# Querying Incomplete Data: Complexity and Tractability via Datalog and First-Order Rewritings 

AMÉLIE GHEERBRANT<br>Université Paris Cité, CNRS, IRIF, F-75013, Paris, France<br>(e-mail: amelie@irif.fr)<br>LEONID LIBKIN<br>School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, UK<br>RelationalAI, Paris, France<br>(e-mail: l@libk.in)<br>ALEXANDRA ROGOVA<br>Université Paris Cite, CNRS, IRIF, F-75013, Paris, France<br>Data Intelligence Institute of Paris (diiP), Inria, Paris, France<br>(e-mail: rogova@irif.fr)<br>CRISTINA SIRANGELO<br>Université Paris Cite, CNRS, IRIF, F-75013, Paris, France DI ENS, ENS, PSL University, CNRS, Inria, Paris, France<br>(e-mail: cristina@irif.fr)

submitted 27 March 2023; revised 27 September 2023; accepted 21 October 2023


#### Abstract

To answer database queries over incomplete data, the gold standard is finding certain answers: those that are true regardless of how incomplete data is interpreted. Such answers can be found efficiently for conjunctive queries and their unions, even in the presence of constraints. With negation added, the problem becomes intractable however. We concentrate on the complexity of certain answers under constraints and on effficiently answering queries outside the usual classes of (unions) of conjunctive queries by means of rewriting as Datalog and first-order queries. We first notice that there are three different ways in which query answering can be cast as a decision problem. We complete the existing picture and provide precise complexity bounds on all versions of the decision problem, for certain and best answers. We then study a well-behaved class of queries that extends unions of conjunctive queries with a mild form of negation. We show that for them, certain answers can be expressed in Datalog with negation, even in the presence of functional dependencies, thus making them tractable in data complexity. We show that in general, Datalog cannot be replaced by first-order logic, but without constraints such a rewriting can be done in first order.


KEYWORDS: incomplete information, certain answers, datalog rewritings, first-order rewritings, functional dependencies, chase

## 1 Introduction

Answering queries over incomplete databases is crucial in many different scenarios such as data integration (Lenzerini 2002), data exchange (Arenas et al. 2014), inconsistency
management (Bertossi 2011), data cleaning (Geerts et al. 2013), ontology-based data access (OBDA) (Bienvenu and Ortiz 2015), and many others. The common thread running through all these applications lies in computing certain answers (Imielinski and Lipski 1984). Intuitively, this produces answers that are true in all possible worlds, that is, complete databases that an incomplete database represents. An incomplete database in itself is a set of tuples with missing information, plus integrity constraints. One can think, for example, of relations with nulls on which keys can be specified. Then, a possible world is obtained by substituting values for nulls so that all the keys are satisfied.

The notion of certain answers is sometimes too restrictive (for example, for some queries no answers are certain). In that case, an alternative is best answers: for them, there is no other tuple that is an answer in more possible worlds. However, computationally one encounters serious problems with both approaches. To start with, computing certain answers and best answers is intractable for first-order queries (Abiteboul et al. 1991; Libkin 2018) (already for data complexity). Finding such answers in restricted subclasses of first-order queries often relies on sophisticated algorithms - not naturally expressible by other queries - that are therefore difficult to implement in a DBMS. We know that restricting to unions of conjunctive queries allows one to overcome this difficulty by using naïve evaluation which computes certain answers in polynomial time (Imielinski and Lipski 1984). This amounts to evaluating queries over incomplete databases as if nulls were usual data values, thus merely using the standard database query engine to compute certain answers.

We address these problems in the present paper whose goal is two-fold.

1. We start by filling gaps in our knowledge of the complexity of answering queries over incomplete databases. Intractable bounds on certain and best answers cited above were obtained under different formulations of query answering as a decision problem. We show that there are three natural ways to represent query answering as a decision problem and classify the complexity of certain and best answers for all of them.
2. We then look at a way of finding query answers by leveraging the existing database technology, namely by finding query rewritings which, when evaluated on the incomplete database, give us certain answers. We show that for a class extending unions of conjunctive queries with a form of negation (but still falling short of all first-order queries), such rewritings can be found in Datalog with negation, thus giving us a tractable complexity bound.

To elaborate on the first point, the two existing decision versions of the query answering problem are as follows: (a) is a tuple in the answer? and (b) is the answer a member of a given family of sets? We add a third: (c) is the answer equal to a given set. We then prove that for certain answers, the complexity is coNP, $\mathrm{P}^{\mathrm{NP}[\log n]}$, and DP-complete for (a), (b), (c). The result for (a) has long been known of course. For best answers, the complexity is uniform: $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete for all variations (the result for (b) was previously known). We shall define these complexity classes in the next section; for the reader not familiar with them, they all lie within the second level of the polynomial hierarchy.

For the second theme of the paper, we look at query rewritings. This is a standard way of leveraging database technology in the case of incomplete or imperfect information,
and such rewritings were heavily used in data integration, data exchange, OBDA, query answering using views, consistent query answering, etc. (Calvanese et al. 2000; 2007; Calì et al. 2013; 2003b). First-order rewritings are particularly useful, as they allow to use the power of standard database query engines. In fact when they exist, the rewritten queries can be implemented in any relational query engine by expressing them in SQL, with no need to implement ad-hoc algorithms. Next best are rewritings into Datalog (with negation): these let us express queries using recursive features of SQL.

As already mentioned, for unions of conjunctive queries (and even some mild restrictions with guarded negation Gheerbrant et al. 2014) certain answers are computed by naïve evaluation without the presence of constraints. Under constraints, even such simple ones as keys, the picture is less complete. Indeed, keys and in general equality-generating dependencies (EGD) change the syntactic shape of a query that makes naive evaluation work.

- Certain answers to a conjunctive query $Q$ (or a union of CQs) on a database $D$ under key constraints $\Sigma$ can be found by naïve evaluation of $Q$ on the result of the chase of $D$ with $\Sigma$. Mathematically, $\operatorname{cert}_{\Sigma}(Q, D)=Q\left(\operatorname{chase}_{\Sigma}(D)\right)$, where on the lefthand side we have certain answers under constraints and on the right-hand side the naïve evaluation of $Q$ over the result of the chase. Here, chase ${ }_{\Sigma}$ refers to the classical textbook chase procedure with keys or more generally functional dependencies. In fact, the above result applies when $\Sigma$ is a set of functional dependencies or equality generating dependencies (EGDs), not just keys.

Unfortunately, the above result does not work when we move outside the class of select-project-join-union queries or unions of CQs. In fact even without constraints, certain answers to a query of the form $Q_{1}-Q_{2}$, where both $Q_{1}$ and $Q_{2}$ are CQs, are not necessarily produced by naïve evaluation. To see why, take a database containing one fact $R(1, \perp)$ where $\perp$ is a null and $Q_{1}$ returning $R$ while $Q_{2}$ is given by a formula $R(x, y) \wedge x=y$. Here, naïve evaluation of $Q_{1}-Q_{2}$ returns $R$ (as 1 is not equal to $\perp$ ), while certain answers is empty (as 0 is a possible value for $\perp$ ).

This motivates our question whether we can extend the class of CQs and their unions to obtain tractable evaluation of certain answers under constraints such as functional dependencies and EGDs. The answer is positive; in fact the query of the form $Q_{1}-Q_{2}$ above will be an example of a query in this class. To start with, the class must be such that finding certain answers for its queries without constraints is already tractable. We know one such class: it consists of arbitrary Boolean combinations of CQs, not just their union. We shall denote it by BCCQ. It was proved in Gheerbrant and Libkin (2015) that certain answers for it can be found in polynomial time (for data complexity), though the procedure was tableau-based and not suitable for implementation in a database system.

This is precisely what we do in the second part of this paper. We establish three main results:

1. For an arbitrary BCCQ $Q$ and a set of EGDs $\Sigma$, one can construct a Datalog (with negation) query $Q^{\prime}$ whose naive evaluation computes certain answers, thereby ensuring their polynomial-time data complexity.
2. There are however simple BCCQs, in fact even CQs, and keys, such that certain answers cannot be expressed as a first-order queries. Therefore, using Datalog was necessary.
3. Without constraints present, certain answers to BCCQs are not only polynomialtime computable as had been shown previously, but also can be expressed by firstorder queries and thus efficiently implemented in SQL databases without using recursion.

The Sections $4.2,4.3$, and 4.4 address these items, respectively.
Note that the material from this paper is based on the two conference papers Gheerbrant and Sirangelo (2019) and Gheerbrant et al. (2022).

## 2 Preliminaries

### 2.1 Incomplete databases and constraints

We represent missing information in relational databases in the standard way using nulls (Abiteboul et al. 1995; Imielinski and Lipski 1984; van der Meyden 1998). Incomplete databases are populated by constants and nulls, coming respectively from two countably infinite sets Const and Null. We denote nulls by $\perp$, sometimes with sub- or superscript. We also allow them to repeat, thus adopting the model of marked nulls, as customary in the context of applications such as OBDA or data integration and exchange.

A relational schema, or vocabulary $\sigma$, is a set of relation names with associated arities. A database $D$ over $\sigma$ associates to each relation name of arity $k$ in $\sigma$, a k-ary relation which is a finite subset of $(\text { Const } \cup \text { Null })^{k}$. Sets of constants and nulls occurring in $D$ are denoted by Const $(D)$ and $\operatorname{Null}(D)$. A database is complete if it contains no nulls, that is $\operatorname{Null}(D)=\emptyset$.

The active domain of $D$ is the set of all values appearing in $D$, that is adom $(D)=$ Const $(D) \cup \operatorname{Null}(D)$.

A valuation $v: \operatorname{Null}(D) \rightarrow$ Const on a database $D$ is a map that assigns constant values to nulls occurring in $D$. By $v(D)$ and $v(\bar{a})$, we denote the result of replacing each null $\perp$ by $v(\perp)$ in a database $D$ or in a tuple $\bar{a}$. The semantics $\llbracket D \rrbracket$ of an incomplete database $D$ is the set $\{v(D) \mid v$ is a valuation on $D\}$ of all complete databases it can represent. Here as is common in research on incomplete data, we use closed-world assumption (Imielinski and Lipski 1984; Reiter 1977) (i.e. everything we do not know to be true is automatically assumed to be false and no new tuple can be added).

An equality-generating dependency (EGD) is a first-order sentence of the form $\forall \bar{x} \quad\left(\varphi(\bar{x}) \rightarrow z=z^{\prime}\right)$, where $\varphi(\bar{x})$ is a conjunction of atoms (without constants), each variable in $\bar{x}$ occurs in some atom of $\varphi$, and $z, z^{\prime}$ are distinct variables in $\bar{x}$. As a special case, a functional dependency (FD) over a relation name $R$ is of the form $\forall \bar{x}, \bar{y}, z, z^{\prime} \quad\left(R(\bar{x}, \bar{y}, z) \wedge R\left(\bar{x}, \bar{y}^{\prime}, z^{\prime}\right) \rightarrow z=z^{\prime}\right)$. Throughout this paper, we will assume that a (possibly empty) set of EGDs $\Sigma$ is associated with the database schema $\sigma$.

A valuation $v$ is consistent with $\Sigma$ (or just consistent, when $\Sigma$ is clear from the context) if $v(D) \vDash \Sigma$. We denote by $\mathrm{V}(D)$ the set of all consistent valuations defined on $D$.

### 2.2 Query answering

An $m$-ary query Q of active domain $\operatorname{adom}(Q) \subseteq$ Const is a map that associates with a database $D$ a subset of $(\operatorname{adom}(D) \cup \operatorname{adom}(Q))^{m}$. To answer an $m$-ary query $Q$ over an
incomplete database $D$, we follow (Lipski 1984) and adopt a slight generalization of the usual intersection-based certain answers notion, defined as $\cap_{v} Q(v(D))$, and furthermore incorporate constraints into query answering.

The set of certain answers to $Q$ over $D$, with respect to a set of constraints $\Sigma$, is

$$
\operatorname{cert}_{\Sigma}(Q, D)=\left\{\bar{a} \in(\operatorname{adom}(D) \cup \operatorname{adom}(Q))^{m} \mid v(\bar{a}) \in Q(v(D)) \text { for all consistent } v\right\}
$$

For queries that explicitly use constants, we shall expand this to allow $\bar{a}$ range over $\operatorname{adom}(D)$ and those constants. The only difference with the usual notion is that we allow answers to contain nulls, to avoid pathological situations when answers known with certainty are not returned (e.g. in a query returning a relation $R$ one would expect $R$ to be returned while the intersection-based certain answer will only return null-free tuples). If the set of constraints $\Sigma$ is empty, we omit it and write simply $\operatorname{cert}(Q, D)$. Of course, every valuation is consistent with the empty set of constraints.

Following Libkin (2018), given a query $Q$, a database $D$, a set of constraints $\Sigma$, and a tuple $\bar{a}$ over $\operatorname{adom}(D) \cup \operatorname{adom}(Q)$, we let the support of $\bar{a}$ be the set of all valuations that witness it:

$$
\operatorname{Supp}_{\Sigma}(Q, D, \bar{a})=\{v \in \mathrm{~V}(D) \mid v(\bar{a}) \in Q(v(D))\}
$$

Again if $\Sigma=\emptyset$ we omit the subscript.
Supports thus measure how close a tuple is to certainty. We consider one answer to be better than another if it has more support. That is, given a database $D$, a $k$-ary query $Q$, and $k$-tuples $\bar{a}, \bar{b}$ over $\operatorname{adom}(D) \cup \operatorname{adom}(Q)$, we let

$$
\bar{a} \triangleleft_{Q, D}^{\Sigma} \bar{b} \Leftrightarrow \operatorname{Supp}_{\Sigma}(Q, D, \bar{a}) \subset \operatorname{Supp}_{\Sigma}(Q, D, \bar{b}) .
$$

The set of best answers to $Q$ over $D$ is defined as the set of answers for which there is no better one:

$$
\operatorname{best}_{\Sigma}(Q, D)=\left\{\bar{a} \mid \neg \exists \bar{b}: \bar{a} \triangleleft_{Q, D}^{\Sigma} \bar{b}\right\} .
$$

As the set of certain answers to $Q$ over $D$ is the set of answers that are witnessed by all valuations, note that it could also be defined using the notion of support. Namely, $\operatorname{cert}_{\Sigma}(Q, D)$ consists of all tuples $\bar{a} \in \operatorname{adom}(D)^{m}$ for which $\operatorname{Supp}_{\Sigma}(Q, D, \bar{a})$ contains all consistent valuations in $\mathrm{V}(D)$.

## Example 2.1

Let $Q(x)=\exists y(R(y) \wedge S(y, x))$ and $D=\left\{R\left(\perp_{1}\right), R(1), S\left(\perp_{2}, \perp_{2}\right)\right\}$.
We have $\operatorname{Supp}\left(Q, D, \perp_{2}\right)=\left\{v \in \mathrm{~V}(D) \mid v\left(\perp_{2}\right)=1\right.$ or $\left.v\left(\perp_{1}\right)=v\left(\perp_{2}\right)\right\}$, $\operatorname{Supp}(Q, D, 1)=\left\{v \in \mathrm{~V}(D) \mid v\left(\perp_{2}\right)=1\right\}$ and $\operatorname{Supp}\left(Q, D, \perp_{1}\right)=\{v \in \mathrm{~V}(D) \mid$ $\left.v\left(\perp_{1}\right)=v\left(\perp_{2}\right)\right\}$.

It follows that $\operatorname{cert}(Q, D)=\emptyset$ and $\operatorname{best}(Q, D)=\left\{\left(\perp_{2}\right)\right\}$.

### 2.3 Naïve evaluation and certain answers

For a query $Q$ written in FO or Datalog, we write $Q(D)$ to mean that such a query is evaluated naïve ly. That is, if $D$ contains nulls, nulls of $D$ are treated as new constants in the domain of $D$, distinct from each other, and distinct from all the other constants in $D$ and $\varphi$. For example, the query $\varphi(x, y)=\exists z(R(x, z) \wedge R(z, y))$, on the database $D=\left\{R\left(1, \perp_{1}\right), R\left(\perp_{1}, \perp_{2}\right), R\left(\perp_{3}, 2\right)\right\}$ selects only the tuple $\left(1, \perp_{2}\right)$.

There are known connections between naïve evaluation and certain answers. If $\Sigma$ is empty and $Q$ is a union of conjunctive queries, then $\operatorname{cert}_{\Sigma}(Q, D)=Q(D)$, see Imielinski and Lipski (1984). If $\Sigma$ contains a set of EGDs, then $\operatorname{cert}_{\Sigma}(Q, D)=Q\left(\operatorname{chase}_{\Sigma}(D)\right)$; cf. Greco et al. (2012). Here chase ${ }_{\Sigma}$ refers to the standard chase procedure with a set of EGDs, see Abiteboul et al. (1995).

### 2.4 Query languages

Here we shall study best and certain answers to first-order (FO) queries, possibly in the presence of constraints, by means of their rewriting in FO and Datalog. FO queries of vocabulary $\sigma$ use atomic relational and equality formulae and are closed under Boolean connectives $\wedge, \vee, \neg$ and quantifiers $\exists, \forall$. We write $\varphi(\bar{x})$ for an FO-formula $\varphi$ with free variables $\bar{x}$. With slight abuse of notation, $\bar{x}$ will denote both a tuple of variables and the set of variables occurring in it. The set of constants used by $\varphi$ is denoted by adom $(\varphi)$. We interpret FO-formulas under active domain semantics, that is quantified variables range over $\operatorname{adom}(D) \cup \operatorname{adom}(\varphi)$. Thus, an FO-formula $\varphi(\bar{x})$ represents a query (of active domain $\operatorname{adom}(\varphi))$ mapping each database $D$ into the set of tuples $\{\bar{t}$ over $\operatorname{adom}(D) \cup$ $\operatorname{adom}(\varphi) \mid D \models \varphi(\bar{t})\}$.

