq x$ for all $x \in X\}$). Since $f(\uparrow_{\mathrm{cpl}} x) \subseteq \uparrow_{\mathrm{cpl}} f(x)$ (as f is monotone and sends complete objects to complete objects), we have $LB_{\mathrm{cpl}}(\uparrow_{\mathrm{cpl}} f(x)) \subseteq LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$.

For the converse, suppose $c \in LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$. First assume that $f(x) \in \mathcal{C}'$. Then, by the complete saturation property, $f(x) = f(c_1)$ for some $c_1 \in \uparrow_{\mathrm{cpl}} x$. This implies $f(x) \in f(\uparrow_{\mathrm{cpl}} x)$ and thus $c \preceq f(x)$. In particular, this implies that $c \preceq c'$ for every $c' \in \uparrow_{\mathrm{cpl}} f(x)$ and thus $c \in LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$.

Hence, it remains to consider the case when $f(x) \notin \mathcal{C}'$ and $c \in LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$. There are two subcases. The first is when $c \preceq f(x)$. In that case clearly c is a lower bound for $\uparrow_{\mathrm{cpl}} f(x)$. The second subcase is when c is not less than $f(x)$, with respect to $\preceq$. In that case, by the complete saturation property, we get $c_1 \succeq x$ so that $c_1 \in \mathcal{C}$ and $f(c_1) \mid c$. But then $f(c_1)$ is a complete object in $f(\uparrow_{\mathrm{cpl}} x)$ which contradicts the assumption that c is in $LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$. Hence, this case is impossible, and thus in all the cases we have $c \in LB_{\mathrm{cpl}}(f(\uparrow_{\mathrm{cpl}} x))$, hence proving (6) and (5). This completes the proof. □

Proof of Proposition 3

We do the proof for the relational case first. Assume we have a homomorphism $D \xrightarrow{h} D'$. Then, if R is a database over constants and we have a homomorphism $D' \xrightarrow{g} R$, then there is a homomorphism $D \xrightarrow{g \circ h} R$, and thus $[\![D']\!] \subseteq [\![D]\!]$.

Conversely, assume that $[\![D']\!] \subseteq [\![D]\!]$. Let the universe of D' contain nulls $\perp_1, \ldots, \perp_n$ and constants $c_1, \ldots, c_m$. Let $c'_1, \ldots, c'_n$ be constants not mentioned in D and D'. Consider a mapping f such that $f(\perp_i) = c'_i$ for $1 \leq i \leq n$, and $f(c_j) = c_j$ for $1 \leq j \leq m$. Then $D' \xrightarrow{f} f(D')$ (in fact, f is an isomorphism between D' and $f(D')$ viewed as first-order structures). Since $f(D')$ contains no nulls, this implies $f(D') \in [\![D']\!]$, and thus $f(D') \in [\![D]\!]$. Therefore, $D \xrightarrow{g} f(D')$ for some homomorphism g. Since f is an isomorphism, this implies that $f^{-1} \circ g$ is a homomorphism from D to D'.

The proof for XML trees (and even more general models considered in Section 5) is the same: the direction from $T \xrightarrow{h} T'$ to $[\![T']\!] \subseteq [\![T]\!]$ only requires the closure of homomorphisms under composition. For the converse, we proceed as above, and define $f' = (id, f)$ to be a homomorphism of trees (i.e., the first component is the identity on tree nodes). The rest of the proof then applies verbatim. □

Proof of Proposition 4

Let D, D' be two Codd tables. Assume that $D \preceq D'$. Then, by Proposition 3, we have a homomorphism $D \xrightarrow{h} D'$. Since for every tuple $\bar{t}$ in D we have $t \sqsubseteq h(t)$, this implies $D \sqsubseteq^\flat D'$.

Conversely, assume that $D \sqsubseteq^\flat D'$. Let $D = \{\bar{t}_1, \ldots, \bar{t}_n\}$ and $D' = \{\bar{s}_1, \ldots, \bar{s}_m\}$. Since $D \sqsubseteq^\flat D'$, there is a function $f : \{1, \ldots, n\} \to \{1, \ldots, m\}$ such that $\bar{t}_i \sqsubseteq \bar{s}_{f(i)}$ for every $i \leq n$. If $\bar{t}_i = (a_1, \ldots, a_k)$ and $\bar{s}_{f(i)} = (b_1, \ldots, b_k)$, then for every $j \leq k$, if a_j is a constant, then $b_j = a_j$. If $a_j \in \mathcal{N}$, then set $h(a_j) = b_j$. Since each null occurs just once in D, this function h is well defined, and by construction it is a homomorphism between D and D'. Thus, $D \preceq D'$, proving the proposition. □

15

Proof of Proposition 5

We shall slightly abuse the notation and write $x \otimes y$ to denote x if $x = y \in \mathcal{C}$ or $\perp_{xy}$ otherwise, and $R \otimes R'$ to denote $\{t \otimes t' \mid t \in R, \ t' \in R'\}$. It will always be clear from the context what the operation applies to.

First, note that the mappings $h(\perp_{xy}) = x$ and $h'(\perp_{xy}) = y$ are homomorphisms $R \otimes R' \xrightarrow{h} R$ and $R \otimes R' \xrightarrow{h'} R'$ and thus $R \otimes R'$ is a lower bound for $\{R, R'\}$.

Conversely, suppose we have a naive table D such that $D \preceq R$ and $D \preceq R'$, i.e., there are homomorphisms $D \xrightarrow{h} R$ and $D \xrightarrow{h'} R'$. If $\perp$ is a null present in D, define $g(\perp) = h(\perp) \otimes h'(\perp)$. Then, if $(v_1, \ldots, v_m)$ is tuple in D, then $(h(v_1), \ldots, h(v_m))$ is a tuple in R and $(h'(v_1), \ldots, h'(v_m))$ is a tuple in R', and thus

$$\big(h(v_1) \otimes h'(v_1), \ldots, h(v_m) \otimes h'(v_m)\big) \ = \ (g(v_1), \ldots, g(v_m))$$

is a tuple in $R \otimes R'$, which shows that g is a homomorphism $D \xrightarrow{g} R \otimes R'$. This shows that $R \otimes R'$ is a glb of R and R'.

