# Complexity of Inconsistency-Tolerant Query Answering in Datalog+/- under Cardinality-Based Repairs 

Thomas Lukasiewicz<br>Department of Computer Science<br>University of Oxford, UK

Enrico Malizia<br>Department of Computer Science<br>University of Exeter, UK

Andrius Vaicenavičius<br>Department of Computer Science<br>University of Oxford, UK


#### Abstract

Querying inconsistent ontological knowledge bases is an important problem in practice, for which several inconsistencytolerant query answering semantics have been proposed, including query answering relative to all repairs, relative to the intersection of repairs, and relative to the intersection of closed repairs. In these semantics, one assumes that the input database is erroneous, and the notion of repair describes a maximally consistent subset of the input database, where different notions of maximality (such as subset and cardinality maximality) are considered. In this paper, we give a precise picture of the computational complexity of inconsistencytolerant (Boolean conjunctive) query answering in a wide range of Datalog ${ }^{ \pm}$languages under the cardinality-based versions of the above three repair semantics.


## Introduction

In many ontology-based applications in practice, such as ontology-based data extraction from the Web, or ontologybased integration of different data sources, it is very likely that the data are inconsistent with the ontology, and thus inconsistency-tolerant semantics for ontology-based query answering are urgently needed. Among the most prominent ontology languages are description logics (DLs) and existential rules from the context of Datalog ${ }^{ \pm}$.

The most widely accepted semantics for querying inconsistent ontological knowledge bases is perhaps consistent query answering, which was first developed for relational databases (Arenas, Bertossi, and Chomicki 1999) and then generalized as the ABox repair (AR) semantics for several DLs (Lembo et al. 2010). Consistent query answering is based on the concept of repair, which is a maximal consistent subset of the input database. A fact/query is entailed by an ontological knowledge base in consistent query answering, if it is (classically) entailed by all the repairs (under the ontology). Several other repair semantics for querying inconsistent ontological knowledge bases have recently been developed as alternatives to consistent query answering (Lembo et al. 2010; Bienvenu 2012; Lukasiewicz, Martinez, and Simari 2012; Bienvenu and Rosati 2013). In particular, in (Lembo et al. 2010), besides consistent query answering, three other inconsistency-tolerant query answering

[^0]semantics are proposed, including the intersection of repairs (IAR) semantics, in which an answer is considered to be valid, if it can be inferred from the intersection of the repairs (and the ontology). The intersection of closed repairs (ICR) (Bienvenu 2012) is another semantics, in which an answer is valid, if it can be inferred from the intersection of the closure of the repairs (and the ontology).

There are several reasons for the practical relevance of the IAR and the ICR semantics, and thus for motivating an indepth analysis of their computational properties. First, they are two natural semantics that identify "surer" answers than consistent query answering, and so they can also be seen as under-approximations of the latter. Investigating their complexity helps to understand whether such approximations have actually lower complexities, which is the case for different languages and one complexity measure considered in this paper. Second, recent work on explanations in the context of inconsistency-tolerant query answering shows that explanations are much easier to define and compute for the IAR semantics (Bienvenu, Bourgaux, and Goasdoué 2016). Third, a crucial advantage of the IAR and the ICR semantics is that their intersection of (closed) repairs can be materialized, while the consistent query answering semantics exists only virtually (the intersection of (closed) repairs can be computed offline, and then standard querying algorithms can be employed online)-indeed, this has been used to implement the IAR semantics (Lembo et al. 2015), and for ICR, it has been remarked in (Bienvenu and Bourgaux 2016).

The complexity of consistent query answering when the ontology is described via one of the main DLs is well-understood. Rosati (2011) studied the data and combined complexity for a wide spectrum of DLs, while Bienvenu (2012) identified cases for simple ontologies (within the DL-Lite family) for which tractable data complexity results can be obtained. In (Lukasiewicz, Martinez, and Simari 2012; 2013; Lukasiewicz et al. 2015), the data and different types of combined complexity of consistent query answering have been studied for ontologies described via existential rules and negative constraints, and such investigation has recently been extended to the IAR and ICR semantics in (Lukasiewicz, Malizia, and Molinaro 2018a).

Alternative notions of maximality for repairs, however, such as cardinality-maximal repairs (Lopatenko and Bertossi 2007), rather than subset-maximal ones, have been
explored less. Bienvenu, Bourgaux, and Goasdoué (2014) analyzed the data and the combined complexity of query answering under the AR and IAR semantics over the language DL-Lite $\mathcal{R}_{\mathcal{R}}$ for different notions of maximal repairs, among which the notion of maximum cardinality repairs.

This paper continues this line of research on cardinalitymaximal consistent query answering, and we analyze the complexity of the above three inconsistency-tolerant query answering semantics for a wide range of Datalog ${ }^{ \pm}$languages and for several different complexity measures:
$\triangleright$ We consider different popular inconsistency-tolerant semantics, namely, the AR, the IAR, and the ICR semantics, with maximum cardinality database repairs.
$\triangleright$ We consider the most popular Datalog ${ }^{ \pm}$languages: linear, guarded, sticky, and acyclic existential rules, along with "weak" generalizations, as well as full (i.e., nonexistential) restrictions, and full rules in general.
$\triangleright$ We analyze the data, fixed-program combined, boundedarity combined, and combined complexity.

Detailed proofs are given in a forthcoming extended paper.