To evaluate FO-formulas with free variables, we use assignments $\nu$ from variables to constants in the active domain. Note that with a little abuse of notation, we write $D \models$ $\varphi(\bar{t})$ for $D \models_{\nu} \varphi(\bar{x})$ under the assignment $\nu$ sending $\bar{x}$ to $\bar{t}$.

Here it is important to note that the query associated with $\varphi$ is a mapping defined on all databases $D$, possibly with nulls. If $D$ contains nulls, $D \models \varphi(\bar{t})$ is to be intended "naïvely," that is nulls of $D$ are treated as new constants in the domain of $D$, distinct from each other, and distinct from all the other constants in $D$ and $\varphi$. For example, the query $\varphi(x, y)=\exists z(R(x, z) \wedge R(z, y))$, on the database $D=\left\{R\left(1, \perp_{1}\right), R\left(\perp_{1}, \perp_{2}\right), R\left(\perp_{3}, 2\right)\right\}$ selects only the tuple $\left(1, \perp_{2}\right)$.

Conjunctive queries (CQs) are given by the $\exists, \wedge$-fragment of FO, and their unions (UCQs) by the $\exists, \wedge, \vee$-fragment of FO; these are also captured by the positive fragment of relational algebra (select-project-union-join queries).

We also consider Boolean combination of conjunctive queries (BCCQs), that is, the closure of conjunctive queries under operations $q \cap q^{\prime}, q \cup q^{\prime}$, and $q-q^{\prime}$.

A Datalog rule (Abiteboul et al. 1995) is an expression of the form $R_{1}\left(u_{1}\right) \leftarrow$ $R_{2}\left(u_{2}\right), \ldots, R_{n}\left(u_{n}\right)$ where $n \geq 1, R_{1}, \ldots, R_{n}$ are relation names and $u_{1}, \ldots, u_{n}$ are tuples of appropriate arities. Each variable occurring in $u_{1}$ must occur in at least one of $u_{2}, \ldots, u_{n}$. A Datalog program is a finite set of Datalog rules. The head of the rule is the expression $R_{1}\left(u_{1}\right)$; and $R_{2}\left(u_{2}\right), \ldots, R_{n}\left(u_{n}\right)$ forms the body. The semantics is the standard fixed-point semantics.

As the language of our rewritings, we shall be using FO, but also a fragment of stratified Datalog with negation in bodies that can be seen in two different ways.

1. A program is evaluated in two steps. First, we can have a Datalog program $P$ defining new idb predicates $S_{1}, \ldots, S_{\ell}$. Then, we ask an FO query over the schema extended with these predicates $S_{1}, \ldots, S_{\ell}$.
2. We evaluate a stratified Datalog with negation program in which the first stratum has no negation (but may have recursion) and the second stratum has no recursion (but may have negation).

From the rewritings, we produce it will be clear that they fall in these classes. The key point about them is that they can be implemented in recursive SQL and that they both have PTIME data complexity, making their evaluation feasible. Note that recursive SQL as it is currently implemented, for example, in PostgreSQL 8.4, is actually Turing complete (Gierth 2011; Coelho 2013).

### 2.5 Complexity classes

In order to study the complexity of best and certain answer computation, we shall need two classes in the second level of the polynomial hierarchy. Both of these contain NP and coNP and are contained in $\Sigma_{2}^{p} \cap \Pi_{2}^{p}$. The class DP consists of languages $L_{1} \cap L_{2}$ where $L_{1} \in$ NP and $L_{2} \in$ CoNP. This class has appeared in database applications (Fagin et al. 2005; Barceló et al. 2014). The class $\mathrm{P}^{\mathrm{NP}[\log n]}$ consists of problems that can be solved in polynomial time with a logarithmic number of calls to an NP oracle (Buss and Hay 1991). Equivalently, it can be described as the class of problems solved in P with an NP oracle where calls to the oracle are done in parallel, that is, independent of each other. This class has appeared in the context of AI, modal logic, OBDA (Gottlob 1995; Eiter and Gottlob 1997; Calvanese et al. 2006; Bienvenu and Bourgaux 2016), data exchange (Arenas et al. 2013).

## 3 Complexity of best and certain answers

We start by looking at complexity of certain and best answers of first-order queries and answer a few questions that are (perhaps somewhat surprisingly) missing in the literature. In this case, we look at arbitrary first-order queries; thus, we do not mention constraints since $\operatorname{cert}_{\Sigma}(Q, D)=\operatorname{cert}(\Sigma \rightarrow Q, D)$ and likewise for best answers. In the subsequent sections, when we consider rewritings for sublanguages of first order, we shall again mention constraints explicitly since queries of the form $\Sigma \rightarrow Q$ will normally not belong to the same syntactic class as $Q$ itself. In this context, we will refer to $\triangle \operatorname{ANSWER}_{\Sigma}(Q)$.
As is common in database theory, we look at complexity in terms of complexity classes, which necessitates looking at decision versions of problems. The most common one that is found, stated here for $\triangle \in\{$ Certain, Best $\}$, is the following problem:

Problem: $\triangle \operatorname{Answer}(Q)$
Input: A database $D$, a tuple $\bar{a}$
Question: Is $\bar{a} \in \triangle(Q, D)$ ?

We are thus interested in data complexity: the query is fixed. We do not study combined complexity in this paper. In the remainder, we thus often omit the query and write $\triangle$ Answer instead of $\triangle \operatorname{AnSWER}(Q)$. Recall that in this case, for a language $L$, we say that the problem $\triangle$ Answer is $C$-complete in data complexity for a complexity class $C$, if $\triangle \operatorname{Answer}(Q)$ is solvable in $C$ for every $Q \in L$, and there exists a specific $Q_{0} \in L$ so that $\triangle \operatorname{Answer}\left(Q_{0}\right)$ is hard for $C$. We know from Abiteboul et al.
(1991) that CertainAnswer $(Q)$ is conP-complete in data complexity for first-order queries.

For best answers, it is a different version of the decision problem for which the complexity is known. Specifically, Libkin (2018) considered the problem of checking whether the set $\triangle(Q, D)$ belongs to a specified family of sets:

$$
\begin{array}{ll}
\text { Problem: } & \triangle \text { Answer }^{\epsilon}(Q) \\
\text { Input: } & \text { A database } D, \text { a family } \mathcal{X} \text { of sets of tuples } \\
\text { Question: } & \text { Is } \triangle(Q, D) \in \mathcal{X} ?
\end{array}
$$

For this decision version, the complexity of the problem was shown to be $\mathrm{P}^{\mathrm{NP}[\log n]}$ complete. This version looks a bit artificial, but we include it for the sake of completeness, because it has appeared in the literature.

However, this presentation of a decision suggests another rather natural presentation of a decision version, namely asking if a given set is $\triangle(Q, D)$ :

Problem: $\triangle$ Answer $^{=}(Q)$
Input: A database $D$, a set $X$ of tuples
Question: Is $\triangle(Q, D)=X$ ?

Our current state of knowledge is the complexity of CERTAINANSWER (CONPcomplete) and BESTANSWER ${ }^{\in}\left(\mathrm{P}^{\mathrm{NP}[\log n]}\right.$-complete). Thus, we now fill the gap and classify complexities of all the problems - for data complexity - in the case of FO queries.

We start by showing that all the alternatives for best answers - BESTANSWER ${ }_{\Sigma}$, BestAnswer ${ }^{=}$, and BestAnswer ${ }^{\epsilon}$ - are computationally equivalent.

## Theorem 3.1

For FO queries, the problems BestAnswere, BestAnswer $^{\epsilon}$, and BestAnswer ${ }^{=}$are $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete in data complexity.

## Proof

The upper bound for BestAnswer ${ }^{=}$immediately follows from the upper bound for BestAnswer $^{\in}$ (take the family $\mathcal{X}$ to be a singleton $\{X\}$ ). As for BestAnswer ${ }_{\Sigma}$, we only need a slight modification of the upper bound proof in Libkin (2018). To check whether $\bar{a} \in \operatorname{best}(Q, D)$ we proceed as follows. Since the query is fixed, and has therefore fixed arity $k$, in polynomial time we can enumerate all the $k$-tuples of adom $(D)$. Then, using parallel calls to the NP oracle, we can check for each such tuple $\bar{b}$ whether $\operatorname{Supp}(Q, D, \bar{a}) \subseteq \operatorname{Supp}(Q, D, \bar{b})$ and whether $\operatorname{Supp}(Q, D, \bar{b}) \subseteq \operatorname{Supp}(Q, D, \bar{a})$. With this information, in polynomial time we know whether $\bar{a} \triangleleft_{Q, D} \bar{b}$ for some $\bar{b}$.

Assuming $\Sigma$ empty, we prove the two remaining lower bounds, reducing from the same $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete problem (Wagner 1990): given an undirected graph $G$, is its chromatic number $\chi(G)$ odd? With each undirected graph $G=\langle N, E\rangle$ with nodes $N$ and edges $E$, we associate a database $D_{G}$ over binary relations $L, E$ and unary relations $C, O$ as
follows. We use a null $\perp_{n}$ in $D_{G}$ for each node $n$ of $G$. For each edge $\left\{n, n^{\prime}\right\}$ of $G$, we have pairs $\left(\perp_{n}, \perp_{n^{\prime}}\right)$ and $\left(\perp_{n^{\prime}}, \perp_{n}\right)$ in the relation $E$ of $D_{G}$. In relation $C$, we have $m$ constants $\left\{c_{1}, \ldots, c_{m}\right\}$ (intuitively representing possible colors), where $m$ is the number of nodes of $G$. Relation $O$ of $D_{G}$ is defined as $\left\{c_{i} \mid i\right.$ is odd $\}$, and $L$ is a linear ordering on them, that is, $\left(c_{i}, c_{j}\right) \in L$ iff $i \leq j$, for $i, j \leq m$.

Remark that any valuation $v$ of $D_{G}$ that maps each null into a constant of $C$ represents an assignment of colors in $\left\{c_{1}, \ldots, c_{m}\right\}$ to nodes of $G$. Then, we define a query

$$
\phi(x)=C(x) \wedge \forall y, z(E(y, z) \rightarrow L(y, x)) \wedge \forall y(L(y, x) \rightarrow \exists z E(y, z)) \wedge \neg \exists y E(y, y)
$$

For any valuation $v, \phi(c)$ holds in $v\left(D_{G}\right)$ iff (1) $c=c_{j}$ for some $j=1 . . m$ (ensured by the first conjunct). (2) For such a $c_{j}$, the valuation $v$ maps each null into $\left\{c_{1}, \ldots, c_{j}\right\}$ (second conjunct), that is $v$ represents an assignment of colors to nodes of $G$, using at most the first j colors. (3) Each color $\left\{c_{1}, \ldots, c_{j}\right\}$ is used by $v$, that is $v$ represents an assignment of colors to nodes of $G$, using precisely the first j colors (third conjunct). (4) There are no loops in $E$ (fourth conjunct).

Thus, for a valuation $v$, the formula $\phi\left(c_{j}\right)$ is true in $v\left(D_{G}\right)$ iff $v$ represents a coloring of G using precisely the first j colors $\left\{c_{1}, \ldots, c_{j}\right\}$ (which in the sequel we refer to as an exact $j$-coloring of $G$ ).

Next, we define:

$$
Q(x)=C(x) \wedge(\phi(x) \vee \exists y(O(y) \wedge L(x, y) \wedge \phi(y)))
$$

For a valuation $v$, we have that $Q\left(c_{i}\right)$ holds in $v\left(D_{G}\right)$ iff either $v$ represents an exact $i$-coloring of $G$; or $v$ represents an exact $j$-coloring of $G$ with $j$ odd, and $i \leq j$. In other words, valuations representing exact $j$-colorings, with $j$ even, support only the maximal color $c_{j}$; while valuations representing exact $j$-colorings, with $j$ odd, support all colors $\left\{c_{1} \ldots c_{j}\right\}$.

With this in place, we can conclude the reduction for the BESTANSWER ${ }_{\Sigma}$ problem:
Claim. $c_{1} \in \operatorname{best}\left(Q, D_{G}\right)$ iff the chromatic number of $G$ is odd.
First, we prove the above claim. Let $\chi_{G}$ be the chromatic number of $G$. Then, there exist no exact colorings of $G$ which are prefixes of $\left\{c_{1}, \ldots c_{\chi_{G}}\right\}$, while $\left\{c_{1}, \ldots c_{\chi_{G}}\right\}$ is an exact coloring of $G$.

Assume first that $\chi_{G}$ is even. Then, there exist no valuations representing the exact coloring $\left\{c_{1}\right\}$. Thus, the support of $c_{1}$ is the set of valuation representing an exact coloring $\left\{c_{1} \ldots c_{j}\right\}$ of $G$ with $j$ odd and $j>\chi_{G}$. This support is not maximal. In fact, the support of $c_{\chi_{G}}$ is:

- the valuations representing the exact coloring $\left\{c_{1} \ldots c_{\chi_{G}}\right\}$ (there exists at least one);
- the valuations representing an exact coloring $\left\{c_{1} \ldots c_{j}\right\}$ of $G$ with $j$ odd and $j>\chi_{G}$.

This support strictly contains the support of $c_{1}$; in fact valuations in the first item cannot be also in the second.

Assume now that $\chi_{G}$ is odd. Then, the support of $c_{1}$ is the set of valuations representing an exact coloring $\left\{c_{1} \ldots c_{j}\right\}$ of $G$ with j odd and $j \geq \chi_{G}$. We show that this set is maximal; that is, no color $c_{k}$ can have a support strictly containing it.

- if $k \leq \chi_{G}$ then the support of $c_{k}$ is the set of valuations representing an exact coloring $\left\{c_{1} \ldots c_{j}\right\}$ of $G$ with j odd, and $j \geq \chi_{G}$. So same support as $c_{1}$.
- if $k>\chi_{G}$, the support of $c_{k}$ cannot contain the valuations representing $\left\{c_{1}, \ldots c_{\chi_{G}}\right\}$. There exists at least one such valuation and it belongs to the support of $c_{1}$. Thus, the support of $c_{k}$ does not contain the support of $c_{1}$.

We now move to BestAnswer ${ }^{=}$. With any undirected graph $G$, we associate a relational structure $D_{G}^{\prime}$ obtained from $D_{G}$ by adding a new color $c_{0}$ in $C$ with $L\left(c_{0}, c_{i}\right)$ for every $0 \leq i \leq m$. We define a restriction $\psi$ of the original formula $\phi$ by disallowing $c_{0}$ in colorings: to obtain $\psi$ it suffices to replace $L(y, x)$ in $\phi$ by $L(y, x) \wedge y \neq c_{0}$, and $C(x)$ by $C(x) \wedge x \neq c_{0}$. Thus, it is still true that $\psi\left(c_{j}\right)$ is true in $v\left(D_{G}^{\prime}\right)$ iff $v$ represents a coloring of G using precisely $\left\{c_{1}, \ldots, c_{j}\right\}$.

We define a new query:

$$
\begin{aligned}
Q^{\prime}(x):= & O(x) \wedge(\psi(x) \vee \exists y(O(y) \wedge L(x, y) \wedge \psi(y)) \\
& \vee \\
& \neg O(x) \wedge(\psi(x) \vee \exists y(O(y) \wedge x+2<y \wedge \psi(y)) \\
& \vee \\
& \neg O(x) \wedge \exists y(x \neq y \wedge L(x, y)) \wedge \forall y \forall z\left(E(y, z) \rightarrow\left(y=c_{0} \wedge z=c_{0}\right)\right)
\end{aligned}
$$

Note that $x+2<y$ is used as a shorthand, as it is definable in our language. $Q^{\prime}\left(c_{i}\right)$ holds in $v\left(D_{G}^{\prime}\right)$ iff

- either $i$ is odd and $v$ represents an exact $j$-coloring of $G$, with $j$ odd and $i \leq j$;
- or $i$ is even and:
- either $v$ represents an exact coloring $\left\{c_{1} \ldots c_{j}\right\}$ of $G$ with j odd, and $i+2<j$;
- or $v$ represents an exact coloring $\left\{c_{1} \ldots c_{i}\right\}$ of $G$;
- or $i<m$ and $v\left(\perp_{j}\right)=c_{0}$ for all $1 \leq j \leq m$;

The following claim completes the reduction for BEStANSWER ${ }^{=}$:
Claim. $\left\{c_{i} \mid i\right.$ is even $\}=\operatorname{best}\left(Q^{\prime}, D_{G}^{\prime}\right)$ iff $\chi(G)$ is even.
In the following, we call $v_{0}$ the unique valuation such that $v_{0}\left(\perp_{j}\right)=c_{0}$ for all $1 \leq$ $j \leq m$. First assume that $\chi_{G}$ is even. For all $0<i \leq m \operatorname{odd}$, $\operatorname{Supp}\left(c_{i}\right)$ is not maximal as $\operatorname{Supp}\left(c_{0}\right) \supset \operatorname{Supp}\left(c_{i}\right) \cup\left\{v_{0}\right\}$. Hence, $\operatorname{best}\left(Q^{\prime}, D_{G}^{\prime}\right) \subseteq\left\{c_{i} \mid i\right.$ is even $\}$, so we show $\left\{c_{i} \mid i\right.$ is even $\} \subseteq \operatorname{best}\left(Q^{\prime}, D_{G}^{\prime}\right)$. The inclusion holds whenever $c_{i} \geq \chi(G)$, as $\operatorname{Supp}\left(c_{i}\right)$ contains all valuations representing exact colorings $\left\{c_{1} \ldots c_{i}\right\}$ of G , while no other $\operatorname{Supp}\left(c_{j}\right)$ with $i \neq j$ contains them. Now take $c_{i}<\chi(G)$ with $i$ even, then $\operatorname{Supp}\left(c_{i}\right)$ contains $v_{0}$ together with all exact odd colorings (if there are any). First assume that there exists odd exact colorings of $G$, so there are $\chi(G)+1$ ones and valuations representing them are not contained in $\operatorname{Supp}(\chi(G))$. Also, $v_{0} \notin \operatorname{Supp}\left(c_{k}\right)$ with $k$ odd and $k<\chi(G)$. It follows that $\operatorname{Supp}\left(c_{i}\right)$, which is the union of $v_{0}$ and of all valuations representing odd exact colorings, is maximal. Now assume that there is no exact odd coloring. This corresponds to the special case $\chi(G)=m$ where $\operatorname{Supp}\left(c_{m}\right)$ contains only the exact colorings $\left\{c_{1} \ldots c_{m}\right\}$ of $G$, but not $v_{0}$; while $\operatorname{Supp}\left(c_{j}\right)=\emptyset$ whenever $j$ odd. In such a case, $\operatorname{Supp}\left(c_{i}\right)=\left\{v_{0}\right\}$ is also maximal.