Proof of Theorem 3

For the first statement, we use the fact that every countable partial order can be embedded into the lattice of cores of graphs [25]. We use this with the order $\langle \mathbb{Q}, < \rangle$ (which in fact can be embedded into infinitely many intervals in the lattice of cores [24]). In particular, there is a family of graphs $\mathcal{G}_\mathbb{Q} = \{G_q \mid q \in \mathbb{Q}\}$ such that each G_q is a core and $G_q \prec G_{q'}$ iff $q < q'$. Since each G_q is finite, the union of all the domains of G_q's is countable, and thus we can assume that the nodes of all the G_q's come from $\mathcal{N}$, i.e., we can view graphs in $\mathcal{G}_\mathbb{Q}$ as naïve binary tables.

Now assume that every subset $\mathcal{X}$ of $\mathcal{G}_\mathbb{Q}$ has a glb. Then every subset $\mathcal{X}$ will have a lub as well (which is the glb of $\{G_q \in \mathcal{G}_\mathbb{Q} \mid q > q' \text{ for all } G_{q'} \in \mathcal{X}\}$). Hence, the Dedekind-MacNeille completion of $\langle \mathcal{G}_\mathbb{Q}, \prec \rangle$, which is isomorphic to the completion of $\langle \mathbb{Q}, < \rangle$ is embeddable in the preorder of naïve tables. Since the latter is countable but the former is not, this implies that there are families $\mathcal{X}$ without a glb (in fact the same cardinality argument shows that the number of sets without a glb is uncountable).

For the second statement, let P_n be a directed path of length n, and let C_n be a directed cycle of length n. We shall assume that their nodes come from $\mathcal{N}$, i.e., these are naïve tables. We first notice that there exist homomorphisms $P_n \xrightarrow{h} P_m$ whenever $n < m$, but not the other way around (since P_i's are rigid). For directed cycles, we have a homomorphism $C_{2^m} \xrightarrow{g_m} C_{2^{m-1}}$ for every m. Indeed, let C_{2^m} contain edges $(\perp_1, \perp_2), (\perp_2, \perp_3), \ldots, (\perp_{2^m-1}, \perp_{2^m}), (\perp_{2^m}, \perp_1)$, and let $C_{2^{m-1}}$ contain edges $(\perp'_1, \perp'_2), (\perp'_2, \perp'_3), \ldots, (\perp'_{2^{m-1}-1}, \perp'_{2^{m-1}}), (\perp'_{2^{m-1}}, \perp'_1)$. Then g_m is given by letting,

$$g_m(\perp_i) = g_m(\perp_{i+2^{m-1}}) = \perp'_i, \quad \text{for each } i \text{ with } 1 \leq i \leq 2^{m-1}.$$

Since each C_n is a core, we do not have homomorphisms going the other way (note that it is essential that we use directed cycles here: all undirected cycles of even length are bipartite, and thus have the same core, namely an undirected edge). Therefore, we have

$$P_1 \prec P_2 \prec P_3 \prec \ldots \prec P_n \prec \ldots \ldots \prec C_{2^m} \prec \ldots \prec C_8 \prec C_4 \prec C_2$$

We claim that the family $\{C_{2^m} \mid m > 0\}$ does not have a glb. Assume that it does, and a graph G is a glb. In particular, $P_n \prec G$ for all n, and $G \prec C_{2^m}$ for all m. There are two cases.

- Case 1: G is acyclic. Let k be the length of the longest path in G; by acyclicity, k is finite. Then by rigidity $P_{k+1} \not\preceq G$, and thus G is not a glb for $\{C_{2^m} \mid m > 0\}$.

- Case 2: G has cycles. Let k be the smallest length of a cycle in G. Since there is no homomorphism from C_k into C_{2^k} (in fact into any C_m for $m > k$), this means that there is no homomorphism from G to C_{2^k}, i.e., $G \not\preceq C_{2^k}$, and thus G is not a glb for $\{C_{2^m} \mid m > 0\}$.

Hence, $\{C_{2^m} \mid m > 0\}$ has no glb. In fact the same argument shows that for every $p \geq 1$, the set $\{C_{2^m \cdot p} \mid m > 0\}$ has no glb, completing the proof. $\square$

Proof of Proposition 8

First, assume that $D \preceq_{\mathrm{CWA}} D'$, i.e., there is a one-to-one homomorphism h from D to D'. Since $t \sqsubseteq h(\bar{t})$ for every homomorphism h, we have $D \sqsubseteq^\flat D'$. Since every tuple in D' is in $h(D)$, we obtain $D \sqsubseteq^\natural D'$. Now suppose the relation $\sqsubseteq^{-1}$ fails Hall's condition, that is, for some set S' of tuples in D', the set of tuples $S = \{\bar{t} \in D \mid \bar{t} \sqsubseteq \bar{t}'$ for some $\bar{t}' \in S'\}$ satisfies $|S| < |S'|$. Observe that, given any two tuples $\bar{t}$ in D and $\bar{t}'$ in D', if $\bar{t} \not\sqsubseteq \bar{t}'$, then there is no homomorphism $D \xrightarrow{g} D'$ such that $g(\bar{t}) = \bar{t}'$. Thus, there is no homomorphism $D \xrightarrow{g} D'$ such that its image contains the entire set S', and, in particular, there is no onto homomorphism from D to D'. This contradiction shows that Hall's condition must be satisfied.

Conversely, assume that $D \sqsubseteq^\natural D'$ and $\sqsubseteq^{-1}$ satisfies Hall's condition. For each $\bar{t}'$ in D', let $S_{\bar{t}'} = \{\bar{t} \in D \mid \bar{t} \sqsubseteq \bar{t}'\}$ (i.e., $\sqsubseteq^{-1}(\bar{t}')$). Hall's condition implies that the the family $\{S_{\bar{t}'} \mid \bar{t}' \in D'\}$ has a family of distinct representatives. Let $s_{\bar{t}'}$ be the distinct representative of $S_{\bar{t}'}$. Notice that $s_{\bar{t}'} \sqsubseteq \bar{t}'$.