## Preliminaries

We now briefly recall some basics on Datalog ${ }^{ \pm}$(Calì, Gottlob, and Lukasiewicz 2012) and the complexity classes that we will encounter in our analysis in this paper.
General. We assume sets $\mathbf{C}, \mathbf{N}$, and $\mathbf{V}$ of constants, labeled nulls, and variables, respectively. A term $t$ is a constant, null, or variable. We also assume a set of predicates, each with an arity, i.e., a non-negative integer. An atom has the form $p\left(t_{1}, \ldots, t_{n}\right)$, where $p$ is an $n$-ary predicate, and $t_{1}, \ldots, t_{n}$ are terms. Conjunctions of atoms are often identified with the sets of their atoms. An instance $I$ is a (possibly infinite) set of atoms $p(\mathbf{t})$, where $\mathbf{t}$ is a tuple of constants and nulls. A database $D$ is a finite instance without nulls. A homomorphism is a mapping $h: \mathbf{C} \cup \mathbf{N} \cup \mathbf{V} \rightarrow \mathbf{C} \cup \mathbf{N} \cup \mathbf{V}$ that is the identity on $\mathbf{C}$ and maps $\mathbf{N}$ to $\mathbf{C} \cup \mathbf{N}$. A conjunctive query ( CQ ) $Q$ has the form $\exists \mathbf{Y} \phi(\mathbf{X}, \mathbf{Y})$, where $\phi(\mathbf{X}, \mathbf{Y})$ is a conjunction of atoms without nulls. The answer to $Q$ over an instance $I$, denoted $Q(I)$, is the set of all tuples $\mathbf{t}$ over $\mathbf{C}$ for which there is a homomorphism $h$ such that $h(\phi(\mathbf{X}, \mathbf{Y})) \subseteq I$ and $h(\mathbf{X})=\mathbf{t}$. A Boolean $C Q(\mathrm{BCQ}) Q$ is a $\mathrm{CQ} \exists \mathbf{Y} \phi(\mathbf{Y})$, i.e., all variables are existentially quantified; $Q$ is true over $I$, denoted $I \models Q$, if $Q(I) \neq \emptyset$, i.e., there is a homomorphism $h$ with $h(\phi(\mathbf{Y})) \subseteq I$.
Dependencies. A tuple-generating dependency (TGD) $\sigma$ is a first-order formula $\forall \mathbf{X} \forall \mathbf{Y} \varphi(\mathbf{X}, \mathbf{Y}) \rightarrow \exists \mathbf{Z} p(\mathbf{X}, \mathbf{Z})$, where $\mathbf{X} \cup \mathbf{Y} \cup \mathbf{Z} \subseteq \mathbf{V}, \varphi(\mathbf{X}, \mathbf{Y})$ is a conjunction of atoms, and $p(\mathbf{X}, \mathbf{Z})$ is an atom, all without nulls; $\varphi(\mathbf{X}, \mathbf{Y})$ is the body of $\sigma$, denoted body $(\sigma)$, while $p(\mathbf{X}, \mathbf{Z})$ is the head of $\sigma$, denoted head $(\sigma)$. For clarity, we consider single-atom-head TGDs; however, our results can be extended to TGDs with a conjunction of atoms in the head. An instance $I$ satisfies $\sigma$, written $I \models \sigma$, if the following holds: whenever there exists a homomorphism $h$ such that $h(\varphi(\mathbf{X}, \mathbf{Y})) \subseteq I$, then there exists $\left.h^{\prime} \supseteq h\right|_{\mathbf{x}}$, where $\left.h\right|_{\mathbf{x}}$ is the restriction of $h$ on $\mathbf{X}$, such that $h^{\prime}(p(\mathbf{X}, \mathbf{Z})) \in I$. A negative constraint $(N C) \nu$ is a first-order formula $\forall \mathbf{X} \varphi(\mathbf{X}) \rightarrow \perp$, where $\mathbf{X} \subseteq \mathbf{V}, \varphi(\mathbf{X})$
is a conjunction of atoms without nulls, called the $\operatorname{body}$ of $\nu$, denoted $\operatorname{body}(\nu)$, and $\perp$ denotes the truth constant false. An instance $I$ satisfies $\nu$, written $I \models \nu$, if there is no homomorphism $h$ such that $h(\varphi(\mathbf{X})) \subseteq I$. Given a set $\Sigma$ of TGDs and NCs, $I$ satisfies $\Sigma$, written $I \models \Sigma$, if $I$ satisfies each TGD and NC of $\Sigma$. For brevity, we omit the universal quantifiers in front of TGDs and NCs, and use the comma (instead of $\wedge$ ) for conjoining body atoms. Given a class of TGDs $\mathbb{C}$, we denote by $\mathbb{C}_{\perp}$ the formalism obtained by combining $\mathbb{C}$ with arbitrary NCs. Finite sets of TGDs and NCs are also called programs, and TGDs are also called existential rules.
Knowledge Bases. A knowledge base is a pair $(D, \Sigma)$, where $D$ is a database, and $\Sigma$ is a program. For programs $\Sigma$, $\Sigma_{T}$ and $\Sigma_{N C}$ are the subsets of $\Sigma$ containing the TGDs and NCs of $\Sigma$, respectively. The set of models of $K B=(D, \Sigma)$, denoted mods $(K B)$, is the set of instances $\{I \mid I \supseteq D, I \models$ $\Sigma\}$. We say that $K B$ is consistent, if $\operatorname{mods}(K B) \neq \emptyset$, otherwise $K B$ is inconsistent. The answer to a CQ $Q$ relative to $K B$ is the set of tuples $\operatorname{ans}(Q, K B)=\bigcap\{Q(I) \mid$ $I \in \operatorname{mods}(K B)\}$. The answer to a BCQ $Q$ is true, denoted $K B \vDash Q$, if ans $(Q, K B) \neq \emptyset$. The decision version of the $C Q$ answering problem is as follows: given a knowledge base $K B$, a CQ $Q$, and a tuple of constants $\mathbf{t}$, decide whether $\mathbf{t} \in \operatorname{ans}(Q, K B)$. Since CQ answering can be reduced in LogSpace to BCQ answering, we focus on BCQs. Following Vardi (1982), the combined complexity of BCQ answering considers the database, the set of dependencies, and the query as part of the input. The bounded-arity combined (or ba-combined) complexity assumes that the arity of the underlying schema is bounded by an integer constant. The fixed-program combined (or fp-combined) complexity considers the sets of TGDs and NCs as fixed; the data complexity also assumes the query fixed.
Decidability Paradigms. The main (syntactic) conditions on TGDs that guarantee the decidability of BCQ answering are guardedness (Calì, Gottlob, and Kifer 2013), stickiness (Calì, Gottlob, and Pieris 2012), and acyclicity, each having a "weak" counterpart: weak guardedness (Calì, Gottlob, and Kifer 2013), weak stickiness (Calì, Gottlob, and Pieris 2012), and weak acyclicity (Fagin et al. 2005).

A TGD $\sigma$ is guarded, if there exists an atom in the body that contains (or "guards") all the body variables of $\sigma$. The class of guarded TGDs, denoted G, is the family of all possible sets of guarded TGDs. A key subclass of guarded TGDs are linear TGDs with just one body atom (which is a guard), and the corresponding class is denoted L. Weakly guarded TGDs extend guarded TGDs by requiring only "harmful" body variables to appear in the guard, and the associated class is denoted WG. It is easy to verify that $L \subset G \subset W G$.

Stickiness is inherently different from guardedness, and its central property is as follows: variables that appear more than once in a body (i.e., join variables) are always propagated (or "stick") to the inferred atoms. A set of TGDs with the above property is sticky, and the corresponding class is denoted S . Weak stickiness is a relaxation of stickiness where only "harmful" variables are taken into account. A set of TGDs that enjoys weak stickiness is weakly sticky, and the associated class is denoted WS. Observe that $S \subset$ WS.

A set $\Sigma$ of TGDs is acyclic, if its predicate graph is acyclic, and the underlying class is denoted $\mathrm{A} . \Sigma$ is weakly acyclic, if its dependency graph enjoys a certain acyclicity condition, which actually guarantees the existence of a finite canonical model; the associated class is denoted WA. Clearly, A $\subset$ WA. Observe also that WA $\subset$ WS.

Another key fragment of TGDs which deserves our attention are the so-called full TGDs, i.e., TGDs without existentially quantified variables, and the corresponding class is denoted F. If full TGDs enjoy linearity, guardedness, stickiness, or acyclicity, then we obtain the classes LF, GF, SF, and $A F$, respectively. Observe that $F \subset W A$ and $F \subset W G$.