We assume now $\chi(G)$ is odd and show $\left\{c_{i} \mid i\right.$ is even $\} \neq \operatorname{best}\left(Q^{\prime}, D_{G}^{\prime}\right)$. First notice that $\operatorname{Supp}\left(c_{1}\right)$ is maximal whenever $\chi(G)=1$, as neither $\operatorname{Supp}\left(c_{0}\right)$ nor any $\operatorname{Supp}\left(c_{i}\right)$ with

| Type of problem | $\bar{a} \in$ Answer | $X=$ Answer | Answer $\in \mathcal{X}$ |
| :--- | :---: | :---: | :---: |
| Certain Answer | coNP-complete ${ }^{1}$ | DP-complete | $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete |
| Best Answer | $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete | $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete | $\mathrm{P}^{\mathrm{NP}[\log n]}$-complete ${ }^{2}$ |

Fig. 1. Summary of data complexity results for FO queries.
${ }^{1}$ Abiteboul et al. (1991); ${ }^{2}$ Libkin (2018).
$i \geq 2$ contain valuations representing the exact $\left\{c_{1}\right\}$ colorings. So we assume $\chi(G) \geq 3$, from which it follows that there exists a constant $c_{\chi(G)-3}$ in the active domain which support contains $v_{0}$ together with all valuations representing exact odd colorings. As $\operatorname{Supp}\left(c_{\chi(G)-1}\right)$ contains exactly the same set of valuations, to the exclusion of those representing $\left\{c_{1} \ldots c_{\chi(G)}\right\}$ colorings, it follows that $\operatorname{Supp}\left(c_{\chi(G)-3}\right) \supset \operatorname{Supp}\left(c_{\chi(G)-1}\right)$ and so $c_{\chi(g)-1} \notin \operatorname{best}\left(Q^{\prime}, D_{G}^{\prime}\right)$.

Now that we showed that all three formulations of best answers actually collapse computationally, another natural question arises. Does a similar result hold for certain answers ? It is well known that data complexity of CertainAnswer ${ }_{\Sigma}$ is coNP-complete for FO queries (Abiteboul et al. 1991). We complete the picture as follows and summarize results in Figure 1.

## Theorem 3.2

For FO queries, CertainAnswer $=$ is $D P$-complete and CertainAnswer ${ }^{\epsilon}$ is $\mathrm{P}^{\mathrm{NP}[\log n]}$ complete in data complexity.

## Proof

To prove membership of CertainAnswer $=$ in $D P$, notice that for a query $Q$, this problem is the intersection of two languages $L_{1} \cap L_{2}$ where $L_{1}=\{(D, X) \mid X \subseteq \operatorname{cert}(Q, D)\}$ and $L_{2}=\{(D, X) \mid \bar{X} \subseteq \overline{\operatorname{cert}(Q, D)}\}$. $L_{1}$ is known to be in CoNP : we guess a tuple $\bar{a} \in X$ and a valuation $v \in V(D)$ with $v(\bar{a}) \notin Q(v(D))$. Similarly, $L_{2}$ is in NP : we guess a tuple $\bar{b} \in \bar{X}$ and a valuation $v^{\prime} \in V(D)$ with $v^{\prime}(\bar{b}) \in Q(v(D))$.

To prove membership of CertainAnswer ${ }^{\in}$ in $\mathrm{P}^{\mathrm{NP}[\log n]}$, suppose the query $Q$ is k-ary, and we are given a family of sets of k-ary tuples $\mathcal{X}=\left\{X_{1}, \ldots, X_{n}\right\}$ and a database $D$. For each $X_{i} \in \mathcal{X}$, we use the NP oracle to decide in parallel whether $X_{i}=\operatorname{cert}(Q, D)$ (for each $X_{i}$, the two calls to the oracle do not depend on each other and they can also be done in parallel).

For $D P$-hardness, we reduce from the problem of checking whether $\chi(G)$, the chromatic number of an undirected graph $G$, equals 4 (Rothe 2003) and for $\mathrm{P}^{\mathrm{NP}[\log n]}$-hardness, we reduce from the related problem of checking whether $\chi(G)$ is odd. With such a graph $G$, we associate the same database $D_{G}$ as in the proof of Theorem 3.1. Using the exact coloring formula $\varphi$ in the proof of Theorem 3.1, we define a query

$$
Q(x):=C(x) \wedge \forall y(\varphi(y) \rightarrow L(x, y))
$$

We claim that $\operatorname{cert}(Q, D)=\left\{c_{1}, \ldots, c_{n}\right\}$ iff $\chi(G)=n$, which entails $\operatorname{cert}(Q, D)=$ $\left\{c_{1}, \ldots, c_{4}\right\}$ iff $\chi(G)=4$ and $\operatorname{cert}(Q, D) \in\left\{\left\{c_{1}, \ldots, c_{j}\right\} \mid j\right.$ is odd and $\left.1 \leq j \leq|G|\right\}$ iff $\chi(G)$ is odd. Recall that $v\left(D_{G}\right) \models \varphi\left(c_{i}\right)$ iff $c_{i}$ is a color in $\left\{c_{1}, \ldots, c_{|G|}\right\}$ and $v$ represents
an exact $i$-coloring of the graph. Now $v\left(D_{G}\right) \models Q\left(c_{j}\right)$ iff $c_{j}$ is a color and there is no $i<j$ such that $v$ represents an exact $i$-coloring of the graph, which holds exactly whenever $c_{j} \in\left\{c_{1}, \ldots, c_{\chi(G)}\right\}$.

## 4 Query rewritings for tractable fragments

Considering arbitrary FO queries brought us an intrinsic intractability result for all variants of the considered decision problems. This motivates restricting to well-behaved fragments such as CQs and UCQs. Recall that conjunctive queries (CQs) are given by the $\exists, \wedge$-fragment of FO, and their unions (UCQs) by the $\exists, \wedge, \vee$-fragment of FO. We extend them with a mild form of negation (since adding negation leads to coNP-hardness of certain answers). This mild form comes in the shape of Boolean combination of conjunctive queries ( $B C C Q s$ ), that is, the closure of conjunctive queries under operations $q \cap q^{\prime}$, $q \cup q^{\prime}$, and $q-q^{\prime}$.

If there are no constraints in $\Sigma$, finding certain answers to BCCQs is known to be tractable (Gheerbrant and Libkin 2015), though by tableau-based techniques that are hard to implement in a database system. We now extend this in two ways. First, we show that tractability is preserved even in the presence of EGDs (and thus functional dependencies and keys). Second, we show that certain answers can be obtained by rewriting into a fragment of Datalog as described in Section 2. In particular, it means that certain answers can be found by a query expressible in recursive SQL (and even in SQL in the absence of constraints).

For Best Answere a polynomial-time evaluation algorithm (in data complexity) already exists (Libkin 2018). The resolution-based procedure is however in sharp contrast with naïve evaluation, which allows to compute certain answers to unions of conjunctive queries via usual model checking. We thus show how to apply our query rewriting techniques to the best answers problem.

### 4.1 A normal form for queries: neutralizing variable repetition

Towards our rewritings, we start by putting each conjunctive query in a normal form which eliminates repetition of variables, by introducing new equality atoms.

Definition 4.1 (NRV normal form)
A conjunctive query $Q$ is in non-repeating variable normal form (NRV normal form) whenever it is of the form $Q(\bar{x})=\exists \bar{w}(q(\bar{w}) \wedge e(\bar{x}, \bar{w}))$ where variables in $\bar{x} \bar{w}$ are pairwise distinct, and:

- $q(\bar{w})$ is a conjunction of relational atoms without constants, where each free variable in $\bar{w}$ has at most one occurrence in $q$,
- $e(\bar{x}, \bar{w})$ is a conjunction of equality atoms, possibly using constants, where each variable of $\bar{x}$ is involved in at least one equality.

We say that $q(\bar{w})$ is the relational subquery of $Q$, and $e(\bar{x}, \bar{w})$ is the equality subquery of $Q$. A BCCQ is in NRV normal form if it is a Boolean combination of CQs in NRV normal form.

## Example 4.2

The query $Q(x)$ from Example 2.1 is equivalent to $\exists w_{1} w_{2} w_{3}\left(R\left(w_{1}\right) \wedge S\left(w_{2}, w_{3}\right) \wedge w_{1}=\right.$ $w_{2} \wedge w_{3}=x$ ), which is in NRV normal form.

Clearly every conjunctive query $Q$ is equivalent to a query in NRV normal form; moreover, $Q$ can be easily rewritten in NRV normal form (in linear time in the size of the query). Thus in what follows, we assume w.l.o.g. that conjunctive queries are given in NRV normal form. Intuitively, the NRV normal form allows us to separate the two ingredients of a conjunctive query: the existence of facts in some relations of the database, on the one side, and a set of equality conditions on data values occurring in these facts, on the other side. The existence of facts does not depend on the valuation of nulls and thus can be directly tested on the incomplete database. Instead equality atoms in an NRV normal form imply conditions that valuations need to satisfy in order for the query to hold. We can thus first concentrate on the support of equality subqueries. This will be encoded in FO and then integrated in the rewriting of the whole conjunctive query.

We introduce a notion of equivalence of database elements w.r.t. to a set of equalities. Intuitively, equivalent elements of a tuple $\bar{t}$ are the ones which should be collapsed into a single value in order for a valuation of $\bar{t}$ to satisfy all the given equalities.

## Definition 4.3

Given a database $D$, a conjunction of equality atoms $\gamma(\bar{y})$ and an assignment $\nu: \bar{y} \cup$ $\operatorname{adom}(\gamma) \rightarrow \operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ preserving constants, we say that $u, u^{\prime} \in \operatorname{adom}(D) \cup$ $\operatorname{adom}(\gamma)$ are equivalent w.r.t. $\gamma$ and $\nu$ and write $u \equiv_{\gamma}^{\nu} u^{\prime}$, if either $u=u^{\prime}$ or $\left(u, u^{\prime}\right)$ belongs to the reflexive symmetric transitive closure of $\{(\nu(x), \nu(w)) \mid x=w \in \gamma\}$.

The relation $\equiv_{\gamma}^{\nu}$ is clearly an equivalence relation over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$, where each element outside the range of $\nu$ forms a singleton equivalence class.

## Example 4.4

Let $\gamma$ be $x_{1}=x_{2} \wedge x_{2}=x_{3} \wedge x_{4}=x_{5} \wedge x_{6}=1$. Let $\nu$ assign $\perp_{i}$ to $x_{i}$ for $i \leq 5$, and $\perp_{5}$ to $x_{6}$. The equivalence classes of $\equiv_{\gamma}^{\nu}$ are $\left\{\perp_{i} \mid i \leq 3\right\}$ and $\left\{1, \perp_{4}, \perp_{5}\right\}$, plus one singleton for each other element of the active domain.

In what follows, we denote by $\sim_{\gamma}$ the reflexive symmetric transitive closure of $\{(x, w) \mid x=w \in \gamma\}$. Note that this is an equivalence relation among variables and constants of $\gamma$. We will provide two syntactic encodings of this relation, one in Datalog and one in FO.

### 4.2 Datalog rewriting for certain answers for BCCQs with EGDs

Recall that, given a query $Q$, a database $D$, and a tuple $\bar{a} \operatorname{over} \operatorname{adom}(D) \cup \operatorname{adom}(Q)$ we let the support of $\bar{a}$ be the set of all valuations that witness it:

$$
\operatorname{Supp}(Q, D, \bar{a})=\{v \in \mathrm{~V}(D) \mid v(\bar{a}) \in Q(v(D))\}
$$

In order to look for rewritings of BCCQs, a key observation is that $\bar{a}$ is a certain answer to $Q$ iff $\operatorname{Supp}(\neg Q, D, \bar{a})=\emptyset$. When $Q$ is a BCCQ, so is $\neg Q$, thus we look for ways of expressing (non-)emptiness of the support for BCCQs.

We start by concentrating on the support of equality subqueries. This will be encoded in Datalog and then integrated, as a key ingredient, in the rewriting of the whole query. We
let $\gamma(\bar{y})$ be an arbitrary set of equality atoms among variables $\bar{y}$ and possibly constants. Intuitively, we will be interested in the case that $\gamma(\bar{y})$ is the equality subquery $e(\bar{x}, \bar{w})$ of a CQ in NRV normal form (thus notice that in the Datalog program below $\bar{y}$ encompasses variables $\bar{x} \bar{w}$ of an equality subquery).

Remark that we can always write an EGD so that no variable in its body occurs more than once; it suffices to add to the body a set of variable equalities. Thus, we assume that EGDs in $\Sigma$ are of the form $\forall \bar{u}\left((\varphi(\bar{u}) \wedge \psi) \rightarrow z=z^{\prime}\right)$ where $z, z^{\prime}$ are in $\bar{u}$, the conjunction of atoms $\varphi(\bar{u})$ contains no constants, no variable occurs twice in $\varphi(\bar{u})$, and $\psi$ is a set of equalities among variables of $\bar{u}$. Remark also that membership in the set $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ can be expressed by a UCQ formula that we call $\operatorname{Dom}(x)$. We encode equivalence of database elements in $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ w.r.t. a set of equalities $\gamma(\bar{y})$ using the following Datalog program ${ }^{1}$ :

$$
\begin{aligned}
& \operatorname{equiv}_{\gamma}(\bar{y}, z, z) \leftarrow \wedge_{i} \operatorname{Dom}\left(y_{i}\right), \operatorname{Dom}(z) \\
& \operatorname{equiv}_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \leftarrow z=y_{k}, z^{\prime}=y_{l}, \wedge_{i} \operatorname{Dom}\left(y_{i}\right) \text { for each }\left(y_{k}=y_{l}\right) \in \gamma \\
& \operatorname{equiv}_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \leftarrow \operatorname{equiv}_{\gamma}(\bar{y}, z, u), \text { equiv }_{\gamma}\left(\bar{y}, u, z^{\prime}\right) \\
& \text { equiv }_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \leftarrow \operatorname{equiv}_{\gamma}\left(\bar{y}, z^{\prime}, z\right) \\
& \text { equiv }_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \leftarrow \varphi(\bar{u}) \wedge_{\left(w=w^{\prime}\right) \in \psi(\bar{u})} \text { equiv }_{\gamma}\left(\bar{y}, w, w^{\prime}\right) \\
& \quad \text { for each } \operatorname{EGD} \forall \bar{u}\left((\varphi(\bar{u}) \wedge \psi(\bar{u})) \rightarrow z=z^{\prime}\right) \in \Sigma
\end{aligned}
$$

Intuitively, if $\bar{t}$ is a tuple of database elements assigned to $\bar{y}$, equivalent elements of $D$ are the ones which should be collapsed into a single value in order for a valuation of $D$ to satisfy all the equalities $\gamma(\bar{t})$ and the EGDs. For fixed $\gamma$ and $\bar{t}$, the relation $\left\{\left(s, s^{\prime}\right) \mid D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)\right\}$ is an equivalence relation over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ where each element of $\operatorname{adom}(D)$ neither in $\bar{t}$ nor in adom $(\gamma)$ forms a singleton equivalence class.

The formula equiv is a key ingredient in our rewriting; as formalized in the following lemma, it selects precisely the pairs of elements that a consistent valuation needs to collapse to satisfy a set of equalities.

## Lemma 4.5

Let $\gamma(\bar{y})$ be a conjunction of equality atoms, $D$ a database, and $\nu(\bar{y})=\bar{t}$ an assignment over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$. Assume $v$ is a consistent valuation of nulls, then $v(D) \models \gamma(v(\bar{t}))$ if and only if $v(s)=v\left(s^{\prime}\right)$ for all $s, s^{\prime}$ such that $D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$.