We now construct an onto homomorphism $D \xrightarrow{h} D'$ as follows. For each $s_{\bar{t}'} = (a_1, \ldots, a_k)$ and $\bar{t}' = (b_1, \ldots, b_k)$, if a_i is a null, we let $h(a_i) = b_i$, for $i \le k$. Since $s_{\bar{t}'} \sqsubseteq \bar{t}'$, this implies $h(s_{\bar{t}'}) = \bar{t}'$. If $\bar{t}$ is a tuple not of the form $s_{\bar{t}'}$, then we know from $D \sqsubseteq^\natural D'$ that there is a tuple $\bar{t}'' \in D'$ such that $\bar{t} \sqsubseteq \bar{t}''$. Then again if $\bar{t} = (a_1, \ldots, a_k)$ and $\bar{t}'' = (b_1, \ldots, b_k)$, then for each null a_i we let $h(a_i) = b_i$; as before, $h(\bar{t}) = \bar{t}''$. Since each null occurs at most once, this definition of h is correct, and we showed that it is a homomorphism. Since we have $h(s_{\bar{t}'}) = \bar{t}'$ for every $\bar{t}' \in D'$, it is also an onto homomorphism. Thus, $D \preceq_{\mathrm{CWA}} D'$, proving the proposition.

Remark The proof also shows that when $\sqsubseteq^{-1}$ satisfies Hall's condition, $D \sqsubseteq^\flat D'$ alone (which is a weaker condition than $D \sqsubseteq^\natural D'$) implies $D \preceq_{\mathrm{CWA}} D'$. In particular, the following are equivalent: (a) $D \preceq_{\mathrm{CWA}} D'$; (b) $D \sqsubseteq^\flat D'$ and $\sqsubseteq^{-1}$ satisfies Hall's condition; and (c) $D \sqsubseteq^\natural D'$ and $\sqsubseteq^{-1}$ satisfies Hall's condition. $\square$

Proof of Theorem 4

First notice that $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}'$ is a lower bound of $\mathfrak{D}$ and $\mathfrak{D}'$. Indeed, we have a homomorphism $\iota : \mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \to \mathbf{M}_\lambda$. Consider the mapping h that sends each $\perp_{xy}$ to x. Then (ι, h) is a homomorphism from $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}'$ to $\mathfrak{D}$. Indeed, for each ν in $\mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'}$, we have

$$h(\rho \otimes \rho'(\nu)) \;=\; h(\rho(\iota(\nu)) \otimes \rho'(\iota'(\nu))) \;=\; \rho(\iota(\nu)).$$

The proof of the existence of h' such that $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}' \xrightarrow{(\iota', h')} \mathfrak{D}'$ is the same. Hence, $\mathfrak{D} \wedge_{\mathcal{K}} \mathfrak{D}'$ is a lower bound of $\mathfrak{D}$ and $\mathfrak{D}'$.

For the converse, we need an auxiliary result first. Define $\mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}$ as the restriction of $\mathbf{M}_\lambda \times \mathbf{M}'_{\lambda'}$ to pairs (ν, ν') so that $\lambda(\nu) = \lambda'(\nu')$. We claim that it is a glb of $\mathbf{M}_\lambda$ and $\mathbf{M}'_{\lambda'}$ in the class of Σ-colored structures (i.e., structures where each node is assigned one label in Σ). Clearly, it is a lower bound, as the projections serve as homomorphisms.

Now suppose we have a structure $\mathbf{M}^*_{\lambda^*}$ and homomorphisms $\mathbf{M}^*_{\lambda^*} \xrightarrow{g} \mathbf{M}_\lambda$ and $\mathbf{M}^*_{\lambda^*} \xrightarrow{g'} \mathbf{M}'_{\lambda'}$ (recall that the $\longmapsto$ notation refers to homomorphisms in the graph-theoretic sense, i.e., defined on all elements). In particular, for each ν in $\mathbf{M}^*$, we have $\lambda^*(\nu) = \lambda(g(\nu)) = \lambda'(g'(\nu))$. Hence, (g, g') is a homomorphism from $\mathbf{M}^*_{\lambda^*}$ to $\mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}$, which shows that the latter serves as a glb in the preorder of Σ-colored structures. In particular, for an arbitrary $\mathcal{K}$, we shall have $\mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'} \preceq \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}$ since by definition, $\mathbf{M}_\lambda \sqcap_{\mathcal{K}} \mathbf{M}'_{\lambda'}$ (when it exists) is a Σ-colored structure.

We next define $\mathfrak{D} \wedge_\Sigma \mathfrak{D}'$ as $\langle \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle$. Here $\rho \otimes \rho'\big((\nu, \nu')\big) = \rho(\nu) \otimes \rho'(\nu')$; this is of course well-defined since $\lambda(\nu) = \lambda'(\nu')$ for every $(\nu, \nu') \in \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}$. We claim that $\mathfrak{D} \wedge_\Sigma \mathfrak{D}'$ is a glb of $\mathfrak{D}$ and $\mathfrak{D}'$ in the class of all generalized databases of schema $\mathbb{S}$. The proof of this statement is similar to the proof of Proposition 5. We first note that there are homomorphisms from $\mathfrak{D} \wedge_\Sigma \mathfrak{D}'$ to both $\mathfrak{D}$ and $\mathfrak{D}'$. Indeed, we have $\mathfrak{D} \wedge_\Sigma \mathfrak{D}' \xrightarrow{(\pi_1, h_1)} \mathfrak{D}$ where π_1 is the first projection, and $h_1(\perp_{xy}) = x$ for each $\perp_{xy}$, and likewise we have a homomorphism from $\mathfrak{D} \wedge_\Sigma \mathfrak{D}'$ to $\mathfrak{D}'$ using second projections. Suppose now we have $\mathfrak{D}^* = \langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle$ such that $\mathfrak{D}^* \preceq \mathfrak{D}$ and $\mathfrak{D}^* \preceq \mathfrak{D}'$; that is, we have homomorphisms

$$\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \xrightarrow{(g, h)} \langle \mathbf{M}_\lambda, \rho \rangle \quad \text{and} \quad \langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \xrightarrow{(g', h')} \langle \mathbf{M}'_{\lambda'}, \rho' \rangle.$$

Then we have a homomorphism $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \xrightarrow{(g^*, h^*)} \langle \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle$ defined as follows. First, $g^*(\nu) = (g(\nu), g'(\nu))$. Since the labels of $g(\nu)$ in $\mathbf{M}_\lambda$ and $g'(\nu)$ in $\mathbf{M}'_{\lambda'}$ are the same, we can define $h^*(\nu) = \rho(g(\nu)) \otimes \rho'(g'(\nu))$. It is routine to verify that this defines a homomorphism and thus $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \preceq \langle \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle$, showing that the latter is a glb in the class of all generalized databases.