For a summary of the complexity of (classical) BCQ answering over these Datalog ${ }^{ \pm}$languages, see Table 1.
Complexity Classes. The complexity class $\mathrm{AC}^{0}$ is the class of all decision problems that can be solved by uniform families of Boolean circuits of polynomial size and constant depth. PSPACE (resp., P, EXP, 2EXP) is the class of all problems that can be decided in polynomial space (resp., polynomial time, exponential time, double exponential time) on a deterministic Turing machine. NP and NEXP are the classes of all problems that are decidable in polynomial and exponential time on a nondeterministic Turing machine, respectively, and co-NP and co-NEXP are their complementary classes, where 'yes' and 'no' instances are interchanged. The class $\Sigma_{2}^{P}$ is the class of all problems that can be decided in nondeterministic polynomial time using an NP oracle, and $\Pi_{2}^{\mathrm{P}}$ is the complement of $\Sigma_{2}^{\mathrm{P}}$. The class $\Theta_{2}^{\mathrm{P}}$ (resp., $\Theta_{3}^{\mathrm{P}}$ ) is the class of all problems that can be decided in polynomial time by a deterministic Turing machine with either a logarithmic number of calls to an NP (resp., $\Sigma_{2}^{p}$ ) oracle, or (equivalently) a constant number of rounds of polynomially many parallel calls to an NP (resp., $\Sigma_{2}^{\mathrm{P}}$ ) oracle. $\mathrm{P}^{\text {NEXP }}$ is the class of all problems that are decidable in deterministic polynomial time using a NEXP oracle. The above complexity classes and their inclusion relationships (which are all currently believed to be strict) are: $\mathrm{AC}^{0} \subseteq \mathrm{P} \subseteq \mathrm{NP}, \mathrm{co}-\mathrm{NP} \subseteq \Theta_{2}^{\mathrm{P}} \subseteq \Sigma_{2}^{\mathrm{P}}, \Pi_{2}^{\mathrm{P}} \subseteq$ $\Theta_{3}^{\mathrm{P}} \subseteq \mathrm{PSPACE} \subseteq \mathrm{EXP} \subseteq$ NEXP, $\mathrm{co}-\mathrm{NEXP} \subseteq \mathrm{P}^{\mathrm{NEXP}} \subseteq 2 \mathrm{EXP}$.

## Inconsistency-Tolerant Semantics

In classical $B C Q$ answering, for an inconsistent knowledge base $K B$ (i.e., $\operatorname{mods}(K B)=\emptyset$ ), every query is entailed, as everything follows from a contradiction. Clearly, the answers obtained in such cases are not meaningful.

Several inconsistency-tolerant semantics have been proposed. We now recall three prominent ones for ontologybased query answering, namely, the ABox repair $(A R)$, the intersection of repairs (IAR), and the intersection of closed repairs (ICR) semantics (Lembo et al. 2010; Bienvenu 2012); all three are based on the notion of repair, which is a maximal consistent subset of the given database. Below we define these inconsistency-tolerant semantics for a generic concept of repair maximality.

Given a knowledge base $K B=(D, \Sigma)$, a selection $D^{\prime}$ of $K B$ is a database such that $D^{\prime} \subseteq D$. A selection $D^{\prime}$ of $K B$ is consistent, if $\operatorname{mods}\left(\left(D^{\prime}, \Sigma\right)\right) \neq \emptyset$. Consistent selections of knowledge bases can be ordered according to some criteria to select the more desired ones. Given a preorder $\preccurlyeq$ over a set $\mathcal{S}$ of databases, for two elements $D^{\prime}$ and $D^{\prime \prime} \in \mathcal{S}, D^{\prime} \prec D^{\prime \prime}$
denotes that $D^{\prime} \preccurlyeq D^{\prime \prime}$ and $D^{\prime \prime} \npreceq D^{\prime}$. A database $D \in \mathcal{S}$ is $\preccurlyeq$-maximal in $\mathcal{S}$ iff there is no $D^{\prime} \in \mathcal{S}$ such that $D \prec D^{\prime}$.
Definition 1. A $\preccurlyeq$-repair of a knowledge base $K B$ is a consistent selection of $K B$ that is $\preccurlyeq$-maximal in the set of all the consistent selections of $K B$.

We now define the three different inconsistency-tolerant semantics for BCQ answering. In what follows, $R e p_{\preccurlyeq}(K B)$ denotes the set of all $\preccurlyeq$-repairs of a knowledge base $K B$. For a knowledge base $K B=(D, \Sigma)$, the closure $C n(K B)$ of $K B$ is the set of all ground atoms, built from constants in $D$ and $\Sigma$, entailed by $D$ and the TGDs of $\Sigma$.
Definition 2. Let $K B$ be a knowledge base, let $Q$ be a BCQ, and let $\preccurlyeq$ be an order over the consistent selections of $K B$.

- KB entails $Q$ under the ABox repair semantics and order $\preccurlyeq(\preccurlyeq-A R)$, denoted by $K B \neq \preccurlyeq-A R Q$, if, for all $D^{\prime} \in$ $R e p_{\preccurlyeq}(K B),\left(D^{\prime}, \Sigma\right) \models Q$.
- KB entails $Q$ under the intersection of repairs semantics and order $\preccurlyeq(\preccurlyeq-I A R)$, denoted by $K B \not \models \preccurlyeq-I A R Q$, if $\left(D^{*}, \Sigma\right) \models Q$, where $D^{*}=\bigcap\left\{D^{\prime} \mid D^{\prime} \in \operatorname{Rep} p_{\preccurlyeq}(K B)\right\}$.
- $K B$ entails $Q$ under the intersection of closed repairs semantics and order $\preccurlyeq(\preccurlyeq-I C R)$, denoted by $K B \models \preccurlyeq-I C R$ $Q$, if $\left(D_{I}, \Sigma\right) \models Q$, where $D_{I}=\bigcap\left\{C n\left(\left(D^{\prime}, \Sigma\right)\right) \mid D^{\prime} \in\right.$ $\left.R e p_{\preccurlyeq}(K B)\right\}$.
Different orders over the set of consistent selections give rise to different inconsistency-tolerant semantics for BCQ answering, because they select different repairs.

Inclusion-maximal repairs have often been considered in the literature. An interesting class of repairs are those selected by the cardinality order ' $\leq$ ' (Bienvenu, Bourgaux, and Goasdoué 2014). In this order, consistent selections of larger cardinality are preferred over ones of smaller cardinality (without looking at inclusion-wise relationships between the selections: only the cardinality counts). Hence, a $\leq$-repair of a knowledge base $K B$ is a consistent selection of $K B$ of maximum cardinality. In this paper, we consider only the ' $\leq$ ' order, therefore, we often call $\leq$-repairs simply repairs, and by $\operatorname{Rep}(K B)$, we mean $R e p_{\leq}(K B)$.
Cardinality-maximal repairs are very appropriate when it is known (or believed) that all the facts in the database have the same (possibly small) probability of being erroneous. In these cases, larger repairs are preferred, because fewer facts are dropped (Bienvenu, Bourgaux, and Goasdoué 2014). When facts in the database have different likelihoods of being erroneous, then other concepts of repairs can also be taken into consideration (Bienvenu 2012).
It can be shown that the following semantic relationships hold between the above three different semantics.
Proposition 3. Let $K B$ be a knowledge base, let $Q$ a $B C Q$. $(K B \vDash \leq-I A R Q) \Rightarrow(K B \models \leq-I C R Q) \Rightarrow(K B \models \leq-A R Q)$.