Proof
$\Rightarrow$ Assume $v(D) \models \gamma(v(\bar{t}))$ and let $s, s^{\prime}$ such that $D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. We prove $v(s)=v\left(s^{\prime}\right)$. We proceed by induction on the derivation of equiv $\left(\bar{t}, s, s^{\prime}\right)$ by the fixpoint evaluation of the Datalog program. Assume $\operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ is derived at the first iteration, and then, it follows from one of the first two rules. If it is derived by the fist rule then $s=s^{\prime}$ and therefore $v(s)=v\left(s^{\prime}\right)$ trivially. Assume equiv $\left(\bar{t}, s, s^{\prime}\right)$ is derived using the second rule then there exists $\left(y_{k}=y_{l}\right) \in \gamma$, and $s=t_{k}$ and $s^{\prime}=t_{l}$; now since $v(D) \models \gamma(v(\bar{t}))$, we have $v\left(t_{k}\right)=v\left(t_{l}\right)$. Now assume that equiv $\left(\bar{t}, s, s^{\prime}\right)$ is derived at some subsequent step. If it follows from the second rule, then it follows from

[^0]$e q u i v_{\gamma}(\bar{t}, s, p)$ and $e q u i v_{\gamma}\left(\bar{t}, p, s^{\prime}\right)$ derived at previous steps, thus by the induction hypothesis $v(s)=v(p)=v\left(s^{\prime}\right)$. Similarly if equiv$\left(\bar{t}, s, s^{\prime}\right)$ is derived by the third rule, then it follows from equiv $\left(\bar{t}, s^{\prime}, s\right)$, and thus by induction $v(s)=v\left(s^{\prime}\right)$. The last case is that $e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ is derived by the last rule for some EGD $\left.(\varphi(\bar{u}) \wedge \psi(\bar{u})) \rightarrow z=z^{\prime}\right)$. So there exists a mapping $\mu$ of $\bar{u}$ such that $D \models \phi(\mu(\bar{u}))$ and $D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, \mu(w), \mu\left(w^{\prime}\right)\right)$ for each $\left.\left(w=w^{\prime}\right) \in \psi(\bar{u})\right)$ and $s=\mu(z)$ and $s^{\prime}=\mu\left(z^{\prime}\right)$; so by the induction hypothesis $v(\mu(w))=v\left(\mu\left(w^{\prime}\right)\right)$. It follows that $v(D) \models \phi(v(\mu(\bar{u}))) \wedge \psi(v(\mu(\bar{u})))$ and because $v(D)$ satisfies the EGDs ( $v$ being consistent), $v(\mu(z))=v\left(\mu\left(z^{\prime}\right)\right)$, thus $v(s)=v\left(s^{\prime}\right)$.
$\Leftarrow$ Assume $\forall s, s^{\prime}, D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ implies $v(s)=v\left(s^{\prime}\right)$. We show that $v(D) \models$ $\gamma(v(\bar{t}))$. In fact for each $\left(y_{k}=y_{l}\right) \in \gamma$, we have $D \vDash \operatorname{equiv}_{\gamma}\left(\bar{t}, t_{k}, t_{l}\right)$ (derived by the second rule). By our hypothesis $v\left(t_{k}\right)=v\left(t_{l}\right)$, thus $\gamma(v(\bar{t}))$ holds.

## Example 4.6

Let $\gamma$ and $\nu$ be as in Example 4.4, then Lemma 4.5 implies that a valuation $v(D) \models \gamma(v(\bar{t}))$ iff $v\left(\perp_{i}\right)=v\left(\perp_{j}\right)$ for all $i, j=1 . .3$, and $v\left(\perp_{i}\right)=1$ for all $i=4,5$.

Formulas we write in the remainder are over signature $\sigma \cup N u l l$, where $\sigma$ is the database schema. In any incomplete database $D$ over $\sigma \cup N u l l, N u l l$ is always interpreted by the set of nulls occurring in $D$ (in accordance with the semantics of the SQL construct IS NULL). That is we allow rewritings to test whether a database element is null or not.

For $\gamma(\bar{y})$ a conjunction of equality atoms, using equiv $_{\gamma}$ we define a new formula $\operatorname{comp}_{\gamma}(\bar{y})$ stating the existence of a consistent valuation that collapses all equivalent elements of a tuple:

$$
\operatorname{comp}_{\gamma}(\bar{y}):=\forall z z^{\prime}\left(e q u i v_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \wedge \neg \operatorname{Null}(z) \wedge \neg \operatorname{Null}\left(z^{\prime}\right) \rightarrow z=z^{\prime}\right)
$$

## Proposition 4.7

Let $\gamma(\bar{y})$ be a conjunction of equality atoms, $D$ a database, and $\nu(\bar{y})=\bar{t}$ an assignment over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$, then $D \models \operatorname{comp}_{\gamma}(\bar{t})$ if and only if there exists a consistent valuation $v$ of nulls such that $v(D) \models \gamma(v(\bar{t})$. Moreover if such valuation exists, there exists one further satisfying $v(s)=v\left(s^{\prime}\right)$ iff $D=e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$, for all $s, s^{\prime} \in \operatorname{adom}(D) \cup$ $\operatorname{adom}(\gamma)$.

## Proof

$\Rightarrow$ Assume $D \models \operatorname{comp}_{\gamma}(\bar{t})$. Then, $\forall c, c^{\prime}$ constants, $D \models e q u i v_{\gamma}\left(\bar{t}, c, c^{\prime}\right)$ implies $c=c^{\prime}$. As $\left\{\left(s, s^{\prime}\right) \mid D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right)\right\}$ is an equivalence relation over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$, its equivalence classes form a partition of this set. In each equivalence class, there is at most one constant, so we define a valuation $v$ mapping all nulls of a class to the unique constant of that class (or to a new fresh constant if the class does not contain any). Note that $v$ has the property that $v(s)=v\left(s^{\prime}\right)$ iff $D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$; this allows to prove that $v$ is a consistent valuation.

In fact, consider an arbitrary EGD $\left.(\varphi(\bar{u}) \wedge \psi(\bar{u})) \rightarrow z=z^{\prime}\right)$ in $\Sigma$ and assume $v(D) \models \varphi(\mu(\bar{u})) \wedge \psi(\mu(\bar{u}))$ for some $\mu$; we prove $\mu(z)=\mu\left(z^{\prime}\right)$. Since $\varphi$ is a conjunction of atoms with no constants and no repeated variables, there exists $\mu^{\prime}$ such that $D \vDash \varphi\left(\mu^{\prime}(\bar{u})\right)$ and $v\left(\mu^{\prime}(\bar{u})\right)=\mu(\bar{u})$. Take any equality $w=w^{\prime}$ in $\psi(\bar{u})$, we show that $D \vDash \operatorname{equiv}_{\gamma}\left(\bar{t}, \mu^{\prime}(w), \mu^{\prime}\left(w^{\prime}\right)\right)$. In fact because $v(D) \models \psi(\mu(\bar{u}))$ we have $\mu(w)=\mu\left(w^{\prime}\right)$,
and therefore $v\left(\mu^{\prime}(w)\right)=v\left(\mu^{\prime}\left(w^{\prime}\right)\right)$. By definition of $v$ then $D \vDash e q u i v_{\gamma}\left(\bar{t}, \mu^{\prime}(w), \mu^{\prime}\left(w^{\prime}\right)\right)$. By the last rule of the Datalog program defining equiv ${ }_{\gamma}$, we thus have that $D \models$ equiv $\left(\bar{t}, \mu^{\prime}(z), \mu^{\prime}\left(z^{\prime}\right)\right)$; then again by definition of $v$, we have $v\left(\mu^{\prime}(z)\right)=v\left(\mu^{\prime}\left(z^{\prime}\right)\right.$, and therefore $\mu(z)=\mu\left(z^{\prime}\right)$. This shows that $v(D)$ satisfies all the EGDs; thus, $v$ is consistent.

We have proved that $v$ satisfies the characterization of Lemma 4.5, and can conclude that $v(D) \models \gamma(v(\bar{t}))$.
$\Leftarrow$ Assume $v$ is a consistent valuation and $v(D) \models \gamma(v(\bar{t}))$. Let $s, s^{\prime}$ such that $D \models$ equiv $\left(\bar{t}, s, s^{\prime}\right) \wedge \neg \operatorname{Null}\left(s^{\prime}\right) \wedge \neg \operatorname{Null}(s)$. By Lemma 4.5 we have $v(s)=v\left(s^{\prime}\right)$. Moreover, $s, s^{\prime}$ are both constants, then $s=s^{\prime}$. Hence, $D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right) \wedge \neg N u l l\left(s^{\prime}\right) \wedge \neg N u l l(s) \rightarrow$ $s=s^{\prime}$, that is, $D \models \operatorname{comp}_{\gamma}(\bar{t})$.

We are now ready to define a formula capturing the inclusion of supports between two conjunctions of equality atoms, which will be a crucial ingredient in our rewriting. Let $\gamma(\bar{x})$ and $\gamma^{\prime}(\bar{y})$ be conjunctions of equality atoms with $\operatorname{adom}(\gamma)=\operatorname{adom}\left(\gamma^{\prime}\right)$. We define:

$$
i^{i m p l} y_{\gamma, \gamma^{\prime}}(\bar{x}, \bar{y}):=\forall z z^{\prime}\left(e q u i v_{\gamma^{\prime}}\left(\bar{y}, z, z^{\prime}\right) \rightarrow \text { equiv }_{\gamma}\left(\bar{x}, z, z^{\prime}\right)\right)
$$

Using Proposition 4.7 and Lemma 4.5, we obtain:

## Proposition 4.8

Let $\gamma(\bar{x}), \gamma^{\prime}(\bar{y})$ be conjunctions of equality atoms with $\operatorname{adom}(\gamma)=\operatorname{adom}\left(\gamma^{\prime}\right), D$ a database and $\nu(\bar{y})=\bar{t}, \nu^{\prime}(\bar{y})=\overline{t^{\prime}}$ assignments over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$. Then $D \models i m p l y_{\gamma, \gamma^{\prime}}\left(\bar{t}, \bar{t}^{\prime}\right) \vee$ $\neg \operatorname{comp}_{\gamma}(\bar{t})$ iff for all consistent valuations $v$, one has $v(D) \models \gamma(v(\bar{t}))$ implies $v(D) \models$ $\gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$.

## Proof

$\Rightarrow$ Assume $D \models \operatorname{imply}_{\gamma, \gamma^{\prime}}\left(\bar{t}, \bar{t}^{\prime}\right) \vee \neg \operatorname{comp}_{\gamma}(\bar{t})$. If $D \models \neg \operatorname{comp}_{\gamma}(\bar{t})$ then by Proposition 4.7, there is no consistent valuation $v$ such that $v(D) \models \gamma(v(\bar{t}))$ and so the implication trivially holds. Now assume $D \models$ imply $_{\gamma, \gamma^{\prime}}\left(\bar{t}, \bar{t}^{\prime}\right)$, that is:

$$
D \models \forall z z^{\prime}\left(\operatorname{equiv}_{\gamma^{\prime}}\left(\overline{t^{\prime}}, z, z^{\prime}\right) \rightarrow \operatorname{equiv}_{\gamma}\left(\bar{t}, z, z^{\prime}\right)\right)
$$

We want to show that for all consistent valuations $v$ of nulls such that $v(D) \vDash \gamma(v(\bar{t}))$, one also has $v(D) \models \gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$. Consider a valuation $v$ of nulls such that $v(D) \models \gamma(v(\bar{t}))$. Using Lemma 4.5, it is enough to show that $\forall s, s^{\prime}$ such that $D=e q u i v_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right)$ one has $v(s)=v\left(s^{\prime}\right)$. So assume $D \models$ equiv $_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right)$, by our assumption it follows that $\left.D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)\right)$, and therefore, again by Lemma 4.5, $v(s)=v\left(s^{\prime}\right)$.
$\Leftarrow$ Assume for all consistent valuations $v$ of nulls such that $v(D) \models \gamma(v(\bar{t})$, one also has $v(D) \models \gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$. By Proposition 4.7, if there is no such valuation, then $D \models \neg c o m p_{\gamma}(\bar{t})$. So assume now there is one such valuation. This entails that in particular, there exists a consistent $v^{*}$ satisfying $v^{*}(D) \models \gamma\left(v^{*}(\bar{t})\right)$ and $v^{*}(s)=v^{*}\left(s^{\prime}\right)$ iff $D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. By our assumption we have $v^{*}(D) \models \gamma^{\prime}\left(v^{*}\left(\overline{t^{\prime}}\right)\right)$. Hence, by Lemma $4.5, \forall s, s^{\prime}$ such that $D \models$ equiv $_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right)$ one has $v^{*}(s)=v^{*}\left(s^{\prime}\right)$. By the properties of $v^{*}$ mentioned above, this implies $D \models e q u i v_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. We have thus shown that $D \models i m p l y_{\gamma, \gamma^{\prime}}\left(\bar{t}, \bar{t}^{\prime}\right)$.

So far, we have dealt with equality subqueries and we have characterized the emptiness and inclusion of their supports (cf. Propositions 4.7 and 4.8 , respectively). We can now use this machinery to characterize the support of a BCCQ. We start by expressing membership in the support of an individual CQ:

## Lemma 4.9

Let $D$ be a database, $v$ a consistent valuation of $D$ and $Q(\bar{x})$ a conjunctive query in NRV normal form, with relational subquery $q(\bar{w})$ and equality subquery $\gamma(\bar{x}, \bar{w})$. Then, $v \in \operatorname{Supp}(Q, D, \bar{r})$ if and only there exists $\bar{s}$ such that $D \models q(\bar{s}) \wedge \operatorname{comp}_{\gamma}(\bar{r} \bar{s})$ and $v(D) \models$ $\gamma(v(\bar{r} \bar{s}))$.

## Proof

$\Rightarrow$ Assume $v \in \operatorname{Supp}(Q, D, \bar{r})$, that is $v$ is consistent and $v(\bar{r}) \in Q(v(D))$ and so there exists an assignment $\mu$ of $\bar{x} \bar{w}$ over $\operatorname{adom}(v(D)) \cup \operatorname{adom}(Q)$ such that $\mu(\bar{x})=v(\bar{r})$ and $v(D) \models q(\mu(\bar{w})) \wedge \gamma(\mu(\bar{x} \bar{w}))$.

Recall that $q(\bar{w})$ is a conjunction of relational atoms, with no constants and where each one of the free variables $\bar{w}$ has at most one occurrence in $q$. Thus, there exists a mapping $\nu$ of $\bar{w}$ over adom $(D)$ such that $D \models q(\nu(\bar{w}))$ and $v(\nu(\bar{w}))=\mu(\bar{w})$. We let $\bar{s}=\nu(\bar{w})$. Recall that $\bar{x}$ and $\bar{w}$ do not share variables, so we can extend the mapping $\nu$ by setting $\nu(\bar{x})=\bar{r}$. We thus have that $\nu(\bar{x} \bar{w})=\bar{r} \bar{s}$ and $v(\bar{r} \bar{s})=\mu(\bar{x} \bar{w})$. It follows that $v$ is a consistent valuation for which $v(D) \models \gamma(v(\bar{r} \bar{s}))$; then by Proposition 4.7, we also have $D \models \operatorname{comp}_{\gamma}(\bar{r} \bar{s})$.
$\Leftarrow$ Since $D \models q(\bar{s})$, we have $v(D) \models q(v(\bar{s}))$. Moreover by the hypothesis $v(D) \models$ $\gamma(v(\bar{r} \bar{s}))$, thus $v(D)=Q(v(\bar{r}))$.

In the remainder, we consider BCCQs $Q(\bar{x}):=Q_{1}(\bar{x}) \vee \ldots \vee Q_{n}(\bar{x})$ in NRV disjunctive normal form (DNF) where for all $1 \leq i \leq n$ :

$$
Q_{i}:=Q_{i_{0}}(\bar{x}) \wedge \neg Q_{i_{1}}(\bar{x}) \wedge \ldots \wedge \neg Q_{i_{m}}(\bar{x})
$$

and for all $1 \leq j \leq m$ :

$$
Q_{i_{j}}:=\exists \bar{w}_{i_{j}} q_{i_{j}}\left(\bar{w}_{i_{j}}\right) \wedge \gamma_{i_{j}} \text { with } \gamma_{i_{j}}:=e_{i_{j}}\left(\bar{x} \bar{w}_{i_{j}}\right)
$$

For convenience, we assume w.l.o.g every conjunction of literals to be of the same length $m$. We can also assume without loss of generality that for each $i$ we have $\operatorname{adom}\left(\gamma_{i_{j}}\right)=$ $\operatorname{adom}\left(\gamma_{i_{0}}\right)$ for all $j$. In fact, we can always pad any $\gamma_{i_{j}}$ with dummy equalities $c=c$ to extend its active domain.

Given a disjunct $Q_{i}$ in a BCCQ in DNF, we now define poss $Q_{i}$, encoding the set of possible answers to $Q_{i}$, and cons $_{Q_{i}}$, checking the compatibility of an answer with the negative literals in $Q_{i}$.

$$
\begin{gathered}
\operatorname{poss}_{Q_{i}}(\bar{x} \bar{w}):=q_{i_{0}}(\bar{w}) \wedge \operatorname{comp}_{\gamma_{i_{0}}}(\bar{x} \bar{w}) \wedge \operatorname{cons}_{Q_{i}}(\bar{x} \bar{w}) \\
\operatorname{cons}_{Q_{i}}(\bar{x} \bar{w}):=\bigwedge_{1 \leq j \leq m} \forall \bar{w}^{\prime}\left(\left(q_{i_{j}}\left(\bar{w}^{\prime}\right) \wedge \operatorname{comp}_{\gamma_{i_{j}}}\left(\bar{x} \bar{w}^{\prime}\right)\right) \rightarrow \neg \operatorname{imply}_{\gamma_{i_{0}}, \gamma_{i_{j}}}\left(\bar{x} \bar{w}, \bar{x} \bar{w}^{\prime}\right)\right)
\end{gathered}
$$

Using these new formulae, we show that the non-emptiness of $\operatorname{Supp}(Q(\bar{x}), D, \bar{r})$ can be expressed as the existence of a possible answer.