Now suppose we have a generalized database $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle$ such that $\mathbf{M}^*_{\lambda^*}$ is in $\mathcal{K}$ and both $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \preceq \mathfrak{D}$ and $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \preceq \mathfrak{D}'$ hold. Then, by the above, we have

$$\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \xrightarrow{(g,h)} \langle \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle \tag{7}$$

for some homomorphism (g, h). We also have

$$\langle \mathbf{M}_\lambda \sqcap_\mathcal{K} \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle \xrightarrow{(g',id)} \langle \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle \tag{8}$$

since $\mathbf{M}_\lambda \sqcap_\mathcal{K} \mathbf{M}'_{\lambda'} \preceq \mathbf{M}_\lambda \sqcap_\Sigma \mathbf{M}'_{\lambda'}$ (the glb in a class cannot be greater than an unrestricted glb). Furthermore, since $\mathbf{M}^*_{\lambda^*}$ is in $\mathcal{K}$ and both $\mathbf{M}^*_{\lambda^*} \preceq \mathbf{M}_\lambda$ and $\mathbf{M}^*_{\lambda^*} \preceq \mathbf{M}'_{\lambda'}$ hold, we have $\mathbf{M}^*_{\lambda^*} \xmapsto{f} \mathbf{M}_\lambda \sqcap_\mathcal{K} \mathbf{M}'_{\lambda'}$. Combining this with (7) and using the fact that in (8) the homomorphism on values is the identity, we conclude $\langle \mathbf{M}^*_{\lambda^*}, \rho^* \rangle \xmapsto{(f,h)} \langle \mathbf{M}_\lambda \sqcap_\mathcal{K} \mathbf{M}'_{\lambda'}, \rho \otimes \rho' \rangle$, showing that the latter is a glb of $\langle \mathbf{M}_\lambda, \rho \rangle$ and $\langle \mathbf{M}'_{\lambda'}, \rho' \rangle$ in class $\mathcal{K}$. $\qquad\square$

Proof of Theorem 5

Let $\mathfrak{D}$ be a source instance (in particular, it is an instance without nulls), and let $\mathbb{M}$ be a mapping containing rules of the form $\mathfrak{I} \to \mathfrak{I}'$. We need to prove two facts: first, that $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ is a solution, and second, that whenever $\mathfrak{I}''$ is in $\mathbb{M}(\mathfrak{D})$ and $\mathfrak{D}'$ is a solution, we have $\mathfrak{I}'' \preceq \mathfrak{D}'$. Indeed, this implies universality of $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$: if $\mathfrak{D}'$ is an arbitrary solution in $\mathcal{K}$, then $\mathfrak{I}'' \preceq \mathfrak{D}'$ for all $\mathfrak{I}'' \in \mathbb{M}(\mathfrak{D})$ and hence $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D}) \preceq \mathfrak{D}'$; therefore, there is a homomorphism from $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ to $\mathfrak{D}'$, proving universality.

We first show that $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ is a solution. Assume, for a rule $\mathfrak{I} \to \mathfrak{I}' \in \mathbb{M}$ that there is a homomorphism $\mathfrak{I} \xrightarrow{(h_1,h_2)} \mathfrak{D}$. Consider the instance $h_2(\mathfrak{I}')$. Since $h_2(\mathfrak{I}') \preceq \bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$, we have a homomorphism $h_2(\mathfrak{I}') \xrightarrow{(g,h')} \bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$. Since $\mathfrak{D}$ is a source (i.e., has no nulls), the domains of h' and h_2 are disjoint, and thus $h = h' \circ h_2$ is a homomorphism that extends h_2. We then have $\mathfrak{I}' \xrightarrow{(id,h_2)} h_2(\mathfrak{I}') \xrightarrow{(g,h')} \bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ implying $\mathfrak{I}' \xrightarrow{(g,h)} \bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ with h extending h_2. This proves that $\bigvee_\mathcal{K} \mathbb{M}(\mathfrak{D})$ is a solution.

Finally, we need to show that $\mathfrak{I}'' \preceq \mathfrak{D}'$ for each $\mathfrak{I}'' \in \mathbb{M}(\mathfrak{D})$ and a solution $\mathfrak{D}'$. If $\mathfrak{I}''$ is in $\mathbb{M}(\mathfrak{D})$, then for some $\mathfrak{I} \to \mathfrak{I}'$ in $\mathbb{M}$ and some homomorphism $\mathfrak{I} \xrightarrow{(h_1,h_2)} \mathfrak{D}$ we have $\mathfrak{I}'' = h_2(\mathfrak{I}')$. Since $\mathfrak{D}'$ is a solution, there exists a homomorphism $\mathfrak{I}' \xrightarrow{(g_1,g_2)} \mathfrak{D}'$ such that g_2 coincides with h_2 on nulls common to $\mathfrak{I}$ and $\mathfrak{I}'$. Note that since the image of h_2 contains only constants (as $\mathfrak{D}$ is a source), g_2 is a disjoint union of h_2 and a homomorphism h' defined on nulls in $\mathfrak{I}'$ not present in $\mathfrak{I}$. Since we have a homomorphism $\mathfrak{I}' \xrightarrow{(id,h_2)} h_2(\mathfrak{I})$, this implies that (g_1, h') is a homomorphism from $h_2(\mathfrak{I}')$ to $\mathfrak{D}'$. Thus, $h_2(\mathfrak{I}') \preceq \mathfrak{D}'$, i.e., $\mathfrak{I}'' \preceq \mathfrak{D}'$. This completes the proof. $\qquad\square$

Proof of Proposition 10

Consider trees T_1 and T_2, both having their roots labeled a, and one child each: in T_1 it is labeled b, and in T_2 it is labeled c.

Now consider the following trees. In T', we have the root labeled a, with two children labeled b and c. In T'', we have the root labeled d, with two children, both labeled a, the first of which has a single child labeled b, and the other a single child labeled c. It is easy to see that $T_1, T_2 \preceq T'$ and $T_1, T_2 \preceq T''$.