## Overview of Complexity Results

We give a precise picture of the complexity of BCQ answering from existential rules under the $\leq-A R, \leq-I A R$, and $\leq-I C R$ inconsistency-tolerant semantics. Our results are summarized in Tables 2 and 3, and they range from $\Theta_{2}^{P}-$ completeness to 2EXP-completeness.

|  | Data | $f p$-comb. | $b a-\mathbf{c o m b}$. | Comb. |
| :---: | :---: | :---: | :---: | :---: |
| $\mathrm{L}_{\perp}, \mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$ | in $\mathrm{AC}^{0}$ | NP | NP | PSPACE |
| $\mathrm{G}_{\perp}$ | P | NP | EXP | 2EXP |
| $\mathrm{WG}_{\perp}$ | EXP | EXP | EXP | 2EXP |
| $\mathrm{S}_{\perp}, \mathrm{SF}_{\perp}$ | in $\mathrm{AC}^{0}$ | NP | NP | EXP |
| $\mathrm{F}_{\perp}, \mathrm{GF}_{\perp}$ | P | NP | NP | EXP |
| $\mathrm{A}_{\perp}$ | in $\mathrm{AC}^{0}$ | NP | NEXP | NEXP |
| $W S_{\perp}, W_{\perp}$ | P | NP | 2EXP | 2EXP |

Table 1: Complexity of BCQ answering (Lukasiewicz et al. 2015). All non-"in" entries are completeness results.

|  | Data | $f p$-comb. | $b a-$ comb. | Comb. |
| :---: | :---: | :---: | :---: | :---: |
| $\mathrm{L}_{\perp}, \mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$ | $\Theta_{2}^{\text {p+* }}$ | $\Pi_{2}^{\mathrm{p}}$ | $\Theta_{3}^{\mathrm{P}}$ | PSPACE |
| $\mathrm{G}_{\perp}$ | $\Theta_{2}^{\text {P }}$ * | $\Pi_{2}^{\text {p }}$ | EXP | 2EXP |
| $\mathrm{WG}_{\perp}$ | EXP | EXP | EXP | 2EXP |
| $\mathrm{S}_{\perp}, \mathrm{SF}_{\perp}$ | $\Theta_{2}^{\text {P }}{ }^{*}$ | $\Pi_{2}^{\mathrm{p}}$ | $\Theta_{3}^{\text {P }}$ | EXP |
| $\mathrm{F}_{\perp}, \mathrm{GF}_{\perp}$ | $\Theta_{2}^{\text {p+ }}$ | $\Pi_{2}^{\text {p }}$ | $\Theta_{3}^{\text {P}}$ | EXP |
| $\mathrm{A}_{\perp}$ | $\Theta_{2}^{\text {p+ }}$ | $\Pi_{2}^{\text {p }}$ | $\mathrm{P}^{\text {NEXP }}$ | $\mathrm{P}^{\mathrm{NEXP}}$ |
| WS ${ }_{\perp},{ }^{+} A_{\perp}$ | $\Theta_{2}^{\text {P }}{ }^{+}$ | $\Pi_{2}^{p}$ | 2EXP | 2EXP |

Table 2: Complexity of $\leq-A R$ BCQ answering. All entries are completeness results. ${ }^{+}$Different proof of membership in (Bienvenu, Bourgaux, and Goasdoué 2014). *Different proof of hardness for $L_{\perp}, G_{\perp}, S_{\perp}$, and $W S_{\perp}$ in (Bienvenu, Bourgaux, and Goasdoué 2014).

Compared to (Lukasiewicz, Malizia, and Molinaro 2018a; 2018b), where subset maximality is considered for the inconsistency-tolerant semantics, we observe that using the maximum cardinality comes at a cost in several cases. This increased cost is due to the computational effort needed to evaluate the size of the largest consistent selections of the database. This extra effort does not always influence the overall complexity of the problem. Indeed, in some circumstances it is masked out by the complexity of classical BCQ answering or the complexity associated with the evaluation of the inconsistency-tolerant semantics.

In detail, $\leq-A R$-BCQ answering (see Table 2 ) is complete for $\Theta_{2}^{\mathrm{P}}$ (resp., $\Pi_{2}^{\mathrm{P}}$ ) in the data (resp., $f p$-combined) complexity for all languages of existential rules, but for $\mathrm{WG}_{\perp}$, where it is EXP-complete. In the data complexity, except for the $W G_{\perp}$ case, for which already classical BCQ answering over consistent ontologies is EXP-complete, the complexity of computing the size of the maximal-cardinality repairs dominates the complexity of the task. For the $f p$-combined complexity case, on the other hand, there is no increase in the complexity compared to the case of $A R$ semantics with subset-maximal repairs (see Lukasiewicz et al. 2015).

The $b a$-combined complexity for $\leq-A R$ - BCQ answering is among $\Theta_{3}^{\mathrm{P}}$ (for $\mathrm{L}_{\perp}, \mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}, \mathrm{S}_{\perp}, \mathrm{SF}_{\perp}, \mathrm{F}_{\perp}$, and $\mathrm{GF}_{\perp}$ ), EXP (for $G_{\perp}$ and $W G_{\perp}$ ), $\mathrm{P}^{\text {NEXP }}$ (for $\mathrm{A}_{\perp}$ ), and 2EXP (for $\mathrm{WS}_{\perp}$ and $W A_{\perp}$ ). If we compare these results with those for the subset-maximal $A R$ semantics, we observe that there is an increase in the complexity only for the languages whose classical reasoning in the $b a$-combined complexity is in NP

|  | Data | $f p$-comb. | $b a$-comb. | Comb. |
| :---: | :---: | :---: | :---: | :---: |
| $\mathrm{L}_{\perp}, \mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+*}$ | $\Theta_{2}^{\mathrm{P}}$ | $\Theta_{3}^{\mathrm{P}}$ | PSPACE |
| $\mathrm{G}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+*}$ | $\Theta_{2}^{\mathrm{P}}$ | EXP | 2 EXP |
| $\mathrm{WG}_{\perp}$ | EXP | EXP | EXP | 2 EXP |
| $\mathrm{S}_{\perp}, \mathrm{SF}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+*}$ | $\Theta_{2}^{\mathrm{P}}$ | $\Theta_{3}^{\mathrm{P}}$ | EXP |
| $\mathrm{F}_{\perp}, \mathrm{GF}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+}$ | $\Theta_{2}^{\mathrm{P}}$ | $\Theta_{3}^{\mathrm{P}}$ | EXP |
| $\mathrm{A}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+}$ | $\Theta_{2}^{\mathrm{P}}$ | $\mathrm{P}^{\text {NEXP }}$ | $\mathrm{P}^{\mathrm{NEXP}}$ |
| $\mathrm{WS}_{\perp}, \mathrm{WA}_{\perp}$ | $\Theta_{2}^{\mathrm{P}+*}$ | $\Theta_{2}^{\mathrm{P}}$ | 2 EXP | 2 EXP |