## Proposition 4.10

Let $D$ be a database and $Q(\bar{x})$ a DNF BCCQ in NRV normal form, then $\operatorname{Supp}(Q(\bar{x}), D, \bar{r}) \neq \emptyset$ if and only if $D \models \bigvee_{1 \leq i \leq n} \exists \bar{w} \operatorname{poss}_{Q_{i}}(\bar{r} \bar{w})$.

## Proof

$\Leftarrow$ Let $D \models \bigvee_{1 \leq i \leq n} \exists \bar{w} \operatorname{poss}_{Q_{i}}(\bar{r} \bar{w})$, then there exists $1 \leq i \leq n$ and an assignment $\nu$ with $\nu(\bar{w})=\bar{s}, D \models q_{i_{0}}(\bar{s}) \wedge \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s})$ and for all $1 \leq j \leq m$, for all $\bar{s}^{\prime}$ such that $D \models q_{i_{j}}\left(\overline{s^{\prime}}\right) \wedge \operatorname{comp}_{\gamma_{i_{j}}}\left(\bar{r} \overline{s^{\prime}}\right)$, one has $D \models \neg i m p l y_{\gamma_{i_{0}}, \gamma_{i_{j}}}\left(\bar{r} \bar{s}, \bar{r} \bar{s}^{\prime}\right)$. Since $D \models \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s})$, by Proposition 4.7 there exists a consistent valuation $v^{*}$, such that $v^{*}(D) \models \gamma_{i_{0}}\left(v^{*}(\bar{r} \bar{s})\right)$ and for all $s, s^{\prime} \in \operatorname{adom}(D) \cup \operatorname{adom}\left(\gamma_{i_{0}}\right)$, we have $v^{*}(s)=v^{*}\left(s^{\prime}\right)$ iff $D \models e^{\prime}$ quiv $_{\gamma_{i_{0}}}\left(\bar{r} \bar{s}, s, s^{\prime}\right)$.

Moreover, we can prove the following claim:

## Claim 4.11

For all conjunction of equalities $\gamma^{\prime}(\bar{y})$ with $\operatorname{adom}\left(\gamma^{\prime}\right)=\operatorname{adom}\left(\gamma_{i_{0}}\right)$ and all $\bar{t}$ over $\operatorname{adom}(D) \cup \operatorname{adom}\left(\gamma_{i_{0}}\right)$, one has $v^{*}(D) \models \gamma^{\prime}\left(v^{*}(\bar{t})\right)$ iff for all consistent valuations $v$, $v(D) \models \gamma_{i_{0}}(v(\bar{r} \bar{s}))$ implies $v(D) \models \gamma^{\prime}(v(\bar{t}))$.

## Proof of the Claim

$\Rightarrow$ Assume $v^{*}(D) \models \gamma^{\prime}\left(v^{*}(\bar{t})\right)$ and let $v$ be a consistent valuation such that $v(D) \models$ $\gamma_{i_{0}}(v(\bar{r} \bar{s}))$. We want to show $v(D) \models \gamma^{\prime}(v(\bar{t}))$. By Lemma 4.5, it is enough to show that $\forall s, s^{\prime} \in \operatorname{adom}(D) \cup \operatorname{adom}\left(\gamma_{i_{0}}\right), D \models$ equiv $_{\gamma^{\prime}}\left(\bar{t}, s, s^{\prime}\right)$ implies $v(s)=v\left(s^{\prime}\right)$. So let $s, s^{\prime}$ be such that $D \models e q u i v_{\gamma^{\prime}}\left(\bar{t}, s, s^{\prime}\right)$. As $v^{*}(D) \models \gamma^{\prime}\left(v^{*}(\bar{t})\right.$, by Lemma 4.5, $v^{*}(s)=v^{*}\left(s^{\prime}\right)$. By the properties of $v^{*}$, so $D \models e \operatorname{equiv}_{\gamma_{i_{0}}}\left(\bar{r} \bar{s}, s, s^{\prime}\right)$. As $v(D) \models \gamma_{i_{0}}(v(\bar{r} \bar{s}))$, by Lemma 4.5 it follows that $v(s)=v\left(s^{\prime}\right)$.
$\Leftarrow$ Assume for all consistent valuations $v, v(D) \models \gamma_{i_{0}}(v(\bar{r} \bar{s}))$ implies $v(D) \models \gamma^{\prime}(v(\bar{t}))$. By Lemma 4.5, $v^{*}$ being consistent $v^{*}(D) \models \gamma^{\prime}\left(v^{*}(\bar{t})\right)$.

Now fix some arbitrary $j \geq 1$ and $\overline{s^{\prime}}$ with $D \models q_{i_{j}}\left(\overline{s^{\prime}}\right) \wedge \operatorname{comp}_{\gamma_{i_{j}}}\left(\bar{r} \overline{s^{\prime}}\right)$. By Proposition 4.8, it follows from $D \models \neg i m p l y_{\gamma_{i_{0}}, \gamma_{i_{j}}}\left(\bar{r} \bar{s}, \bar{r} \bar{s}^{\prime}\right) \wedge \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s})$ that there exists a consistent valuation $v^{\prime}$ with $v^{\prime}(D) \models \gamma_{i_{0}}\left(v^{\prime}(\bar{r} \bar{s})\right)$ but $v^{\prime}(D) \not \vDash \gamma_{i_{j}}\left(v^{\prime}\left(\bar{r} \bar{s}^{\prime}\right)\right)$. By the above claim $v^{*}(D) \not \vDash \gamma_{i_{j}}\left(v^{*}\left(\bar{r} \bar{s}^{\prime}\right)\right)$. In summary, we have:
(i) $D \models q_{i_{0}}(\bar{s}) \wedge \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s})$ and $v^{*}(D) \models \gamma_{i_{0}}\left(v^{*}(\bar{r} \bar{s})\right)$ and so by Lemma 4.9, we have $v^{*} \in \operatorname{Supp}\left(Q_{i_{0}}(\bar{x}), D, \bar{r}\right)$, that is, $v^{*}(D) \models Q_{i_{0}}\left(v^{*}(\bar{r})\right)$.
(ii) For all $1 \leq j \leq m$ and assignment $\nu^{\prime}$ with $\nu^{\prime}(\bar{w})=\overline{s^{\prime}}$, if $D \models q_{i_{j}}\left(\overline{s^{\prime}}\right) \wedge \operatorname{comp}_{\gamma_{i_{j}}}\left(\bar{r} \overline{s^{\prime}}\right)$ then $v^{*}(D) \not \models \gamma_{i_{j}}\left(v^{*}\left(\bar{r} \bar{s}^{\prime}\right)\right)$ and so by Lemma 4.9, we have $v^{*} \notin \operatorname{Supp}\left(Q_{i_{j}}(\bar{x}), D, \bar{r}\right)$, that is, for all $1 \leq j \leq m, v^{*}(D) \models \neg Q_{j}\left(v^{*}(\bar{r})\right)$.

This means we have $v^{*} \in \operatorname{Supp}\left(Q_{i_{0}}(\bar{x}) \wedge \neg Q_{i_{1}}(\bar{x}) \wedge \ldots \wedge \neg Q_{i_{m}}(\bar{x}), D, \bar{r}\right)$ for all $1 \leq i \leq n$ and so $v^{*} \in \operatorname{Supp}(Q(\bar{x}), D, \bar{r})$.
$\Rightarrow$ Let $v \in \operatorname{Supp}(Q(\bar{x}), D, \bar{r})$, so $v$ is consistent and there is some $1 \leq i \leq n$ with: (i) $v \in \operatorname{Supp}\left(Q_{i_{0}}, D, \bar{r}\right)$, (ii) for all $1 \leq j \leq m, v \notin \operatorname{Supp}\left(Q_{i_{j}}, D, \bar{r}\right)$. Using Lemma 4.9 (i) implies that there exists $\bar{s}$ such that $D \models q_{i_{0}}(\bar{s}) \wedge \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s})$ and $v(D) \models \gamma_{i_{0}}(v(\bar{r} \bar{s}))$. Again by Lemma 4.9, (ii) implies that for all $1 \leq j \leq m$ and $\overline{s^{\prime}}$, if $D \models q_{i_{j}}\left(\overline{s^{\prime}}\right) \wedge \operatorname{comp}_{\gamma_{i_{j}}}\left(\bar{r} \bar{s}^{\prime}\right)$ then $v(D) \not \vDash \gamma_{i_{j}}\left(v\left(\bar{r} \overline{s^{\prime}}\right)\right)$. This entails by Proposition 4.8 that $D \models \operatorname{comp}_{\gamma_{i_{0}}}(\bar{r} \bar{s}) \wedge \neg i m p l y_{\gamma_{i_{0}}, \gamma_{i_{j}}}\left(\bar{r} \bar{s}, \bar{r} \overline{s^{\prime}}\right)$. This shows $D \models \bigvee_{1 \leq i \leq n} \exists \bar{w}$ $\operatorname{poss}_{Q_{i}}(\bar{r} \bar{w})$.

Now that we have defined the formula expressing for a BCCQ $Q$ non-emptiness of $\operatorname{Supp}(Q(\bar{x}), D, \bar{r})$ (Proposition 4.10), we can easily define a rewriting for the problem $\operatorname{CertainAnswer}_{\Sigma}(Q)$. To do so, we rely on the fact that $\bar{r} \in \operatorname{cert}_{\Sigma}(Q, D)$ iff $\operatorname{Supp}(\neg Q, D, \bar{r})=\emptyset$.

## Theorem 4.12 (Datalog rewriting)

Let D be a database whose schema contains a set of equality-generating dependencies $\Sigma$, and let $Q(\bar{x})$ be a BCCQ in NRV normal form. Let $Q^{\prime}=Q_{1}^{\prime}(\bar{x}) \vee \ldots \vee Q_{n}^{\prime}(\bar{x})$ be $\neg Q$ in DNF normal form. Then, $\bar{r} \in \operatorname{cert}_{\Sigma}(Q, D)$ if and only if $D \models \rho(\bar{r})$ where $\rho(\bar{x})=\bigwedge_{1 \leq i \leq n} \forall \bar{w}$ $\neg$ poss $_{Q_{i}^{\prime}}(\bar{x} \bar{w})$.

## Proof

One has that $\bar{r} \in \operatorname{cert}_{\Sigma}(Q, D)$ iff $\operatorname{Supp}\left(Q^{\prime}, D, \bar{r}\right)=\emptyset . Q^{\prime}$ being still a BCCQ, Proposition 4.10 tells us that $\operatorname{Supp}\left(Q^{\prime}, D, \bar{r}\right)=\emptyset$ iff $D \models \bigwedge_{1 \leq i \leq n} \forall \bar{w} \neg \operatorname{poss}_{Q_{i}^{\prime}}(\bar{r} \bar{w})$.

Corollary 4.13
For each fixed BCCQ query $Q$ and a set of EGDs $\Sigma$, the complexity of $\operatorname{Certain}^{\prime} \operatorname{Answer}_{\Sigma}(Q)$ is in PTIME.

### 4.3 Non-rewritability in FO

The basic starting points for our investigation was the fact that $\operatorname{cert}_{\Sigma}(Q, D)=$ $Q\left(\operatorname{chase}_{\Sigma}(D)\right)$ for a CQ $Q$ and a set $\Sigma$ of FDs, for every database $D$. This remained true for unions of CQs, but failed for BCCQs, forcing us to produce a Datalog rewriting to obtain certain answers. But can a first-order rewriting be obtained instead? This would make it possible to produce certain answers using the core of SQL as opposed to its recursive features which do not always perform as well in practice.

In this section, we show that the answer, in general, is negative even for CQs (and thus for BCCQs). In the next section, however we show that such rewritings can be obtained in FO for BCCQs whenever $\Sigma$ is empty.

The main result of this section is the following.

## Theorem 4.14

There exists a Boolean CQ $Q$ and single FD $\Sigma$ over a relational schema of binary and unary relations such that $\operatorname{cert}_{\Sigma}(Q, D)$ is not expressible as an FO query.

## Proof

Consider a schema with one binary relation $E$ and two unary relations $A$ and $B$. The only FD in $\Sigma$ is $\forall x \forall y \forall z(E(x, y) \wedge E(x, z) \rightarrow y=z)$; in other words, the first attribute of $E$ is a key. The query $Q$ is a Boolean CQ $\exists x(A(x) \wedge B(x))$.

To prove inexpressibility of $\operatorname{cert}_{\Sigma}(Q, \cdot)$ in FO, for each $n>0$ we create two databases $D_{n}$ and $D_{n}^{\prime}$. In both of them, $E$ is interpreted as a disjoint union $T_{1} \cup T_{2}$ where $T_{1}$ and $T_{2}$ are balanced binary trees of depth $n$ in which all nodes are distinct nulls. In both $A$ and $B$ are singleton sets. In $D_{n}$, the set $A$ contains a leaf of $T_{1}$ and $B$ contains a leaf of $T_{2}$. In $D_{n}^{\prime}$, both $A$ and $B$ contain leaves of $T_{1}$ such that their only common ancestor in the tree is the root (in other words, they are leaves of subtrees rooted at different children of the root of $T_{1}$ ).

Because of the constraint $\Sigma$, for every valuation $v$ such that the resulting database satisfies it we have that both $v\left(T_{1}\right)$ and $v\left(T_{2}\right)$ are chains. Indeed, consider any node $\perp$ with children $\perp_{1}, \perp_{2}$ in $T_{i}$. If $v\left(\perp_{1}\right) \neq v\left(\perp_{2}\right)$ then the resulting tuples $\left(v(\perp), v\left(\perp_{1}\right)\right)$ and
$\left(v(\perp), v\left(\perp_{2}\right)\right)$ violate the constraint. Thus, $v\left(\perp_{1}\right)=v\left(\perp_{2}\right)$ and applying this construction inductively we see that $v\left(T_{i}\right)$ is a chain. Hence, it has a single leaf, and thus cert ${ }_{\Sigma}\left(Q, D_{n}^{\prime}\right)$ is true, since $A$ and $B$ must be interpreted as that leaf. On the other hand, $\operatorname{cert}_{\Sigma}\left(Q, D_{n}\right)$ is false, since there is a valuation $v$ that sends $T_{1}$ and $T_{2}$ into two disjoint chains, and thus, $A$ and $B$ are interpreted as two distinct elements.

Assume now that $\operatorname{cert}_{\Sigma}(Q, \cdot)$ is rewritable as an FO sentence $\phi$. Then, for every $n>0$, we have $D_{n}^{\prime} \models \phi$ and $D_{n} \models \neg \phi$. We next show that such a sentence cannot exist, thereby proving non-FO-rewritability.

Recall that in a database (with one binary relation, like considered here), a radius $r$ neighborhood of an element $a$ is its restriction to the set of all elements reachable from $a$ by a path of length at most $r$, where the path does not take into account the orientation of edges of $E$ (e.g. if we have $E(a, b)$ and $E(c, a)$ then both $b$ and $c$ are in the radius 1 neighborhood of $a$ ). When two neighborhoods, of elements $a$ and $b$, are isomorphic, it means that there is an isomorphism between them that sends $a$ to $b$. In other words, centers of neighborhoods are viewed as distinguished elements when it comes to defining neighborhoods. It is known that each first-order sentence $\psi$ is Hanf-local (Fagin et al. 1995): that is, there exists a number $r>0$ such that for any two databases $D_{1}$ and $D_{2}$, if there is a bijection $f$ between $D_{1}$ and $D_{2}$ such that the radius $r$ neighborhoods of $a$ in $D_{1}$ and $f(a)$ in $D_{2}$ are isomorphic then $D_{1}$ and $D_{2}$ agree on $\psi$, that is either both satisfy it or both do not.

Now let $r$ be such a number for the sentence $\phi$ we assumed exists. Consider $D_{n}$ and $D_{n}^{\prime}$ and let $T_{1 a}, T_{1 *}$ be the subtrees of the root of $T_{1}$ in $D_{n}$ such that the first contains $A$ while the second contains neither $A$ not $B$, and let $T_{2 b}, T_{2 *}$ be defined similarly for subtrees of the root of $T_{2}$ with respect to $B$. In $D_{n}^{\prime}$ we define $T_{1 a}^{\prime}, T_{1 b}^{\prime}$ as subtrees of the root of the tree containing $A, B$ such that the first contains the $A$ leaf and the second contains the $B$ leaf, while $T_{2 *}^{\prime}, T_{2 * *}^{\prime}$ be the subtrees of the root of the tree having neither $A$ nor $B$ elements. Then, it is easy to see that the following pairs of trees are isomorphic: $T_{1 a}$ and $T_{1 a}^{\prime}, T_{2 b}$ and $T_{1 b}^{\prime}, T_{1 *}$ and $T_{2 *}^{\prime}, T_{2 *}$ and $T_{2 * *}^{\prime}$.

We now define the bijection $f$ as the union of those isomorphisms plus mapping roots of trees $T_{i}$ in $D$ into roots of $T_{i}$ in $D^{\prime}$. It is an immediate observation that if $n>r+1$ (i.e. leaves are not in the radius $r$ neighborhood of children of roots) then $f$ satisfies the condition that neighborhoods of $a$ and $f(a)$ of radius $r$ are isomorphic. This would tell us that $D_{n}$ and $D_{n}^{\prime}$ agree on $\phi$ but we know they do not. This contradiction completes the proof.