Suppose T is a lub of T_1 and T_2. Then there are homomorphisms $T_1 \xmapsto{g_1} T$ and $T_2 \xmapsto{g_2} T$. Suppose the values of g_1 and g_2 coincide on the a-labeled nodes of T_1 and T_2. Then T contains a copy of T', in which case there is no homomorphism from T to T'', i.e., $T \not\preceq T''$, and thus T cannot be an lub of T_1 and T_2. If, on the other hand, the values of g_1 and g_2 are different on the a-labeled nodes of T_1 and T_2, then T contains copies of T_1 and T_2 which are disjoint. In this case, there must be a common ancestor of a-nodes of these copies, which implies that there is no homomorphism from T to T'', i.e., $T \not\preceq T'$. Hence, in this case T cannot be an lub of T_1 and T_2 either. This shows there is no lub for $\{T_1, T_2\}$. $\qquad\square$

Proof of Proposition 11

In the proof we shall use abbreviations $\mathfrak{D}$ for $\langle \mathbf{M}_\lambda, \rho \rangle$ and $\mathfrak{D}'$ for $\langle \mathbf{M}'_{\lambda'}, \rho' \rangle$.

18

For the NP upper bound, we show that, if the sentence φ is fixed (and thus $\mathbb{S}$ is fixed as it provides the vocabulary for φ), then if there is $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$ such that $\mathbf{M}'_{\lambda'} \models \varphi$, then one can find such $\mathbf{M}'_{\lambda'}$ whose size is polynomial in the size of $\mathbf{M}_\lambda$ (with polynomial depending only on φ). We do this by capturing the existence of such a $\langle \mathbf{M}'_{\lambda'}, \rho' \rangle$ by means of satisfiability of an $\exists^* \forall^*$ sentence whose block of existentials will be of size polynomial in the size of $\mathbf{M}_\lambda$.

To construct such a sentence, assume that the universe of $\mathbf{M}_\lambda$ consist of nodes $\nu_1, \ldots, \nu_m$, and let $\mathrm{Diag}^+[\mathbf{M}_\lambda](x_1, \ldots, x_m)$ be the positive diagram of $\mathbf{M}_\lambda$ (i.e., whenever $(\nu_{i_1}, \ldots, \nu_{i_k})$ is in a relation R in $\mathbf{M}$, we add an atom $R(x_{i_1}, \ldots, x_{i_k})$, and whenever $\lambda(\nu_i) = a$ we add an atom $P_a(x_i)$; then $\mathrm{Diag}^+(x_1, \ldots, x_m)$ is the conjunction of all such atoms). We then define

$$\alpha[\mathbf{M}_\lambda](\bar{x}) \;=\; \mathrm{Diag}^+[\mathbf{M}_\lambda] \;\wedge\; \forall x \; \neg \bigvee_{a,b \in \Sigma, a \neq b} (P_a(x) \wedge P_b(x));$$

it says that the labeling in the diagram is proper (i.e., at most one label per node).

Let $\bot_1, \ldots, \bot_k$ be the nulls, and $a_1, \ldots, a_p$ be the constants used in $\mathfrak{D}$. Let $b_1, \ldots, b_k$ be constants different from the a_i's. Let h be a function from $\bot_1, \ldots, \bot_k$ into $\{a_1, \ldots, a_p, b_1, \ldots, b_k\}$.

We define new unary predicates $V_{ir}(\cdot)$ and $W_{jr}(\cdot)$ for $i \leq p$ and $j \leq p+k$ and r ranging from 1 to the maximum arity of a symbol in Σ (i.e., the maximum number of attributes a node can carry). The intended interpretation of $V_{ir}(x)$ (or $W_{jr}(x)$) in the homomorphic image $h(\mathbf{M}_\lambda)$ is that the rth attribute of x is a_i (or the j element in the list $\{a_1, \ldots, a_p, b_1, \ldots, b_k\}$ via applying h).

We then define $\mathrm{Diag}^+[\rho, h](x_1, \ldots, x_m)$ as the conjunction of all the atoms $V_{ir}(x_l)$ if the rth attribute of ν_i is a_i, and $W_{jr}(x_l)$ if the the rth attribute of ν_i is $\bot_q$ such that $h(\bot_q) = a_j$ for $j \leq p$ and b_{j-p} for $j > p$. We then let $\beta[\rho, h](\bar{x})$ be the conjunction of $\mathrm{Diag}^+[\rho, h]$ and sentences

$$\forall x \; \big(P_a(x) \rightarrow \bigwedge_{i \leq p, \; j \leq p+k, \; r > ar(a)} \neg V_{ir}(x) \wedge \neg W_{jr}(x) \big)$$

for each $a \in \Sigma$, saying that attributes with numbers exceeding the arity of a are not defined, and

$$\forall x \; \neg \big((\bigvee_{i,j,r} V_{ir}(x) \wedge W_{jr}(x)) \vee (\bigvee_{i,i',r; \; i \neq i'} V_{ir}(x) \wedge V_{i'r}(x)) \vee (\bigvee_{j,j',r; \; j \neq j'} W_{jr}(x) \wedge W_{j'r}(x)) \big)$$

saying that no node can have two different values assigned to the same attribute.

Now it follows from the construction that the sentence

$$\psi[\mathfrak{D}, h] \;=\; \exists \bar{x} \big(\alpha[\mathbf{M}_\lambda](\bar{x}) \;\wedge\; \beta[\rho, h](\bar{x}) \big)$$

is true in a generalized database $\mathfrak{D}'$ without nulls iff there exists a homomorphism (g, h) from $\mathfrak{D}$ into it (simply given by sending ν_i to the witness of the existential quantifier $\exists x_i$). Thus, $\psi[\mathfrak{D}, h] \wedge \varphi$ is true in generalized databases $\mathfrak{D}'$ without nulls iff they satisfy φ and there is a homomorphism (g, h) from $\mathfrak{D}$ and $\mathfrak{D}'$.

Now assume we have $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$ with $\mathbf{M}'_{\lambda'}$ satisfying φ. In the homomorphism $\mathfrak{D} \xrightarrow{(g,h)} \mathfrak{D}'$ we can always change values in the image of h so that those which are not present in $\mathfrak{D}$ come from $\{b_1, \ldots, b_k\}$ (since $\mathfrak{D}$ has at most k nulls). Therefore, there is $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$ with $\mathbf{M}'_{\lambda'}$ satisfying φ iff there is a $\mathfrak{D}'$ satisfying $\psi[\mathfrak{D}, h] \wedge \varphi$. Note that the latter is an $\exists^* \forall^*$ sentence, and the number of existential quantifiers in it coincides with the size of the universe of $\mathfrak{D}$ plus the number of existential quantifiers in φ (which is fixed). It is known that if an $\exists^* \forall^*$ has a model, then it has one in which the size of the universe does not exceed the number of existentially quantified variables. Hence, to solve $\mathrm{Cons}(\varphi)$, it suffices to guess a structure $\mathbf{M}'_{\lambda'}$ and a function ρ' on it with the universe of $\mathbf{M}'$ being at most linear in the size of the universe of $\mathbf{M}$. Since the schema $\mathbb{S}$ is fixed, the size of the guessed $\mathfrak{D}'$ is polynomial in the size of $\mathfrak{D}$, and since φ is fixed, checking $\mathbf{M}'_{\lambda'} \models \varphi$ can be done in polynomial time. This shows that the algorithm for solving $\mathrm{Cons}(\varphi)$ runs in NP.