Table 3: Complexity of $\leq-I A R$ and $\leq-I C R$ BCQ answering. All entries are completeness results. ${ }^{+}$Different proof of membership for $I A R$ in (Bienvenu, Bourgaux, and Goasdoué 2014). *Different proof of hardness for $I A R$ for $\mathrm{L}_{\perp}$, $\mathrm{G}_{\perp}, \mathrm{S}_{\perp}$, and $W S_{\perp}$ in (Bienvenu, Bourgaux, and Goasdoué 2014).
(see, Lukasiewicz et al. 2015). Intuitively, in order to compute the size of the largest consistent selections, we need to perform a binary search asking an oracle to guess a selection of convenient size and check that it is consistent. This requires a $\Sigma_{2}^{\mathrm{p}}$ oracle, from which it follows that the overall procedure requires a computation in $\Theta_{3}^{\mathrm{P}}$. The combined complexity of $\leq-A R-\mathrm{BCQ}$ answering is among PSPACE (for $\mathrm{L}_{\perp}$, $\mathrm{LF}_{\perp}$, and $\mathrm{AF} \mathcal{L}_{\perp}$ ), EXP (for $\mathrm{S}_{\perp}, \mathrm{SF}_{\perp}, \mathrm{F}_{\perp}$, and $\mathrm{GF}_{\perp}$ ), $\mathrm{P}^{\text {NEXP }}$ (for $A_{\perp}$ ), and $2 \operatorname{EXP}$ (for $G_{\perp}, W S_{\perp}, W A_{\perp}$, and $W G_{\perp}$ ). In these cases, we do not observe any increase in the complexity of the tasks, compared with the complexity $A R$ semantics with subset-maximal repairs (see, Lukasiewicz et al. 2015).

The complexity of $\leq-I A R$ - and $\leq-I C R$-BCQ answering (see Table 3) slightly drops to $\Theta_{2}^{\mathrm{P}}$ in the $f p$-combined complexity for all languages, but $\mathrm{WG}_{\perp}$. As we can see, the complexity of $I A R$ - and $I C R$-BCQ answering is the same when cardinality-maximal repairs are considered, because either the complexity of classical reasoning or the complexity of computing the size of the biggest repairs dominate the complexity of the task. For this reason, since there is no extra computational cost for using $I C R$ instead of $I A R$, in this setting $I C R$ has an advantage over $I A R$ given that $I C R$ is finer approximation of $A R$ than $I A R$.

## Derivation of Complexity Results

In this section, we first derive the membership results and then the hardness results of Tables 2 and 3.

## Membership Results

The following theorem shows that if BCQ answering from knowledge bases over some Datalog ${ }^{ \pm}$language $L$ is in $\mathbf{C}$ in the data (resp., $b a$-combined and combined) complexity, then $\leq-A R$-BCQ and $\leq-I A R$-BCQ answering from knowledge bases over $L$ can be done in polynomial time with logarithmically many calls to an oracle for $\mathrm{NP}^{\mathrm{C}}$ in the data (resp., $b a$-combined and combined) complexity. This result holds for $\leq-I C R$-BCQ answering as well, but only in the data and $b a$-combined complexity. For the combined complexity case of $\leq-I C R$-BCQ answering, we will need a different proof.

Theorem 4. Let L be a Datalog ${ }^{ \pm}$language. If BCQ answering from knowledge bases over $L$ is in $\mathbf{C}$ in the data / bacombined / combined complexity (resp., data / ba-combined complexity), then $\leq-A R$ and $\leq-I A R$ (resp., $\leq-I C R$ ) BCQ answering from knowledge bases over $L$ is in P with an oracle for $\mathrm{NP} \mathbf{C}[O(\log n)]$ in the data / ba-combined / combined complexity (resp., data / ba-combined complexity).

Proof. Let $K B=(D, \Sigma)$ be a knowledge base over $L$, and let $q$ be a query. First, we compute the size max of the largest consistent selections of $K B$. This can be done by a machine in polynomial time calling, logarithmically many times, an oracle in $\mathrm{NP}^{\mathrm{C}}$. Indeed, we can perform a binary search in the range $[0,|D|]$ by asking the oracle whether there is a consistent selection of size at least $k$. The oracle has to guess a subset of size $k$ of the database, and then check whether the guess is consistent. Since we are assuming that classical reasoning in fragment $L$ is in $\mathbf{C}$, the oracle can guess a selection in NP and then check its consistency in $\mathbf{C}$.

Then, we exploit the $\mathrm{NP}^{\mathrm{C}}$ oracle to decide whether the BCQ $q$ is entailed under the considered semantics. In particular, we ask the oracle whether the query is not entailed. The way in which the oracle finds the answer depends on the specific semantics, $A R, I A R$, or $I C R$, considered.
$A R$ : The oracle guesses a consistent selection $D^{\prime}$ of size $\max$, and checks that $\left(D^{\prime}, \Sigma\right) \not \models q$.
$I A R$ : The oracle guesses a database $D^{\star} \subseteq D$ and checks that: (1) $\left(D^{\star}, \Sigma\right) \not \vDash q$, and (2) there exist repairs $D_{\alpha}^{\prime}$ with $\alpha \notin D_{\alpha}^{\prime}$, one for each $\alpha \in D \backslash D^{\star}$ (witnessing that the intersection of the repairs is a subset of $D^{\star}$ ). Point (2) is checked via additional guesses of the various $D_{\alpha}^{\prime}$.
$I C R$ : The oracle verifies the existence of a subset $D^{\prime}$ of $C n(K B)$ (the size of $C n(K B)$ is polynomial in the input, because the program has, in the worst case, bounded arity) such that: (i) for each atom $\alpha \in\left(C n(K B) \backslash D^{\prime}\right)$, there is a repair $D_{\alpha}^{\prime}$ such that $\alpha \notin C n\left(D_{\alpha}^{\prime}\right)$; (ii) $\left(D^{\prime}, \Sigma\right) \not \models q$. The oracle can do this by first guessing such a set $D^{\prime}$ along with the witnesses $D_{\alpha}^{\prime}$, and second checking that the various $D_{\alpha}^{\prime}$ are consistent, their size is max, and $\alpha \notin C n\left(D_{\alpha}^{\prime}\right)$, and that $\left(D^{\prime}, \Sigma\right) \not \vDash q$. The existence of the various repairs $D_{\alpha}^{\prime}$ with the properties above reported guarantees that $D_{I} \subseteq D^{\prime}$, and $\left(D^{\prime}, \Sigma\right) \not \models q$ implies $\left(D_{I}, \Sigma\right) \not \vDash q$ by monotonicity.