As a corollary to the proof, we obtain the following result showing that non-recursive SQL is incapable of computing $\operatorname{cert}_{\Sigma}(Q, D)$ in the setting of Theorem 4.14.

## Corollary 4.15

There exists a Boolean CQ $Q$ and single FD $\Sigma$ over a relational schema of binary and unary relations such that $\operatorname{cert}_{\Sigma}(Q, D)$ is not expressible in the basic SELECT-FROM-WHERE-GROUP BY-HAVING fragment of SQL with arbitrary aggregate functions.

This is due to the fact that queries in this fragment of SQL with grouping and aggregation can be translated into a logic with aggregate functions (Libkin 2003) which itself is known to be Hanf-local (Hella et al. 2001).

### 4.4 FO rewriting for certain answers

We now focus on the special case where $\Sigma$ is empty. First notice that the only Data$\log$ component in our rewriting was the equiv ${ }_{\gamma}$ formula; moreover, notice that without constraints, equiv $v_{\gamma}$ simply computes a reflexive symmetric transitive closure. More precisely, for a given $\bar{t}=\nu(\bar{y})$, one has that equiv $_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ holds in $D$ iff $\left(s, s^{\prime}\right)$ belongs to the reflexive symmetric transitive closure of $\{(\nu(x), \nu(w)) \mid x=w \in \gamma(\bar{y})\}$ over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$.

## Definition 4.16

Given a database $D$, a conjunction of equality atoms $\gamma(\bar{y})$ and an assignment $\nu: \bar{y} \cup$ $\operatorname{adom}(\gamma) \rightarrow \operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ preserving constants, we say that $u, u^{\prime} \in \operatorname{adom}(D) \cup$ $\operatorname{adom}(\gamma)$ are equivalent w.r.t. $\gamma$ and $\nu$ and write $u \equiv_{\gamma}^{\nu} u^{\prime}$, if either $u=u^{\prime}$ or $\left(u, u^{\prime}\right)$ belongs to the reflexive symmetric transitive closure of $\{(\nu(x), \nu(w)) \mid x=w \in \gamma\}$.

As $\Sigma$ is empty, we can rewrite as follows the equiv formula in FO, where $m$ is the number of equivalence classes of $\sim_{\gamma}$ :

$$
\begin{aligned}
& \operatorname{equivF} O_{\gamma}\left(\bar{y}, z, z^{\prime}\right):=z=z^{\prime} \vee \\
& \bigvee_{\substack{u_{1}, v_{1} \ldots u_{m}, v_{m} \in \bar{y} \cup \operatorname{adom}(\gamma) \\
u_{i} \sim_{\gamma} v_{i} \text { for all } 1 \leq i \leq m}}\left(z=u_{1} \wedge z^{\prime}=v_{m} \wedge \bigwedge_{1 \leq i<m} v_{i}=u_{i+1}\right)
\end{aligned}
$$

## Proposition 4.17

Given an incomplete database $D$, a conjunction of equality atoms $\gamma(\bar{y})$ and an assignment $\nu(\bar{y})=\bar{t}$ over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$, given $s, s^{\prime}$ in $\bar{t} \cup \operatorname{adom}(\gamma)$, we have that $D \models e q u i v F O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ if and only if $s \equiv_{\gamma}^{\nu} s^{\prime}$.

Intuitively, this holds because each disjunct of $\operatorname{equivFO}\left(\bar{t}, s, s^{\prime}\right)$ corresponds to a possible derivation of ( $s, s^{\prime}$ ) in the reflexive symmetric transitive closure of $\{(\nu(x), \nu(w)) \mid x=$ $w \in \gamma\}$, and one can prove that there is a bound only depending on $\gamma$ on the number of steps of this derivation.

## Proof

To start with we naturally extend $\nu$ to be the identity on adom $(\gamma)$. Assume first $D \models$ equiv $F O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. If $s=s^{\prime}$ then $s \equiv_{\gamma}^{\nu} s^{\prime}$. Now assume $s \neq s^{\prime}$. Then, there exist variables and/or constants $u_{1}, v_{1} \ldots u_{m}, v_{m} \in \bar{y} \cup$ adom $(\gamma)$ with $u_{i} \sim_{\gamma} v_{i}$ for all $i$, such that $s=\nu\left(u_{1}\right), s^{\prime}=\nu\left(v_{m}\right)$ and $\nu\left(v_{i}\right)=\nu\left(u_{i+1}\right)$ for all $i<m$. Clearly $u_{i} \sim_{\gamma} v_{i}$ implies $\nu\left(u_{i}\right) \equiv_{\gamma}^{\nu} \nu\left(v_{i}\right)$. Then, $\nu\left(u_{i}\right) \equiv_{\gamma}^{\nu} \nu\left(u_{i+1}\right)$ for all $i<m$. We conclude by transitivity that $s=\nu\left(u_{i}\right) \equiv_{\gamma}^{\nu} \nu\left(u_{m}\right) \equiv_{\gamma}^{\nu} \nu\left(v_{m}\right)=s^{\prime}$, and therefore, $s \equiv_{\gamma}^{\nu} s^{\prime}$.

Assume now that $s \equiv_{\gamma}^{\nu} s^{\prime}$. If $s=s^{\prime}$ then clearly $D \models \operatorname{equivFO} O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. Thus, assume $s \neq s^{\prime}$. We proceed by induction on the number of transitive closure steps needed to derive $\left(s, s^{\prime}\right)$ starting for the base relation $\{(\nu(x), \nu(w)) \mid x=w \in \gamma\}$. In the base case, $\left(s, s^{\prime}\right)=(\nu(x), \nu(w))$ for some equality $x=w \in \gamma$. Then, $D$ satisfies the following disjunct of equivFO $\left(\bar{t}, s, s^{\prime}\right)$ : take $u_{1}=x, v_{1}=w, u_{i}=v_{i}=w$ for all $i=2 . . m$ (this is a disjunct since $u_{i} \sim_{\gamma} v_{i}$ for all $\left.i=1 . . m\right)$. The disjunct is satisfied since $s=\nu(x)=\nu\left(u_{1}\right)$, $s^{\prime}=\nu(w)=\nu\left(v_{m}\right)$, and for all $i=1 . . m-1, \nu\left(v_{i}\right)=\nu\left(u_{i+1}\right)=\nu(w)$.

In the general inductive case, there exists $r$ such that $\left(r, s^{\prime}\right)=(\nu(x), \nu(w))$ for some equality $x=w$ (or $w=x) \in \gamma$, with $s \equiv_{\gamma}^{\nu} r$ derived at the previous step. By the induction
hypothesis $D \vDash \operatorname{equiv} F O_{\gamma}(\bar{t}, s, r)$. We can assume $s \neq r$ since otherwise ( $s, s^{\prime}$ ) would be in the base relation. Therefore, $D$ satisfies one of the disjuncts of equivFO $(\bar{t}, s, r)$. Then, there exists a sequence of $m+1$ pairs in $\bar{y} \cup \operatorname{adom}(\gamma)$

$$
\left(u_{1}, v_{1}\right)\left(u_{2}, v_{2}\right) \ldots\left(u_{m}, v_{m}\right)\left(u_{m+1}, v_{m+1}\right)
$$

such that

- $u_{m+1}=x$ and $v_{m+1}=w$,
- $u_{i} \sim_{\gamma} v_{i}, i=1 . . m+1$,
- $s=\nu\left(u_{1}\right), r=\nu\left(v_{m}\right), s^{\prime}=\nu\left(v_{m+1}\right)$,
- $\nu\left(v_{i}\right)=\nu\left(u_{i+1}\right)$, for all $i \leq m$,

We now show that from this sequence of pairs, one can construct another one of exactly $m$ pairs, $\left(u_{i}^{\prime}, v_{i}^{\prime}\right), i=1$.. $m$ still connecting $s$ ans $s^{\prime}$, that is such that:
(a) $u_{i}^{\prime} \sim_{\gamma} v_{i}^{\prime}, i=1 . . m$
(b) $s=\nu\left(u_{1}^{\prime}\right), s^{\prime}=\nu\left(v_{m}^{\prime}\right)$
(c) $\nu\left(v_{i}^{\prime}\right)=\nu\left(u_{i+1}^{\prime}\right)$, for all $i<m$.

The idea is to first cut the sequence $\left(u_{i}, v_{i}\right), i=1 . . m+1$, removing at least one pair, then pad it to size $m$ if necessary.

In order to cut the original sequence, remark that it contains $m+1$ pairs where $m$ is the number of $\sim_{\gamma}$ equivalence classes. Thus, there exist $i<j$ such that $u_{i} \sim_{\gamma} u_{j}$. We remove from the sequence all elements between $u_{i}$ and $v_{j}$ (excluded), the new sequence is

$$
\left(u_{1}, v_{1}\right) \ldots\left(u_{i-1}, v_{i-1}\right)\left(u_{i}, v_{j}\right)\left(u_{j+1}, v_{j+1}\right) \ldots\left(u_{m+1}, v_{m+1}\right)
$$

Note that this sequence satisfies (a) (b) and (c) above since $u_{i} \sim_{\gamma} u_{j} \sim_{\gamma} v_{j}$. Let the new sequence contain $k$ pairs. We know $k \leq m$ because we have removed at least one pair from the original sequence (recall $i<j$ ). If $k<m$, we pad the sequence on the right with $m-k$ pairs $\left(v_{m+1}, v_{m+1}\right)$. The new sequence still satisfies (a), (b), and (c); therefore, the corresponding disjunct of equivFO $O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ is satisfied by $D$.

## Example 4.18

Let $\gamma:=y_{1}=y_{2} \wedge z=x$ be the equality subquery of the query $Q(x)$ in Example 4.2. Up to logical equivalence, equivFO $O_{\gamma}\left(y_{1}, y_{2}, z, x, w, w^{\prime}\right)$ contains precisely the disjuncts $w=w^{\prime}$, $w=y_{1} \wedge w^{\prime}=y_{2}, w=z \wedge w^{\prime}=x, w=y_{1} \wedge w^{\prime}=x \wedge y_{2}=z$, plus all disjuncts obtained from them by applying one or more of the following transformations: switch $w$ and $w^{\prime}$, switch $y_{1}$ and $y_{2}$, switch $x$ and $z$. Let $D$ be the database from Example 2.1, then we have for instance $D \models \operatorname{equivFO}\left(1, \perp_{2}, \perp_{2}, 1, a, a^{\prime}\right)$ and $D \models \operatorname{equivFO}\left(1, \perp_{2}, \perp_{2}, \perp_{2}, a, a^{\prime}\right)$ for all $a, a^{\prime} \in\left\{1, \perp_{2}\right\}$. Similarly $D \models$ equivFO $O_{\gamma}\left(\perp_{1}, \perp_{2}, \perp_{2}, 1, a, a^{\prime}\right)$ for all $a, a^{\prime} \in$ $\left\{1, \perp_{1}, \perp_{2}\right\}$.

As a consequence of Proposition 4.17, for fixed $\gamma$ and $\bar{t}$, the relation $\left\{\left(s, s^{\prime}\right) \mid D \models\right.$ equivFO $\left.O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)\right\}$ is an equivalence relation over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ where each element of $\operatorname{adom}(D)$ neither in $\bar{t}$ nor in $\operatorname{adom}(\gamma)$ forms a singleton equivalence class.

As in Section 4.2, equivFO ${ }_{\gamma}$ selects precisely the pairs of elements of a tuple that a valuation needs to collapse to satisfy a set of equalities.

As a consequence, we can rewrite in FO the formula $\operatorname{poss}_{Q_{i}}$ of Subsection 4.2 encoding the set of possible answers to $Q_{i}$. It is enough to replace each occurrence of the Datalog equiv $\left(\bar{y}, z, z^{\prime}\right)$ program in it by $\operatorname{equiv}_{\gamma}^{F O}\left(\bar{y}, z, z^{\prime}\right)$. We denote by $\operatorname{poss}_{Q_{i}}^{F O}$ the rewriting so obtained.

Theorem 4.19 (FO rewriting)
Let D be a database, $\Sigma=\emptyset$ and let $Q(\bar{x})$ be a BCCQ in NRV normal form. Let $Q^{\prime}=$ $Q_{1}^{\prime}(\bar{x}) \vee \ldots \vee Q_{n}^{\prime}(\bar{x})$ be $\neg Q$ in DNF normal form. Then, $\bar{r} \in \operatorname{cert}_{\Sigma}(Q, D)$ if and only if $D \models \rho(\bar{r})$ where $\rho(\bar{x})=\bigwedge_{1 \leq i \leq n} \forall \bar{w} \neg \operatorname{poss}_{Q_{i}^{\prime}}^{F O}(\bar{x} \bar{w})$.
Note that tractability of BCCQ was already proved in Gheerbrant and Libkin (2015) using tableau-based methods. We now refine complexity as follows.

## Corollary 4.20

For each fixed $\operatorname{BCCQ}$ query $Q$, the complexity of $\operatorname{Certain}^{\operatorname{Answer}}{ }_{\Sigma}(Q)$ is in DLOGSPACE whenever $\Sigma=\emptyset$.

### 4.5 FO rewriting for best answers

Considering arbitrary FO queries brought us an intrinsic intractability result for all variants of best answers. This motivates restricting to unions of conjunctive queries, for which a polynomial-time evaluation algorithm (in data complexity) already exists(Libkin 2018). The resolution-based procedure is however in sharp contrast with naïve evaluation, which allows to compute certain answers to unions of conjunctive queries via usual model checking. We thus initiate a descriptive complexity analysis of the best answers problem, showing that for unions of conjunctive queries, it can essentially be reduced modulo a preprocessing of the query - to (naïve) evaluation of an FO-formula.

Given a union of conjunctive queries $Q$, our starting point towards an FO rewriting for best answers is finding an FO-formula $Q_{\subseteq}(\bar{x}, \bar{y})$ encoding the inclusion of supports, that is selecting tuples $\bar{s}, \bar{t}$ over $\operatorname{adom}(D) \cup \operatorname{adom}(Q)$ iff $\operatorname{Supp}(Q, D, \bar{s}) \subseteq \operatorname{Supp}(Q, D, \bar{t})$. From $Q_{\subseteq}$, one can easily define an FO-formula selecting precisely all best answers to $Q$ on $D$ :

$$
\begin{equation*}
\operatorname{best}_{Q}(\bar{x}):=\forall \bar{y}\left(Q_{\subseteq}(\bar{x}, \bar{y}) \rightarrow Q_{\subseteq}(\bar{y}, \bar{x})\right) \tag{1}
\end{equation*}
$$

As in Section 4.2, we assume all CQs to be in NRV normal form. We can thus first concentrate on the support of equality subqueries. This will be encoded in FO and then integrated in the rewriting of the whole conjunctive query.

We now go back to an arbitrary union of conjunctive queries of vocabulary $\sigma$ in NRVnormal form:

$$
Q(\bar{x}):=\bigvee_{1 \leq i \leq n} Q_{i}(\bar{x})
$$

where each $Q_{i}$ is in NRV normal form with relational subquery $q_{i}\left(\bar{y}_{i}, \bar{z}_{i}\right)$ and equality subquery $e q_{i}\left(\bar{x}, \bar{y}_{i}, \bar{z}_{i}\right)$.

Recall the formula $\operatorname{comp}_{\gamma}$, defined in Section 4.2, stating the existence of a valuation that collapses all equivalent elements of a tuple:

$$
\operatorname{comp}_{\gamma}(\bar{y}):=\forall z z^{\prime}\left(\operatorname{equivFO} O_{\gamma}\left(\bar{y}, z, z^{\prime}\right) \wedge \neg N u l l(z) \wedge \neg N u l l\left(z^{\prime}\right) \rightarrow z=z^{\prime}\right)
$$

Notice that if $D \models \operatorname{comp}_{\gamma}(\bar{t})$ then for each $s \in \operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ there exists at most one constant $c$ such that $D \models \operatorname{equivFO}(\bar{t}, s, c)$. In fact if for constants $c_{1}$ and $c_{2}, D \models$
equivFO $\gamma_{\gamma}\left(\bar{t}, s, c_{1}\right)$ and $D \models \operatorname{equivFO}\left(\bar{t}, s, c_{2}\right)$, by transitivity $D \models \operatorname{equivFO} O_{\gamma}\left(\bar{t}, c_{1}, c_{2}\right)$, implying $c_{1}=c_{2}$.

## Example 4.21

Let $D$ and $\gamma$ be as in Example 4.18. Consider $\operatorname{comp}_{\gamma}\left(y_{1}, y_{2}, z, x\right)$. Given the tuples selected by equivFO $O_{\gamma}$ in Example 4.18, we can conclude that $D \models \operatorname{comp}_{\gamma}\left(1, \perp_{2}, \perp_{2}, 1\right)$.

We now define a formula capturing the inclusion of supports between two conjunctions of equality atoms, which will be a crucial ingredient in our rewriting.

Let $\gamma(\bar{x})$ and $\gamma^{\prime}(\bar{y})$ be conjunctions of equality atoms with $\operatorname{adom}(\gamma)=\operatorname{adom}\left(\gamma^{\prime}\right)$. We define:

$$
i m p l y_{\gamma, \gamma^{\prime}}(\bar{x}, \bar{y}):=\forall z z^{\prime}\left(\operatorname{equivFO}{\gamma^{\prime}}\left(\bar{y}, z, z^{\prime}\right) \rightarrow \operatorname{equivFO} O_{\gamma}\left(\bar{x}, z, z^{\prime}\right)\right)
$$

## Example 4.22

Let $\gamma$ and $D$ be as in Example 4.18. Let $\gamma^{\prime}:=y_{1}^{\prime}=y_{2}^{\prime} \wedge z^{\prime}=x^{\prime}$, then it follows from Example 4.18 that $D \vDash \operatorname{imply}_{\gamma \gamma^{\prime}}\left(\perp_{1} \perp_{2} \perp_{2} 1,1 \perp_{2} \perp_{2} \perp_{2}\right)$ and $D \models$ imply $y_{\gamma \gamma^{\prime}}\left(1 \perp_{2} \perp_{2} 1,1 \perp_{2} \perp_{2} \perp_{2}\right)$.