The remaining two items are easy. First, if $\Sigma = \{a\}$ and $ar(a) = 0$ (i.e., there is no labeling and there are no data values), consistency problem asks whether there exists $\mathbf{M}'$ satisfying φ such that there is a homomorphism from $\mathbf{M}$ to $\mathbf{M}'$. Since each relational structure can be described up to isomorphism by an $\exists^* \forall$ sentence, a particular case of the consistency problem is whether there is a homomorphism from $\mathbf{M}$ into a fixed structure $\mathbf{M}'$. Considering the vocabulary of graphs, and taking $\mathbf{M}'$ to be K_3, we obtain NP-hardness.

Finally, if φ is an existential sentence, $\mathrm{Cons}(\varphi)$ simply returns true if φ is satisfiable and false otherwise (note that since φ is fixed we check this is constant time). Indeed, if φ is satisfiable, consider a structure $\mathbf{M}^1_{\lambda_1}$ that satisfies

it, and define ρ^1 on it arbitrarily, to obtain a generalized database $\mathfrak{D}^1$. Then take the input $\mathfrak{D}$, and consider an arbitrary $\mathfrak{D}^2 \in [\![\mathfrak{D}]\!]$. Then the disjoint union of $\mathfrak{D}^1$ and $\mathfrak{D}^2$ is in $[\![\mathfrak{D}]\!]$ and satisfies φ. This completes the proof. $\square$

Proof of Theorem 6

Suppose $\mathfrak{D} = \langle \mathbf{M}_\lambda, \rho \rangle$ and $\mathfrak{D}' = \langle \mathbf{M}'_{\lambda'}, \rho' \rangle$ are two generalized databases of the same schema, and assume that ρ has Codd interpretation. Recall the definition of $\sqsubseteq$ on tuples over $\mathcal{C} \cup \mathcal{N}$: $(a_1, \ldots, a_m) \sqsubseteq (b_1, \ldots, b_m)$ if, for each i, whenever $a_i \in \mathcal{C}$, then $b_i = a_i$. Now define a binary relation $R(\mathfrak{D}, \mathfrak{D}')$ between the nodes of $\mathbf{M}$ and $\mathbf{M}'$ as follows:

$$(\nu, \nu') \in R(\mathfrak{D}, \mathfrak{D}') \quad \Leftrightarrow \quad \lambda(\nu) = \lambda'(\nu') \text{ and } \rho(\nu) \sqsubseteq \rho'(\nu').$$

We need the following observation.

Lemma 3. $\mathfrak{D} \preceq \mathfrak{D}'$ iff there is a homomorphism $\mathbf{M}_\lambda \overset{h}{\longmapsto} \mathbf{M}'_{\lambda'}$ such that the graph of h is contained in $R(\mathfrak{D}, \mathfrak{D}')$.

Proof of Lemma 3. First, suppose $\mathfrak{D} \preceq \mathfrak{D}'$, i.e., we have a homomorphism $\langle \mathbf{M}_\lambda, \rho \rangle \overset{(h_1, h_2)}{\longrightarrow} \langle \mathbf{M}'_{\lambda'}, \rho' \rangle$. Then, for each node ν in $\mathbf{M}$, we have $\rho(\nu) \sqsubseteq \rho'(h_1(n))$ and thus h_1 is the required homomorphism between $\mathbf{M}_\lambda$ and $\mathbf{M}'_{\lambda'}$. Conversely, suppose we find a homomorphism $\mathbf{M}_\lambda \overset{h}{\longmapsto} \mathbf{M}'_{\lambda'}$ whose graph of h is contained in $R(\mathfrak{D}, \mathfrak{D}')$. Then take any tuple $\bar{a}$ in $\rho(\nu)$ for ν in $\mathbf{M}$, and let $\bar{b}$ be the tuple (of the same length) $\rho'(h(\nu))$. Since $\bar{a} \sqsubseteq \bar{b}$, we define a mapping h' by mapping $\bar{a}$ into $\bar{b}$. By the definition of $\sqsubseteq$, it only maps nulls to nulls or constants, and since ρ is Codd, it is well-defined. It follows immediately that (h, h') is a homomorphism from $\langle \mathbf{M}_\lambda, \rho \rangle$ to $\langle \mathbf{M}'_{\lambda'}, \rho' \rangle$. This proves the lemma.

Let $\mathbf{A}$ and $\mathbf{B}$ be two relational structures of the same vocabulary, with universes A and B respectively, which we assume to be disjoint. Let R be a subset of $A \times B$. We say that a homomorphism $\mathbf{A} \overset{h}{\longmapsto} \mathbf{B}$ is R-compatible if the graph of h is contained in R. We consider the following problem: Given $\mathbf{A}, \mathbf{B}$ and R as above, is there an R-compatible homomorphism from $\mathbf{A}$ to $\mathbf{B}$? We denote this by R-Hom$(\mathbf{A}, \mathbf{B})$.

We need to show the following.

Lemma 4. *For each fixed k, the problem R-Hom$(\mathbf{A}, \mathbf{B})$ can be solved in polynomial time if $\mathbf{A}$ has treewidth at most k.*

For $R = A \times B$ (i.e., no restriction on the homomorphism), this result is of course well-known (cf. [20]). The theorem will immediately follow from Lemmas 3 and 4.

One way of seeing that Lemma 4 is true is to go over the proof of [20, Theorem 13.12], and make sure that everything works when one adds the extra relation R (one needs to redefine the game by restricting the moves of the duplicator to respect R, prove its correctness, and show that the fixed-point construction works). Instead of doing it that way, we use a different technique, following [29] and encode the problem as the combined complexity of evaluating conjunctive queries with a fixed number of variables.