As mentioned above, to characterize the combined complexity of $\leq-I C R-B C Q$ answering, we need a different proof. Indeed, in a combined complexity scenario, the size of $C n(K B)$ is exponential in the size of the input, and the $D^{\prime} \subseteq C n(K B)$ needed to be guessed in the $I C R$ case of the proof above could be too big for an NP machine.

The next result proves all the upper bounds for $\leq-I C R$ BCQ answering in Table 3 for the combined complexity.
Theorem 5. $\leq-I C R B C Q$ answering from knowledge bases over Datalog ${ }^{\ddagger}$ languages $L$ in the combined complexity is in the complexity classes shown in Table 3.

We next analyze the $f p$-combined complexity and show that if BCQ answering from knowledge bases over some Datalog $^{ \pm}$language $L$ is in $\mathbf{D}$ (resp., $\mathbf{C}$ ) in the data (resp.,
$f p$-combined) complexity, then $\leq-A R, \leq-I A R$, and $\leq-I C R$, BCQ answering from knowledge bases over $L$ can be done in polynomial time with logarithmically many calls to an oracle for $\mathrm{NP}^{\mathrm{D}}$ followed by a computation in co-NP ${ }^{\mathrm{C}}$ (for the $\leq-A R$ ) or a computation in $\mathbf{C}$ (for the $\leq-I A R$ and $\leq-I C R$ ). This is shown in a similar way as Theorem 4.
Theorem 6. If BCQ answering from knowledge bases over a Datalog ${ }^{ \pm}$language $L$ is in $\mathbf{D}$ in the data complexity and in $\mathbf{C}$ in the fp-combined complexity, then $\leq-A R$ (resp., $\leq-I A R$ and $\leq-I C R$ ) BCQ answering from knowledge bases over $L$ is possible by a computation in P with an oracle for $\mathrm{NP}^{\mathbf{D}}[O(\log n)]$, followed by a computation in $\mathrm{co}-\mathrm{NP}^{\mathbf{C}}$ (resp., $\mathbf{C}$ ), in the fp-combined complexity.

As a corollary of the theorems above, we obtain all the upper bounds for $\leq-A R$ BCQ answering in Table 2, and all the upper bounds for $\leq-I A R$ and $\leq-I C R$ BCQ answering in Table 3. In particular, the $\Theta_{2}^{\mathrm{P}}$ membership of $\leq-A R$ and $\leq-I A R \mathrm{BCQ}$ answering in the data complexity for the languages whose BCQ reasoning is feasible in polynomial time in the data complexity was already known (Bienvenu, Bourgaux, and Goasdoué 2014). The other results are new.
Corollary 7. $\leq-A R$ (resp., $\leq-I A R$ and $\leq-I C R$ ) BCQ answering from knowledge bases over Datalog ${ }^{ \pm}$languages $L$ in Table 2 (resp., Table 3) in the data, fp-combined, bacombined, and combined complexity belongs to the complexity classes shown in Table 2 (resp., Table 3).

## Hardness Results

The hardness results not explicitly proven follows from the hardness of classical reasoning over consistent ontologies.

We now show that $\leq-A R, \leq-I A R$, and $\leq-I C R$ BCQ answering are $\Theta_{2}^{\mathrm{P}}$-hard in the data complexity for all the Data$\log ^{ \pm}$languages considered.
Theorem 8. For every $C \in\{A R, I A R, I C R\}, \leq-C B C Q$ answering from knowledge bases over $\mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$, and $\mathrm{SF}_{\perp}$ (and hence also for $\mathrm{L}_{\perp}, \mathrm{A}_{\perp}, \mathrm{S}_{\perp}, \mathrm{G}_{\perp}, \mathrm{GF}_{\perp}, \mathrm{F}_{\perp}, \mathrm{WS}_{\perp}$, and $\mathrm{WA}_{\perp}$ ) is $\Theta_{2}^{\mathrm{P}}$-hard in the data complexity.

Proof. We use a reduction from the $\Theta_{2}^{\mathrm{P}}$-complete problem InALLMAxIS (Lopatenko and Bertossi 2007; 2016): given a graph $G$ and a vertex $w$, decide whether $w$ belongs to all the independent sets of $G$ of maximum size.

Let $(G, w)$ be an instance of InAllMAxIS, where $G=$ $(V, E)$ and $w \in V$, and $n=|V|$. From $(G, w)$ we build the knowledge base $K B_{\text {MaxIS }}=\left(D_{M a x I S}, \Sigma_{\text {MaxIS }}\right)$ and the query $q_{\text {MaxIS }}$ as follows.
For each vertex $v \in V$, there is a fact $\operatorname{In}(v)$ in $D_{M a x I S}$. For each edge $\left(v_{1}, v_{2}\right) \in E$, there are $n$ facts $\operatorname{Edge}\left(v_{1}, v_{2}, i\right)$ in $D_{\text {MaxIS }}$, with $1 \leq i \leq n$. $D_{M a x I S}$ contains also the fact distinguished $(w)$. The only dependency in $\Sigma_{\text {MaxIS }}$ is the NC $\operatorname{In}(X) \wedge \operatorname{In}(Y) \wedge \operatorname{Edge}(X, Y, Z) \rightarrow \perp$. Finally, $q_{\text {MaxIS }}=\exists X(\operatorname{In}(X) \wedge$ distinguished $(X))$. Observe that $\Sigma_{\text {MaxIS }}$ is vacuously linear, sticky, acyclic, and full.

We can show that $(G, w)$ is a 'yes'-instance of InAlLMAXIS iff $K B_{\text {MaxIS }}$ entails $q_{M a x I S}$ under any $\leq-C$, for $C \in\{A R, I A R, I C R\}$. Intuitively, repairs contain all facts $\operatorname{Edge}(v 1, v 2, i)$, distinguished $(w)$ and a set of facts $\operatorname{In}(v)$ corresponding to a maximum independent set of $G$.

For $\leq-A R$ BCQ answering, we show its $\Pi_{2}^{\mathrm{P}}$-hard in the fp-combined complexity, for all the Datalog ${ }^{ \pm}$languages considered. This follows from a reduction in (Lukasiewicz et al. 2015), which also applies in the case of maximum cardinality. Note that also a simplification of the reduction used here to prove Theorem 10 can be used to show this result.
Theorem 9. $\leq-A R B C Q$ answering from knowledge bases over $\mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$, and $\mathrm{SF}_{\perp}$ (and hence also for $\mathrm{L}_{\perp}, \mathrm{A}_{\perp}$, $\mathrm{S}_{\perp}, \mathrm{G}_{\perp}, \mathrm{GF}_{\perp}, \mathrm{F}_{\perp}, \mathrm{WS}_{\perp}$, and $\mathrm{WA}_{\perp}$ ) is $\Pi_{2}^{\mathrm{P}}$-hard in the $f p-$ combined complexity.