By combining Propositions 4.8 and 4.7, we get:
Corollary 4.23
Let $\gamma(\bar{y}), \gamma^{\prime}(\bar{y})$ be conjunctions of equality atoms with $\operatorname{adom}(\gamma)=\operatorname{adom}\left(\gamma^{\prime}\right), D$ a database and $\nu(\bar{y})=\bar{t}, \nu^{\prime}(\bar{y})=\bar{t}^{\prime}$ assignments over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$. If $D \models \operatorname{comp}_{\gamma}(\bar{t}) \wedge$ imply $y_{\gamma, \gamma^{\prime}}\left(\bar{t}, \overline{t^{\prime}}\right)$, then $D \models \operatorname{comp}_{\gamma^{\prime}}\left(\overline{t^{\prime}}\right)$.

## Proof

Assume $D \vDash \operatorname{comp}_{\gamma}(\bar{t}) \wedge \operatorname{imply}_{\gamma, \gamma^{\prime}}\left(\bar{t}, \overline{t^{\prime}}\right)$, that is, $D \vDash \forall z z^{\prime}\left(e^{2} u i v F O_{\gamma}\left(\bar{t}, z, z^{\prime}\right) \wedge\right.$ $\left.\neg \operatorname{Null}(z) \wedge \neg \operatorname{Null}\left(z^{\prime}\right) \rightarrow z=z^{\prime}\right)$ and $D \quad \vDash \forall z z^{\prime} \quad\left(e^{\prime} u i v F O_{\gamma^{\prime}}\left(\overline{t^{\prime}}, z, z^{\prime}\right) \rightarrow\right.$ equivFO $\left.\gamma_{\gamma}\left(\bar{t}, z, z^{\prime}\right)\right)$. Now let $s, s^{\prime} \in \operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ with $D \models$ equivFO $O_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right) \wedge$ $\neg \operatorname{Null}(s) \wedge \neg \operatorname{Null}\left(s^{\prime}\right)$. As $D \models \operatorname{imply}_{\gamma, \gamma^{\prime}}\left(\bar{t}, \overline{t^{\prime}}\right)$, it follows that $D \models$ equivFO $O_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$ and so $D \models \neg \operatorname{Null}(s) \wedge \neg \operatorname{Null}\left(s^{\prime}\right) \rightarrow s=s^{\prime}$ now follows from $D \models \operatorname{comp}_{\gamma}(\bar{t})$. Hence $D \models \operatorname{comp}_{\gamma^{\prime}}\left(\overline{t^{\prime}}\right)$.

We are now ready to define the FO-formula encoding the inclusion of supports.

$$
\begin{aligned}
Q_{\subseteq}\left(\bar{x}, \overline{x^{\prime}}\right):= & \bigwedge_{1 \leq i \leq n} \\
& \left(\forall \overline { y } \overline { z } \left(\left(q_{i}(\bar{y}, \bar{z}) \wedge \operatorname{comp}_{e q_{i}}(\bar{x}, \bar{y}, \bar{z})\right) \rightarrow\right.\right. \\
& \left.\left.\bigvee_{1 \leq j \leq n} \exists \bar{y}^{\prime} \bar{z}^{\prime}\left(q_{j}\left(\bar{y}^{\prime}, \bar{z}^{\prime}\right) \wedge i m p l y_{e q_{i}, e q_{j}}\left(\bar{x} \bar{y} \bar{z}, \bar{x}^{\prime} \bar{y}^{\prime} \bar{z}^{\prime}\right)\right)\right)\right)
\end{aligned}
$$

Combining Lemmas 4.5, 4.9, Propositions 4.7, 4.8, and Corollary 4.23, we get:
Proposition 4.24
$D \models Q_{\subseteq}(\bar{s}, \bar{t})$ iff $\operatorname{Supp}(Q, D, \bar{s}) \subseteq \operatorname{Supp}(Q, D, \bar{t})$.
Proof
$\Rightarrow$ Assume $D \models Q_{\subseteq}(\bar{s}, \bar{t})$ and let $v \in \operatorname{Supp}(Q, D, \bar{s})$ be a valuation of $D$. By Lemma 4.9, $\exists i \bar{a} \bar{b} D \models q_{i}(\bar{a} \bar{b}) \wedge \operatorname{comp}_{e q_{i}}(\bar{s} \bar{a} \bar{b})$ and $v(D) \models e q_{i}(v(\bar{s} \bar{a} \bar{b}))$. So by our assumption there exists $j, \bar{a}^{\prime} \bar{b}^{\prime}$ with $D \models q_{j}\left(\bar{a}^{\prime} \bar{b}^{\prime}\right) \wedge i m p l y_{e q_{i}, e q_{j}}\left(\bar{s} \bar{a} \bar{b}, \bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)$ and by Corollary $4.23 D \models \operatorname{comp}_{e q_{j}}\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)$. Now let $t_{1}, t_{2}$ such that $D \models e q u i v_{e q_{j}}\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}, t_{1}, t_{2}\right)$. By $D \vDash \operatorname{imply}_{e q_{i}, e q_{j}}\left(\bar{s} \bar{a} \bar{b}, \bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)$, we have $D \models \operatorname{equiv}_{e q_{i}}\left(\bar{s} \bar{a} \bar{b}, t_{1}, t_{2}\right)$, and by Lemma 4.5,
$v\left(t_{1}\right)=v\left(t_{2}\right)$. But then, again by Lemma 4.5, $v(D) \models e q_{i}\left(v\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)\right)$ and by Lemma 4.9 it follows that $v \in \operatorname{Supp}(Q, D, \bar{t})$.
$\Leftarrow$ Assume $\operatorname{Supp}(Q, D, \bar{s}) \subseteq \operatorname{Supp}(Q, D, \bar{t})$ and let $i, \bar{a}, \bar{b}$ with $D \vDash q_{i}(\bar{a}, \bar{b}) \wedge$ $\operatorname{comp}_{e q_{i}}(\bar{s}, \bar{a}, \bar{b})$. By Proposition 4.7, there exists a valuation $v$ (that we assume w.l.o.g. to be tight) such that $v(D) \models e q_{i}(v(\bar{s} \bar{a} \bar{b}))$ and so by Lemma $4.9 v \in \operatorname{Supp}(Q, D, \bar{s})$. Hence by our assumption, we also have $v \in \operatorname{Supp}(Q, D, \bar{t})$, and so by Lemma 4.9, there exists $j, \bar{a}^{\prime} \bar{b}^{\prime}$ with $D \models q_{j}\left(\bar{a}^{\prime} \bar{b}^{\prime}\right) \wedge \operatorname{comp}_{e q_{j}}\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)$ and $v(D) \models e q_{j}\left(v\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)\right)$.

## Claim 4.25

Let $D$ be a database, $\gamma(\bar{y}), \gamma^{\prime}(\bar{y})$ conjunctions of equality atoms with adom $(\gamma)=$ $\operatorname{adom}\left(\gamma^{\prime}\right), \nu(\bar{y})=\bar{t}, \nu^{\prime}(\bar{y})=\overline{t^{\prime}}$ assignments over $\operatorname{adom}(D) \cup \operatorname{adom}(\gamma)$ and $v^{*}$ a valuation of $D$ w.r.t. $\nu$ and $\gamma$ satisfying the conditions of Lemma 4.7 (i.e., $v^{*}(s)=v^{*}\left(s^{\prime}\right)$ iff $D \models e$ equiv $_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$, for all $\left.s, s^{\prime} \in \operatorname{adom}(D) \cup \operatorname{adom}(\gamma)\right)$. Then $v^{*}(D) \models \gamma^{\prime}\left(v^{*}\left(\bar{t}^{\prime}\right)\right)$ iff for all valuations $v, v(D) \models \gamma\left(v(\bar{t})\right.$ implies $v(D) \models \gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$.

## Proof of the claim

$\Rightarrow$ Assume $v^{*}(D) \models \gamma^{\prime}\left(v^{*}\left(\bar{t}^{\prime}\right)\right)$ and let $v$ be a valuation such that $v(D) \models \gamma(v(\bar{t})$. We want to show $v(D) \models \gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$. By Lemma 4.5 it is enough to show that $\forall s, s^{\prime} \in \overline{t^{\prime}}$, $D \models \operatorname{equiv}_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right)$ implies $v(D) \models v(s)=v\left(s^{\prime}\right)$. So let $s, s^{\prime} \in \overline{t^{\prime}}$ such that $D \models$ equiv $\gamma_{\gamma^{\prime}}\left(\overline{t^{\prime}}, s, s^{\prime}\right)$. As $v^{*}(D) \models \gamma^{\prime}\left(v^{*}\left(\overline{t^{\prime}}\right)\right.$, by Lemma $4.5, v^{*}(D) \models v^{*}(s)=v^{*}\left(s^{\prime}\right)$. Now $v^{*}$ satisfies the conditions of Lemma 4.7, so $D \models \operatorname{equiv}_{\gamma}\left(\bar{t}, s, s^{\prime}\right)$. As $v(D) \models \gamma(v(\bar{t})$, by Lemma 4.5 it follows that $v(D) \models v(s)=v\left(s^{\prime}\right)$.
$\Leftarrow$ Assume for all valuations $v, v(D) \vDash \gamma(v(\bar{t}))$ implies $v(D) \vDash \gamma^{\prime}\left(v\left(\bar{t}^{\prime}\right)\right)$. By Lemma 4.5, $v^{*}$ satisfying the conditions of Lemma 4.7, we have $v^{*}(D) \models \gamma\left(v^{*}(\bar{t})\right)$ and so by our assumption $v^{*}(D) \models \gamma^{\prime}\left(v^{*}\left(\overline{t^{\prime}}\right)\right)$.

As $v$ satisfies the conditions of Lemma 4.7, by Claim 4.25 it follows from $v(D) \models=$ $e q_{j}\left(v\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)\right)$ that $\forall v$ with $v(D) \models e q_{i}(v(\bar{s} \bar{a} \bar{b}))$, also $v(D) \models e q_{i}\left(v\left(\bar{t} \bar{a}^{\prime} \bar{b}^{\prime}\right)\right)$. Now by Proposition $4.8 D \models i m p l y_{e q_{i}, e q_{j}}\left(\bar{s} \bar{a} \bar{b}, \bar{t}^{\prime} \bar{a}^{\prime} \bar{b}^{\prime}\right) \vee \neg \operatorname{comp}_{e q_{i}}(\bar{s}, \bar{a}, \bar{b})$. But $D \models \operatorname{comp}_{e q_{i}}(\bar{s}, \bar{a}, \bar{b})$, so $D \models \exists \bar{y} \bar{z}\left(q_{j}(\bar{y}, \bar{z}) \wedge i m p l y_{e q_{i}, e q_{j}}(\bar{s} \bar{a} \bar{b}, \bar{t} \bar{y} \bar{z})\right)$.

Recall that from $Q_{\subseteq}$ one can easily define a first-order rewriting $\operatorname{best}_{Q}(\bar{x})$ for best answers as in (1).

## Theorem 4.26

Given $Q$ a union of conjunctive queries over schema $\sigma$ and an incomplete database $D$, $\bar{t} \in \operatorname{best}(Q, D)$ iff $D \models \operatorname{best}_{Q}(\bar{t})$.

## Proof

By Proposition 4.24, $D \models$ best $_{Q}(\bar{t})$ if and only if $\forall \operatorname{sinpp}(Q, D, \bar{t}) \subseteq \operatorname{Supp}(Q, D, \bar{s})$ implies $\operatorname{Supp}(Q, D, \bar{s}) \subseteq \operatorname{Supp}(Q, D, \bar{t})$. Notice that this holds exactly whenever $\neg \exists \bar{s}$ with $\operatorname{Supp}(Q, D, \bar{t}) \subset \operatorname{Supp}(Q, D, \bar{s})$, that is, whenever $\bar{t} \in \operatorname{best}(Q, D)$.

Example 4.27
For $Q, D, \gamma, \gamma^{\prime}$ as in Examples 2.1 and 4.22:

$$
\begin{aligned}
& Q_{\subseteq}\left(x, x^{\prime}\right):=\forall y_{1} y_{2} z\left(\left(R\left(y_{1}\right) \wedge S\left(y_{2}, z\right) \wedge \operatorname{comp}_{\gamma}\left(y_{1}, y_{2}, z, x\right)\right)\right. \\
& \rightarrow \\
& \exists y_{1}^{\prime} y_{2}^{\prime} z^{\prime}\left(R\left(y_{1}^{\prime}\right) \wedge S\left(y_{2}^{\prime}, z^{\prime}\right) \wedge\right.\left.\left.\operatorname{imply}_{\gamma, \gamma^{\prime}}\left(y_{1} y_{2} z x, y_{1}^{\prime} y_{2}^{\prime} z^{\prime} x^{\prime}\right)\right)\right)
\end{aligned}
$$

This allows to derive for instance $\operatorname{Supp}(Q, D, 1) \subseteq \operatorname{Supp}\left(Q, D, \perp_{2}\right)$ (as observed in Example 2.1). In fact the subquery $R\left(y_{1}\right) \wedge S\left(y_{2}, z\right) \wedge \operatorname{comp}_{\gamma}\left(y_{1}, y_{2}, z, x\right)$ with free variables $y_{1}, y_{2}, z, x$ selects on $D$ tuples $\left(1, \perp_{2}, \perp_{2}, 1\right),\left(\perp_{1}, \perp_{2}, \perp_{2}, 1\right)$, and no other tuple with last element 1. Moreover, as shown in Example 4.22

$$
\begin{aligned}
& D \models \text { imply } y_{\gamma \gamma^{\prime}}\left(\perp_{1} \perp_{2} \perp_{2} 1,1 \perp_{2} \perp_{2} \perp_{2}\right) \\
& D \models \text { imply } y_{\gamma \gamma^{\prime}}\left(1 \perp_{2} \perp_{2} 1,1 \perp_{2} \perp_{2} \perp_{2}\right)
\end{aligned}
$$

Thus, $D \vDash Q_{\subseteq}\left(1, \perp_{2}\right)$. Similarly, one can show $D \models Q_{\subseteq}\left(\perp_{1}, \perp_{2}\right)$, and therefore $D \models$ $\operatorname{best}_{Q}\left(\perp_{2}\right)$.

As a corollary of Theorem 4.26 , for a union of conjunctive queries $Q$ one can compute best $(Q, D)$ by first computing the formula best $_{Q}(\bar{x})$ from $Q$, then evaluating best ${ }_{Q}$ on $D$. Since data complexity of FO query evaluation is DLogSpace (and in particular $\mathrm{AC}^{0}$ ), this gives the following corollary:

Corollary 4.28
For each fixed union of conjunctive queries $Q$, the data complexity of Bestanswer $_{\Sigma}$ is DLogSpace.

Note that it was known from Libkin (2018) that the data complexity of computing best answers for unions of conjunctive queries is polynomial time. In terms of combined complexity (i.e. when either $Q, D$ and $\bar{a}$ are in the input), the rewriting approach (i.e. the procedure of computing best $_{Q}$ from $Q$ and then evaluating best $Q_{Q}$ on $D$ ), can be easily shown to be in PSPACE. In fact it is well known that a first-order query $\phi$ can be evaluated on a database $D$ in space at most $q r(\phi) \log |D|+\log |\phi|$, where $q r(\phi)$ is the quantifier rank of $\phi$. Note that although best ${ }_{Q}$ has size exponential in $Q$, the quantifier rank of best ${ }_{Q}$ is linear in the size of $Q$. Thus whether $\bar{a} \in \operatorname{best}(Q, D)$ can be checked using space $O(|Q|,|D|)$. Moreover one can easily check that best ${ }_{Q}$ can be computed from $Q$ in space polynomial in the size of $|Q|$. Since space bounded computations can be composed without storing the intermediate output, computing best ${ }_{Q}$ from $Q$ and then evaluating best $_{Q}$ on $D$ can be done overall in PSPACE in the size of $|Q|$ and $|D|$. The rewriting approach thus implies a PSPACE upper bound for the combined complexity of BESTANSWER for unions of conjunctive queries. However, we can show that the problem actually stands in the third level of the polynomial hierarchy.

## Theorem 4.29

For unions of conjunctive queries, combined complexity of BESTANSWER ${ }_{\Sigma}$ is $\Pi_{3}^{p}$ complete. Hardness already holds for conjunctive queries.

## Proof

For membership, first note that one can check in $\Pi_{2}^{p}$ whether $\operatorname{Supp}(Q, D, \bar{a}) \subseteq$ $\operatorname{Supp}(Q, D, \bar{b})$ on input given by a database $D$, a UCQ $Q$, and tuples $\bar{a}$ and $\bar{b}$. In fact in order to check $\operatorname{Supp}(Q, D, \bar{a}) \nsubseteq \operatorname{Supp}(Q, D, \bar{b})$ one guesses a valuation $v$ of $D$, then calls an NP oracle to check $v(\bar{a}) \in Q(v(D))$ and $v(\bar{b}) \notin Q(v(D))$.