Let the universe of $\mathbf{A}$ be $\{\nu_1, \ldots, \nu_m\}$. Define new unary predicates $U_0, U_1, \ldots, U_m$ and a new binary predicate P and expand $\mathbf{A}$ to a structure $\mathbf{A}'$ of the vocabulary expanded with the U_i's and P as follows:

- The universe of $\mathbf{A}'$ is the disjoint union of $\{\nu_1, \ldots, \nu_m\}$ and a set $\{\omega_1, \ldots, \omega_m\}$.
- The interpretation of the predicates of $\mathbf{A}$ is inherited from $\mathbf{A}$.
- The unary predicates are interpreted as follows: U_0 as $\{\nu_1, \ldots, \nu_m\}$, and each U_i as $\{\omega_i\}$, for $1 \le i \le m$.
- The binary predicate P is interpreted as $\{(\nu_i, \omega_i) \mid 1 \le i \le m\}$.

Next we define $\mathbf{B}_R$, an expansion of $\mathbf{B}$ with the $U_0, U_1, \ldots, U_m$, and P as follows. Assume that the universe of $\mathbf{B}$ is $\{\nu'_1, \ldots, \nu'_l\}$. Then:

- The universe of $\mathbf{B}_R$ is the disjoint union of $\{\nu'_1, \ldots, \nu'_l\}$ and a set $\{\omega'_1, \ldots, \omega'_m\}$.
- The interpretation of the predicates of $\mathbf{B}$ is inherited from $\mathbf{B}$.
- The interpretation of U_0 is $\{\nu'_1, \ldots, \nu'_l\}$, and the interpretation of U_i is $\{\omega'_i\}$, for each $1 \le i \le m$.
- The predicate P is interpreted as $\{(\nu'_i, \omega'_j) \mid (\nu_j, \nu'_i) \in R\}$.

Lemma 5. *R-Hom*$(\mathbf{A}, \mathbf{B})$ *has a solution iff there is a homomorphism from* $\mathbf{A}'$ *to* $\mathbf{B}_R$.

Proof of Lemma 5. Assume a homomorphism $\mathbf{A} \stackrel{h}{\longmapsto} \mathbf{B}$ is such that its graph is contained in R. Extend it to a mapping h' from $\mathbf{A}'$ to $\mathbf{B}_R$ by letting $h'(\omega_i) = \omega_i'$. Then h' is a homomorphism. Indeed, we only need to check the preservation of P. Assume that (ν_j, ω_j) is in the interpretation of P in $\mathbf{A}$. Let $\nu_i' = h(\nu_j)$. We know then that $(\nu_j, \nu_i') \in R$ and hence $(h'(\nu_j), h'(\omega_j)) = (\nu_i', \omega_j')$ is the interpretation of P in $\mathbf{B}_R$.

Conversely, suppose we have a homomorphism $\mathbf{A}' \stackrel{h'}{\longmapsto} \mathbf{B}_R$. Then $h'(\omega_j) = \omega_j'$ for each $j \leq m$. Furthermore, the restriction h of h' to U_0 is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. Thus, we need to show that the graph of h is contained in R. Assume that $h(\nu_j) = \nu_i'$. Since (ν_j, ω_j) is in P, we must have $(h'(\nu_j), h'(\omega_j)) = (h(\nu_j), \omega_j') = (\nu_i', \omega_j') \in P$ in $\mathbf{B}_R$. By the definition of P, this means that (ν_j, ν_i') is in R, i.e., $(\nu_j, h(\nu_j))$ is in R, showing that the graph of h is contained in R. This proves the lemma.

Thus, to prove Lemma 4, it remains to show that one can check in polynomial time whether there is a homomorphism from $\mathbf{A}'$ to $\mathbf{B}_R$ (note that vocabulary of the new structures is now the size of $\mathbf{A}$ itself). For this, observe that under the assumption that the treewidth of $\mathbf{A}$ is k, the treewidth of $\mathbf{A}'$ is at most $k + 1$. Indeed, it suffices to assign each ω_i arbitrarily to one of the nodes in the tree decomposition that contains ν_i. Clearly this preserves the properties of tree decomposition and increases the width by at most 1.

Now we apply the results from [29] showing that the canonical conjunctive query $Q_{\mathbf{A}'}$ of $\mathbf{A}'$ can be expressed in the $\exists, \wedge$-fragment of FO with $k + 1$ variables, as the treewidth of $\mathbf{A}'$ is at most $k + 1$. Since there is a homomorphism from $\mathbf{A}'$ to $\mathbf{B}_R$ iff $\mathbf{B}_R \models Q_{\mathbf{A}'}$, the result now follows from the fact [42] that the combined complexity of evaluating conjunctive queries with a fixed number of variables is in PTIME. This concludes the proof of Lemma 4 and the theorem. □

Proof of Theorem 7

We start by giving a precise definition of the notion $\mathfrak{D} \models \varphi$ when φ is a sentence of FO($\mathbb{S}, \sim$). Given a generalized schema $\mathbb{S} = (\sigma, \Sigma, \rho)$, define a relational vocabulary $\tau_\mathbb{S}$ which contains all relations in σ, a unary relation P_a for each $a \in \Sigma$, and relations EQ_{ij} for all $i \neq j$ that do not exceed the maximum arity of $ar(a)$ of $a \in \Sigma$. We turn each generalized database $\mathfrak{D} = \langle \mathbf{M}_\lambda, \rho \rangle$ into a $\tau_\mathbb{S}$ database $\mathfrak{D}_{EQ}$ as follows:

- The universe of $\mathfrak{D}_{eq}$ is the universe of $\mathbf{M}$.
- $\mathfrak{D}_{EQ}$ contains all the σ-relations from $\mathbf{M}$.
- It contains, for each $a \in \Sigma$, a unary relation P_a that consists of the nodes labeled a.
- Whenever ν, ν' are two nodes in $\mathbf{M}$, such that $ar(\lambda(n)) \geq i$ and $ar(\lambda(j)) \geq i$, and the ith attribute in $\rho(\nu)$ equals the jth attribute in $\rho(\nu')$, then (ν, ν') is in EQ_{ij}.