We now show that, for any $C \in\{A R, I A R, I C R\}, \leq-C$ BCQ answering is $\Theta_{3}^{\mathrm{P}}$-hard in the ba-combined complexity, already for the simplest languages here considered.
Theorem 10. For every $C \in\{A R, I A R, I C R\}, \leq-C B C Q$ answering from knowledge bases over $\mathrm{LF}_{\perp}, \mathrm{AF}_{\perp}$, and $\mathrm{SF}_{\perp}$ (and hence also for $\mathrm{L}_{\perp}, \mathrm{S}_{\perp}, \mathrm{GF}_{\perp}$, and $\mathrm{F}_{\perp}$ ) is $\Theta_{3}^{\mathrm{P}}$-hard in the ba-combined complexity.

Proof sketch. We prove the statement via a reduction from the $\Theta_{3}^{\mathrm{P}}$-complete problem Comp-VALID ${ }_{2}$, which is a generalization of the problem Comp-SAT (Lukasiewicz and Malizia 2016): given two sets $A$ and $B$ of quantified Boolean formulas characterized by 2 alternating quantifiers, decide whether the number of valid formulas in $A$ is greater than the number of valid formulas in $B$. The $\Theta_{3}^{\mathrm{P}}$-hardness of CompVALID $_{2}$ holds even if the following restrictions are imposed over the instances (Lukasiewicz and Malizia 2017):

- $|A|=|B|$;
- all formulas of $A$ and $B$ are of the kind $(\forall X)(\exists Y) \phi(X$, $Y)$, where $\phi(X, Y)$ is a non-quantified 3CNF formula;
- all formulas of $A$ and $B$ have the same number of clauses in the non-quantified part;
- all formulas of $A$ and $B$ have the same sets of universally quantified variables and existentially quantified variables;
- sets $A=\left\{\Phi_{1}, \ldots, \Phi_{v}\right\}$ and $B=\left\{\Psi_{1}, \ldots, \Psi_{v}\right\}$ are such that $\Phi_{u+1}$ (resp., $\Psi_{u+1}$ ) being valid implies $\Phi_{u}$ (resp., $\Psi_{u}$ ) being valid as well, for any $u$ (intuitively, all valid formulas have the lowest indices in sets $A$ and $B$ ).
From the last assumption it follows that $(A, B)$ is a 'yes'instance of Comp-VALID ${ }_{2}$ iff there is an index $u$ such that $\Phi_{u} \in A$ is valid and $\Psi_{u} \in B$ is not valid.

Given the restrictions listed above, we can also assume, without simplifying the computational complexity of the problem, that all formulas of $A$ and $B$ are $N Q B F_{2, \forall}$ ones (Greco et al. 2009; 2011; Schaefer 2001). $N Q B F_{2, \forall}$ formulas are quantified Boolean formulas of the kind $\Phi=$ $(\forall X)(\exists Y) \phi(X, Y)$, in which $X=\left\{x_{1}, \ldots, x_{n}\right\}$ and $Y=$ $\left\{y_{1}, \ldots, y_{r}\right\}$ are two disjoint sets of Boolean variables, and $\phi(X, Y)=c_{i(1)} \wedge c_{\bar{i}(1)} \wedge \cdots \wedge c_{i(n)} \wedge c_{\bar{i}(n)} \wedge c_{1} \wedge \cdots \wedge c_{m}$ is a 3 CNF formula, where each universally quantified variable $x_{k} \in X$ occurs only in the two clauses $c_{i(k)}=$ $\left(x_{k} \vee \neg y_{k}\right)$ and $c_{\bar{i}(k)}=\left(\neg x_{k} \vee y_{k}\right)$-intuitively, each variable $x_{k}$ enforces the truth value of a corresponding variable $y_{k}$, which thus plays its role in the formula $\phi(X, Y)$. Indeed, a formula $(\forall X)(\exists Y) c_{1} \wedge \cdots \wedge c_{m}$ can be replaced by $(\forall X)\left(\exists Y_{X}\right)(\exists Y) c_{i(1)} \wedge c_{\bar{i}(1)} \wedge \cdots \wedge c_{i(n)} \wedge c_{\bar{i}(n)} \wedge c_{1}^{\prime} \wedge \cdots \wedge c_{m}^{\prime}$,
where $X=\left\{x_{1}, \ldots, x_{n}\right\}, Y_{X}=\left\{y_{1}^{x}, \ldots, y_{n}^{x}\right\}$ and each $c_{i}^{\prime}$ is obtained by replacing every $x_{j} \in X$ by $y_{j}^{x} \in Y_{X}$ in $c_{i}$. We will refer to the clauses $c_{1}, \ldots, c_{m}$ of an $N Q B F_{2, \forall}$ formula as the main clauses of the formula.

Let $\mathcal{I}=(A, B)$ be an instance of Comp-VALID ${ }_{2}$ satisfying all the restrictions above. We denote by $X$ and $Y$ the sets of the universally and existentially quantified variables, respectively. From $\mathcal{I}$, we build the knowledge base $K B_{C V}(\mathcal{I})=\left(D_{C V}, \Sigma_{C V}\right)$ and query $q_{C V}(\mathcal{I})$ as follows. In what follows, the first two arguments of each fact of $D_{C V}$ specify the identifier of the set (i.e., whether the formula belongs to set $A$ or $B$ ) and the numeric identifier of the formula in the set, respectively. In this way, we can discriminate facts referring to the different formulas in the two sets.

For each variable $x_{i} \in X$, in $D_{C V}$ there are facts:

$$
\operatorname{Val}\left(s, u, x_{i}, t, k\right) \quad \operatorname{Val}\left(s, u, x_{i}, f, k\right)
$$

where $s \in\{a, b\}$ is a constant representing sets $A$ and $B$, respectively, $u \in\{1, \ldots, v\}$ is a numeric constant representing the index of the formula in the sets, $x_{i}$ is a constant representing the respective variable, $t$ and $f$ are constants representing Boolean values true and false, respectively, and $k \in\{1,2,3\}$ is a numeric constant (allowing us to have three copies of the facts).

There are facts $D_{C V}$ that will be used to impose the consistency of the assignments to the literals:

$$
\begin{array}{ll}
\operatorname{SimLit}(s, u, t, t, k) & \operatorname{OppLit}(s, u, t, f, k) \\
\operatorname{SimLit}(s, u, f, f, k) & \operatorname{OppLit}(s, u, f, t, k),
\end{array}
$$

where $s \in\{a, b\}, u \in\{1, \ldots, v\}, t, f$, and $k \in\{1,2,3\}$, are constants with the same meaning as above.