On input given by a database $D$, a UCQ $Q$, and a tuple $\bar{a}$ one can check $\bar{a} \notin$ best $(Q, D)$ as follows. First guess a tuple $\bar{b}$ over $\operatorname{adom}(D)$ of the same arity as $\bar{a}$; then, using two
calls to a $\Sigma_{2}^{p}$ oracle, check that $\operatorname{Supp}(Q, D, \bar{a}) \subseteq \operatorname{Supp}(Q, D, \bar{b})$ and $\operatorname{Supp}(Q, D, \bar{b}) \nsubseteq$ $\operatorname{Supp}(Q, D, \bar{a})$.

For hardness, we reduce from $\forall \exists \forall 3 D N F$, which is known to be $\Pi_{3}^{p}$-complete (Schaefer and Umans 2002). We take as input a $\forall \exists \forall 3 D N F$-formula of the form

$$
F:=\forall z_{1}, \ldots z_{l} \exists x_{1} \ldots x_{k} \forall y_{1} \ldots y_{p} \bigvee_{i=1}^{n} \operatorname{conj} j_{i}
$$

where the each $\operatorname{con}_{i}$ is a conjunction of 3 (not necessarily distinct) literals over variables $z_{1}, \ldots z_{l}, x_{1}, \ldots, x_{k}, y_{1}, \ldots, y_{p}$.

We construct a database $D_{F}$ with adom $\left(D_{F}\right)=\{0,1$, good, bad $\} \cup\left\{i, \bar{i}, \perp_{i}, \bar{\perp}_{i}, \mid i=\right.$ $1 . . k\}$, and a conjunctive query $Q_{F}\left(z_{1}, . . z_{l}, z\right)$ such that $(\overline{0}$, good $) \in \operatorname{best}\left(Q_{F}, D_{F}\right)$ if and only if $F$ is true.
$D_{F}$ is of signature $\left\{S^{4}, C^{2}, A^{2}, B^{3}\right\}$ as follows:

- The extension of $S$ and $A$ and $B$ is fixed and does not depend on $F$ :
- $S$ contains tuple ( $1,1,1$, good), and tuples $\left(b_{1}, b_{2}, b_{3}\right.$, good) and ( $b_{1}, b_{2}, b_{3}$, bad $)$ for every $b_{1}, b_{2}, b_{3} \in\{0,1\}$ with $\left(b_{1}, b_{2}, b_{3}\right) \neq(1,1,1)$. Intuitively, $S$ encodes the possible truth assignment of each disjunct of $F$. Note that only the satisfying assignment (i.e. $(1,1,1)$ ) appears together with the only constant good, all the others appear both with good and bad.
- $A$ contains only two tuples: $(0,1)$ and $(1,0)$. Intuitively, $A$ will be used to encode truth values for pairs of literals $(w, \neg w), w \in y_{1}, \ldots y_{p}, z_{1}, \ldots z_{l}$.
- $B$ contains tuples $(0,0, b a d),(1,1$, bad $)$ and tuples $\left(b_{1}, b_{2}\right.$, good $)$ and $\left(b_{1}, b_{2}\right.$, bad $)$ for every $b_{1}, b_{2} \in\{0,1\}, b_{1} \neq b_{2}$. Intuitively, $B$ encodes assignments for pairs of literals $(w, \neg w), w \in\left\{x_{1}, \ldots x_{k}\right\}$. Note that here inconsistent pairs (i.e. same truth value) are possible, but these are the only ones which do not appear together with constant good.
- The extension of $C$ depends on $F$ and contains tuples $\left\{\left(\perp_{i}, i\right) \mid i=1 . . k\right\}$ and $\left\{\left(\bar{\perp}_{i}, \bar{i}\right) \mid i=\right.$ $1 . . k\}$. Intuitively, a valuation $(b, i)$ (resp. $(b, \bar{i}))$ of one of these tuples, with $b \in\{0,1\}$, will encode truth value $b$ for the literal $x_{i}\left(\right.$ resp,$\left.\neg x_{i}\right)$ of $F$.
$Q_{F}$ is defined as follows. For each variable $w$ of $F$, the conjunctive query $Q_{F}$ will use variables $w$ and $\bar{w}$ (either quantified or free). For a literal $\alpha$ of $F$, the corresponding variable of $Q_{F}$ will be denoted as $\operatorname{enc}(\alpha)$. More precisely if $\alpha=w$ is a positive literal, then, $\operatorname{enc}(\alpha):=w$, otherwise if $\alpha=\neg w$ then $\operatorname{enc}(\alpha):=\bar{w}$.

$$
\begin{aligned}
Q_{F}\left(z_{1}, \ldots z_{l}, z\right):= & \exists x_{1}, \ldots x_{k}, \bar{x}_{1}, \ldots \bar{x}_{k}, y_{1}, \ldots y_{p}, \bar{y}_{1}, \ldots \bar{y}_{p}, \bar{z}_{1}, \ldots \bar{z}_{p} \\
& \bigwedge_{i=1, . . k} B\left(x_{i}, \bar{x}_{i}, z\right) \wedge \bigwedge_{i=1, . . p} A\left(y_{i}, \bar{y}_{i}\right) \wedge \bigwedge_{i=1, . . l} A\left(z_{i}, \bar{z}_{i}\right) \wedge \\
& \bigwedge_{i=1, . . k}\left(C\left(x_{i}, i\right) \wedge C\left(\bar{x}_{i}, \bar{i}\right)\right) \wedge \\
& \bigwedge_{\left(\alpha_{1} \wedge \alpha_{2} \wedge \alpha_{3}\right) \in F} S\left(\operatorname{enc}\left(\alpha_{1}\right), \operatorname{enc}\left(\alpha_{2}\right), \operatorname{enc}\left(\alpha_{3}\right), z\right)
\end{aligned}
$$

We can prove that all tuples of the form ( $\bar{t}$, good) (which we refer to as good tuples) have the same support. This is given by the set of all consistent boolean valuations (i.e.
valuations of $\perp_{i}, \bar{\perp}_{i}$ in $\{0,1\}$ such that $v\left(\perp_{i}\right) \neq v\left(\bar{\perp}_{i}\right)$ for all $\left.i\right)$. Moreover we can prove that if there exists a $(\bar{t}, b a d)$ whose support contains all consistent boolean valuations then the support of $(\bar{t}, b a d)$ strictly contains the support of good tuples. Therefore, any good tuple (including $(\overline{0}$, good $)$ ) is a best answer iff for all tuples $\bar{t}$ there exists a consistent boolean valuation which is not in the support of $(\bar{t}, b a d)$. We can finally show that the last holds iff $F$ is true.

Therefore, under standard complexity theoretic assumptions, our rewriting approach is not optimal in terms of combined complexity, as it is often the case with generic approaches. However, it has the advantage of exploiting standard FO query evaluation, which despite the PSPACE combined complexity is highly optimized in database systems and works well in practice.

## 5 Future work

Our rewriting techniques are closer to a practical implementation than the previous tableau-based method from Gheerbrant and Libkin (2015). This is due to their expressibility in recursive SQL (or even non-recursive in the case of Theorems 4.19 and 4.26). However, while theoretically feasible, an actual implementation will need additional techniques to achieve acceptable performance. To see why, notice that the first rule in the definition of equiv ${ }_{\gamma}$ creates a cross product over the full active domain, that is, the set of all elements that appeared in the database. This of course will be prohibitively large. While this may appear to be a significant obstacle, a similar situation with computing or approximating certain answers is not new in the literature. For instance, the first approximation scheme for certain answers to SQL queries that appeared in Libkin (2016) has done exactly the same, and generated very large Cartesian products even for simple queries with negation. Nonetheless, an alternative was found quickly (Guagliardo and Libkin 2016) that completely avoided the need for such expensive queries, and it was shown to work well on several TPC-H queries. Thus, looking for a practical and implementable rewriting is one of the possible directions for future work.

As another open problem, we note that the query for which we have shown certain answers to be non-rewritable in FO has DLOGSPACE data complexity. Indeed the problem is essentially reachability over trees, which can be easily encoded using deterministic transitive closure (Immerman 1987). To express DLOGSPACE problems, we need a language weaker than Datalog with negation. Thus, it is natural to ask whether a low complexity Datalog fragment would be sufficient to express rewritings of BCCQ, or a separating example that is PTIME-complete can be found.

Another direction would be to investigate how our techniques can be extended to different semantics of incompleteness. We used here the closed-world semantics (Abiteboul et al. 1995; Imielinski and Lipski 1984; van der Meyden 1998), in which data values are the only missing information, but there are other possible semantics, for example needed in order to cope with data inconsistencies (Calì et al. 2003a), where query rewritings could still be found. Quantitative variations of the notion of certainty as proposed in Libkin (2018) could also be investigated.

## References

Abiteboul, S., Hull, R. and Vianu, V. 1995. Foundations of Databases. Addison-Wesley.
Abiteboul, S., Kanellakis, P. C. and Grahne, G. 1991. On the representation and querying of sets of possible worlds. Theoretical Computer Science 78, 1, 158-187.
Arenas, M., Barceló, P., Libkin, L. and Murlak, F. 2014. Foundations of Data Exchange. Cambridge University Press.
Arenas, M., Pérez, J. and Reutter, J. L. 2013. Data exchange beyond complete data. Journal of the ACM 60, 4, 28:1-28:59.
Barceló, P., Libkin, L. and Romero, M. 2014. Efficient approximations of conjunctive queries. SIAM Journal on Computing 43, 3, 1085-1130.
Bertossi, L. E. 2011. Database Repairing and Consistent Query Answering. Synthesis Lectures on Data Management. Morgan \& Claypool Publishers.
Bienvenu, M. and Bourgaux, C. 2016. Inconsistency-tolerant querying of description logic knowledge bases. In Reasoning Web: Logical Foundation of Knowledge Graph Construction and Query Answering - 12th International Summer School 2016, Aberdeen, UK, 5-9 September 2016, Tutorial Lectures, J. Z. Pan, D. Calvanese, T. Eiter, I. Horrocks, M. Kifer, F. Lin and Y. Zhao, Eds. Lecture Notes in Computer Science, vol. 9885. Springer, 156-202.

Bienvenu, M. and Ortiz, M. 2015. Ontology-mediated query answering with data-tractable description logics. In Reasoning Web. Web Logic Rules - 11th International Summer School 2015, Berlin, Germany, 31 July-4 August 2015, Tutorial Lectures, W. Faber and A. Paschke, Eds. Lecture Notes in Computer Science, vol. 9203. Springer, 218-307.
Buss, S. R. and Hay, L. 1991. On truth-table reducibility to SAT. Information and Computation 91, 1, 86-102.
Calì, A., Calvanese, D. and Lenzerini, M. 2013. Rewrite and conquer: Dealing with integrity constraints in data integration. In Seminal Contributions to Information Systems Engineering, 25 Years of CAiSE, J. Bubenko Jr., J. Krogstie, O. Pastor, B. Pernici, C. Rolland and A. Sølvberg, Eds. Springer, 353-359.
Calì, A., Lembo, D. and Rosati, R. 2003a. On the decidability and complexity of query answering over inconsistent and incomplete databases. In Proceedings of the Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 9-12 June 2003, San Diego, CA, USA, F. Neven, C. Beeri and T. Milo, Eds. ACM, 260-271.
Calì, A., Lembo, D. and Rosati, R. 2003b. Query rewriting and answering under constraints in data integration systems. In IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico, 9-15 August 2003, G. Gottlob and T. Walsh, Eds. Morgan Kaufmann, 16-21.

Calvanese, D., Giacomo, G. D., Lembo, D., Lenzerini, M. and Rosati, R. 2006. Epistemic first-order queries over description logic knowledge bases. In Proceedings of the 2006 International Workshop on Description Logics (DL2006), Windermere, Lake District, UK, 30 May-1 June 2006, B. Parsia, U. Sattler and D. Toman, Eds. CEUR Workshop Proceedings, vol. 189. CEUR-WS.org.
Calvanese, D., Giacomo, G. D., Lenzerini, M. and Vardi, M. Y. 2000. What is view-based query rewriting? In Proceedings of the 7th International Workshop on Knowledge Representation meets Databases (KRDB 2000), Berlin, Germany, 21 August 2000, M. Bouzeghoub, M. Klusch, W. Nutt and U. Sattler, Eds. CEUR Workshop Proceedings, vol. 29. CEUR-WS.org, 17-27.
Calvanese, D., Giacomo, G. D., Lenzerini, M. and Vardi, M. Y. 2007. View-based query processing: On the relationship between rewriting, answering and losslessness. Theoretical Computer Science 371, 3, 169-182.
Coelho, F. 2013. https://blog.coelho.net/database/2013/08/17/turing-sql-1.html.
Eiter, T. and Gottlob, G. 1997. The complexity class theta ${ }_{2}$ p: Recent results and applications in AI and modal logic. In Fundamentals of Computation Theory, 11th International

Symposium, FCT'97, Kraków, Poland, 1-3 September 1997, Proceedings, B. S. Chlebus and L. Czaja, Eds. Lecture Notes in Computer Science, vol. 1279. Springer, 1-18.

Fagin, R., Kolaitis, P. G. and Popa, L. 2005. Data exchange: Getting to the core. ACM Transactions on Database Systems 30, 1, 174-210.
Fagin, R., Stockmeyer, L. J. and Vardi, M. Y. 1995. On monadic NP vs. monadic co-np. Information and Computation 120, 1, 78-92.
Geerts, F., Mecca, G., Papotti, P. and Santoro, D. 2013. The LLUnAtiC data-cleaning framework. Proceedings of the VLDB Endowment 6, 9, 625-636.
Gheerbrant, A. and Libkin, L. 2015. Certain answers over incomplete XML documents: Extending tractability boundary. Theory of Computing Systems 57, 4, 892-926.
Gheerbrant, A., Libkin, L., Rogova, A. and Sirangelo, C. 2022. Certain answers of extensions of conjunctive queries by datalog and first-order rewriting. In Proceedings of the 4th International Workshop on the Resurgence of Datalog in Academia and Industry (Datalog-2.0 2022) Co-located with the 16th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2022), Genova-Nervi, Italy, 5 September 2022, M. Alviano and A. Pieris, Eds. CEUR Workshop Proceedings, vol. 3203. CEUR-WS.org, 14-26.
Gheerbrant, A., Libkin, L. and Sirangelo, C. 2014. Naïve evaluation of queries over incomplete databases. ACM Transactions on Database Systems 39, 4, 31:1-31:42.
Gheerbrant, A. and Sirangelo, C. 2019. Best answers over incomplete data: Complexity and first-order rewritings. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, 10-16 August 2019, S. Kraus, Ed. ijcai.org, 1704-1710.
Gierth, A. 2011. https://wiki.postgresql.org/wiki/Cyclic_Tag_System.
Gottlob, G. 1995. NP trees and carnap's modal logic. Journal of the ACM 42, 2, 421-457.
Greco, S., Molinaro, C. and Spezzano, F. 2012. Incomplete Data and Data Dependencies in Relational Databases. Synthesis Lectures on Data Management. Morgan \& Claypool Publishers.
Guagliardo, P. and Libkin, L. 2016. Making SQL queries correct on incomplete databases: A feasibility study. In Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, 26 June-01 July 2016, T. Milo and W. Tan, Eds. ACM, 211-223.
Hella, L., Libkin, L., Nurmonen, J. and Wong, L. 2001. Logics with aggregate operators. Journal of the ACM 48, 4, 880-907.
Imielinski, T. and Lipski, W. J. 1984. Incomplete information in relational databases. Journal of the ACM 31, 4, 761-791.
Immerman, N. 1987. Languages that capture complexity classes. SIAM Journal on Computing 16, 4, 760-778.
Lenzerini, M. 2002. Data integration: A theoretical perspective. In Proceedings of the Twentyfirst ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 3-5 June, Madison, Wisconsin, USA, L. Popa, S. Abiteboul and P. G. Kolaitis, Eds. ACM, 233246.

Libkin, L. 2003. Expressive power of SQL. Theoretical Computer Science 296, 3, 379-404.
Libkin, L. 2016. Sql's three-valued logic and certain answers. ACM Transactions on Database Systems 41, 1, 1:1-1:28.
Libkin, L. 2018. Certain answers meet zero-one laws. In Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, 10-15 June 2018, J. V. den Bussche and M. Arenas, Eds. ACM, 195-207.
Lipski, W. J. 1984. On relational algebra with marked nulls. In Proceedings of the Third ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, 2-4 April 1984, Waterloo, Ontario, Canada, D. J. Rosenkrantz and R. Fagin, Eds. ACM, 201-203.

Reiter, R. 1977. On closed world data bases. In Logic and Data Bases, Symposium on Logic and Data Bases, Centre d'études et de recherches de Toulouse, France, 1977, H. Gallaire and J. Minker, Eds. Advances in Data Base Theory. Plemum Press, New York, 55-76.

Rothe, J. 2003. Exact complexity of exact-four-colorability. Information Processing Letters 87, 1, 7-12.
Schaefer, M. and Umans, C. 2002. Completeness in the polynomial-time hierarchy a compendium. Sigact News - SIGACT 33, 32-49.
van der Meyden, R. 1998. Logical approaches to incomplete information: A survey. In Logics for Databases and Information Systems (the book grow out of the Dagstuhl Seminar 9529: Role of Logics in Information Systems, 1995), J. Chomicki and G. Saake, Eds. Kluwer, 307-356.
Wagner, K. W. 1990. Bounded query classes. SIAM Journal on Computing 19, 5, 833-846.


[^0]:    ${ }^{1}$ Queries we write hereafter can be domain dependent. So it is important to recall that we always use active domain semantics.