So when we write $\mathfrak{D} \models \varphi$ we mean $\mathfrak{D}_{EQ} \models \varphi$ (as $\mathfrak{D}_{EQ}$ is a first-order structure of the same vocabulary as φ). Now we proceed to prove the theorem.

a) Note that each homomorphism $\mathfrak{D} \stackrel{(h_1, h_2)}{\longrightarrow} \mathfrak{D}'$ is also a homomorphism $\mathfrak{D}_{EQ} \stackrel{h_1}{\longmapsto} \mathfrak{D}'_{EQ}$ (recall that the notation $\longmapsto$ means homomorphism in the graph-theoretic sense). Indeed, if $(\nu, \nu') \in EQ_{ij}$ in $\mathfrak{D}$, then the value of the ith attribute of ν and the jth attribute of ν' are the same. Since h_1 preserves labeling, the number of attributes of $h_1(\nu)$ and $h_1(\nu')$ is the same as for ν and ν', resp., and thus their values (obtained by applying h_2) are the same. Hence, $(h_1(\nu), h_1(\nu')) \in EQ_{ij}$.

Now we conclude that for existential positive queries, certain answers can be found by naïve evaluation. Assume φ is existential positive, and let φ evaluate to true on $\mathfrak{D}$, i.e., $\mathfrak{D}_{EQ} \models \varphi$. Assume $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$, that is, there is a homomorphism $\mathfrak{D} \stackrel{(h_1, h_2)}{\longrightarrow} \mathfrak{D}'$. Then $\mathfrak{D}_{EQ} \stackrel{h_1}{\longmapsto} \mathfrak{D}'_{EQ}$, and since φ is existential positive, we have $\mathfrak{D}'_{EQ} \models \varphi$, thus showing $\mathsf{certain}(\varphi, D) = \mathrm{true}$.

Conversely, let φ evaluate to false on $\mathfrak{D}$. Let $\perp_1, \ldots, \perp_k$ be all the nulls used in $\mathfrak{D}$, and let $c_1, \ldots, c_k$ be distinct constants not used in $\mathfrak{D}$. Let $\mathfrak{D}'$ be obtained from $\mathfrak{D}$ by replacing each $\perp_i$ with c_i, for $i \leq k$. Then $\mathfrak{D}_{EQ} = \mathfrak{D}'_{EQ}$ and thus $\mathfrak{D}'_{EQ} \models \neg\varphi$. Since $\mathfrak{D}' \in [\![\mathfrak{D}]\!]$, this shows $\mathsf{certain}(\varphi, D) = \mathrm{false}$. Summing up, $\mathsf{certain}(\varphi, D) = \mathrm{true}$ iff $\mathfrak{D}_{EQ} \models \varphi$, which shows that the complexity bound is that of computing $\mathfrak{D}_{EQ}$ and evaluating a fixed FO sentence over it, i.e., DLogSpace.

b) Let φ be an existential sentence of FO($\mathbb{S}, \sim$), i.e., it is of the form $\exists \bar{x} \, \alpha$ where α is quantifier-free. Assume certain$(\varphi, \mathfrak{D})$ evaluates to false. Then we have a homomorphism $\mathfrak{D} \overset{(h_1, h_2)}{\longrightarrow} \mathfrak{D}'$ such that the range of h_2 is constants, and $\mathfrak{D}'_{EQ} \not\models \varphi$, i.e., $\mathfrak{D}'_{EQ} \models \forall \bar{x} \, \neg \alpha$. Let $\mathfrak{D}''_{EQ}$ be the restriction of $\mathfrak{D}'_{EQ}$ to the range of h_1. Since universal sentences are closed under taking substructures, we conclude $\mathfrak{D}''_{EQ} \not\models \varphi$. In particular, this tells us that there is a homomorphic image $\mathfrak{D}''$ of $\mathfrak{D}$ such that $\mathfrak{D}''_{EQ} \not\models \varphi$. Hence, to check whether certain$(\varphi, \mathfrak{D})$ evaluates to false it suffices to guess a homomorphism on $\mathfrak{D}$ and check, in polynomial time, whether the image of the homomorphism satisfies $\neg \varphi$. This shows that the problem of checking certain$(\varphi, \mathfrak{D})$ = false is in NP, and thus the problem of checking certain$(\varphi, \mathfrak{D})$ = true is in CONP.

To show CONP-hardness, consider a schema $\mathbb{S}$ in which σ contains a binary relation E, the alphabet has two letters a and b, with $ar(a) = 1$ and $ar(b) = 3$. Let G be an arbitrary undirected graph. We create a generalized database $\mathfrak{D}_G$ as follows. For each node of G, we have a node of $\mathfrak{D}_G$ labeled a. Unique attributes of such nodes are distinct nulls (i.e. if G has nodes $\nu_1, \ldots, \nu_m$, then their attributes are $\perp_1, \ldots, \perp_m$, resp.). Whenever there is an edge between ν and ν', we add both (ν, ν') and (ν', ν) to E. Finally, we have an additional isolated node ν_0 of color b with attribute values $1, 2, 3$.

Let now ψ be the FO($\mathbb{S}, \sim$) sentence

$$\forall x \forall y \, \left(P_a(x) \wedge P_b(y) \; \rightarrow \; (=_{11}(x, y) \; \vee \; =_{12}(x, y) \; \vee \; =_{13}(x, y)) \right)$$

saying that the attribute value of every a-labeled node must be present among attribute values of all b-labeled nodes. Consider then the sentence

$$\varphi_0 \;\; = \;\; \psi \; \rightarrow \; \exists x \exists y \, \left(P_a(x) \wedge P_a(y) \wedge E(x, y) \wedge =_{11}(x, y) \right)$$

saying that for each database satisfying ψ, there is an edge with the same attribute values of its endpoints. That is, for each graph in which nulls are assigned values present in all b-nodes, one edge must have the same value assigned to its vertices. This implies that certain$(\varphi_0, \mathfrak{D}_G)$ evaluates to true iff G is not 3-colorable, thus proving CONP-hardness.

c) The undecidability argument is standard and is already known for relational databases [2, 26]: one simply applies (the proof of) Trakhtenbrot's theorem. One lets φ code the computation of a universal Turing machine. The input consists of relations coding a Turing machine M to be simulated, and empty relations, of the vocabulary appropriate for coding traces of M. With this coding, completions will contain the code of M and either a partial trace, or some relations that do not simulate a trace. The query φ_1 returns false iff the second part of description is a halting trace of M on the empty input. Modeling this in FO is standard. One minor technicality is that the part of the input coding M can be extended itself by the OWA. But this is easy to deal with by numbering transitions and having on the side relations coding arithmetic on those numbers, as well as singleton with the maximum number. If in an extension that relation is no longer singleton, φ_1 returns true. This way simulations only happen with transitions of M. Thus, the certain answer on an input coding M is true iff M does not halt on the empty input, proving undecidability. $\qquad \square$