There are facts in $D_{C V}$ that will be used to select possible ways of satisfying clauses in the formulas:

$$
\begin{array}{ll}
\operatorname{ClSat}(s, u, f, t, t, k) & \operatorname{ClSat}(s, u, t, t, t, k), \\
\operatorname{ClSat}(s, u, f, t, f, k) & \operatorname{ClSat}(s, u, t, t, f, k) \\
\operatorname{ClSat}(s, u, f, f, t, k) & \operatorname{ClSat}(s, u, t, f, t, k) \\
& \operatorname{ClSat}(s, u, t, f, f, k)
\end{array}
$$

where $s \in\{a, b\}, u \in\{1, \ldots, v\}, t, f$, and $k \in\{1,2,3\}$, are constants with the same meaning as above.

SimLit, OppLit, and ClSat, are the structural facts, and we denote by $D_{C V}^{S t}$ the sets of structural facts of $D_{C V}$.

Finally, $D_{C V}$ also contains facts that will be used to signal the validity or the non-validity of the formulas:

$$
\operatorname{Non} \operatorname{Valid}(s, u, k) \quad \operatorname{Valid}(s, u),
$$

where $s \in\{a, b\}, u \in\{1, \ldots, v\}$, and $k \in\{1,2\}$, are constants with the same meaning as above. Note that we have two copies for facts Non Valid, and one copy for facts Valid.

In the program $\Sigma_{C V}$, there is no TGD, and there are the following NCs. For notational convenience, in the NCs below, we will use the underscore '_' as a placeholder for a fresh variable not appearing anywhere else in the NC.

The first negative constraint is

$$
\operatorname{Val}\left(S, U, X, t,{ }_{-}\right), \operatorname{Val}\left(S, U, X, f,{ }_{-}\right) \rightarrow \perp .
$$

Next, for each formula in sets $A$ and $B$, there is an NC. This NC is used to check the satisfiability of a formula once an assignment for the variables in $X$ has been provided.

For simplicity in the exposition, let us assume that we are considering the formula $\Phi_{1}=(\forall X)(\exists Y) \phi_{1}(X, Y)$ belonging to set $A$. The following construction for the NC can be generalized to any formula of $A$ or $B$. Since this NC is intricate, we look at its various constituent pieces.

A first piece of the NC checks that at least one of the copies of the structural facts is in the selection:

$$
\operatorname{Config}(a, 1) \equiv \bigwedge_{p(a, 1, \mathbf{c},-) \in D_{C V}^{S t}} p(a, 1, \mathbf{c},-)
$$

By using a fresh variable as the last argument of all the predicates in Config, the presence of just one copy of each structural fact is enough to support the activation of Config.

A second piece of the NC "reads" the assignment on the variables in $X$ encoded in the selection of the database:

$$
\operatorname{Assign} X(a, 1) \equiv \bigwedge_{i=1}^{n} \operatorname{Val}\left(a, 1, x_{i}, T_{i},,_{-}\right)
$$

Below we use the following notation: $l_{j, k}$ is the $k^{\text {th }}$ literal in the $j^{\text {th }}$ main clause, $c_{j}$, and $v_{j, k}$ is the variable of $l_{j, k}$.

A third piece of the NC aims at "copying" the assignment on the variables in $X$ onto the associated variables in $Y$ :

$$
\operatorname{Copy}(a, 1) \equiv \bigwedge_{i=1}^{n} \operatorname{SimLit}\left(a, 1, T_{i}, T_{j, k},{ }_{-}\right)
$$

where $T_{j, k}$ is a variable for the Boolean value of the literal $l_{j, k}$ for which $v_{j, k}=y_{i}$ in $\phi_{1}(X, Y)$. Observe that, in order for Copy to work properly, each variable $y_{i}$ must appear as a positive literal in one of the main clauses of $\phi_{1}(X, Y)$ at least once. This can be assumed without loss of generality, because if $y_{i}$ always appears as a negative literal in all the main clauses of $\phi_{1}(X, Y)$, then we can replace all the occurrences of the negative literal $\neg y_{i}$ with the positive literal $y_{i}$ without altering the satisfiability properties of $\phi_{1}(X, Y)$.

A fourth piece of the NC forces that the facts selected to simulate the assignment on the variables in $Y$ are consistent. In the notation below, $\ell_{j, k} \sim \ell_{j^{\prime}, k^{\prime}}$ means that literals $\ell_{j, k}$ and $\ell_{j^{\prime}, k^{\prime}}$ are both positive or negative, while $\ell_{j, k} \nsim \ell_{j^{\prime}, k^{\prime}}$ means that one literal is positive and the other is negative.

$$
\operatorname{ConsistY}(a, 1) \equiv \bigwedge_{\begin{array}{c}
\forall\left(\ell_{j, k}, \ell_{\left.j^{\prime}, k^{\prime}\right)}\right) \\
\text { s.t. } v_{j, k}=v_{j^{\prime}, k^{\prime}} \wedge \\
\ell_{j, k} \sim \ell_{j^{\prime}, k^{\prime}}
\end{array}} \operatorname{SimLit}\left(a, 1, T_{j, k}, T_{j^{\prime}, k^{\prime}},,_{-}\right)
$$

where $T_{j, k}$ is a variable with the same meaning as above.
The last piece of the NC checks $\phi_{1}(X, Y)$ 's satisfiability:

$$
\operatorname{Satisfied}(a, 1) \equiv \bigwedge_{j=1}^{m} C l S a t\left(a, 1, T_{j, 1}, T_{j, 2}, T_{j, 3},-\right)
$$

these predicates are only for the main clauses of $\phi_{1}(X, Y)$.

To conclude, the NC associated with $\Phi_{1} \in A$ is:
$\operatorname{Config}(a, 1), \operatorname{AssignX}(a, 1), \operatorname{Copy}(a, 1)$,
$\operatorname{Consist} Y(a, 1), \operatorname{Satisfied}(a, 1), \operatorname{NonValid}(a, 1,-) \rightarrow \perp$.
The last NC is: $\operatorname{Non} \operatorname{Valid}\left(S, U,{ }_{-}\right), \operatorname{Valid}(S, U) \rightarrow \perp$.
The query is:

$$
q_{C V}(\mathcal{I})=(\exists U)(\operatorname{Valid}(a, U), \operatorname{Non} \operatorname{Valid}(b, U,-))
$$

$K B_{C V}(\mathcal{I})$ has no TGDs, and bounded arity predicates.
It can be shown that $\mathcal{I}$ is a 'yes'-instance of CompVALID $_{2}$ iff $K B_{C V}(I)$ entails $q_{C V}(I)$ under $\leq-C$ semantics, for any $C \in\{A R, I A R, I C R\}$.

The following result shows that $\leq-A R-\mathrm{BCQ} \leq-I A R$ BCQ , and $\leq-I C R-\mathrm{BCQ}$ answering for $\mathrm{A}_{\perp}$ are $\mathrm{P}^{\mathrm{NEXP}}$-hard in the $b a$-combined complexity, proving all $\mathrm{P}^{\text {NEXP }}$-hardness results in Tables 2 and 3, including those for the more general combined complexity.
Theorem 11. For any $C \in\{A R, I A R, I C R\}, \leq-C B C Q$ answering for $\mathrm{A}_{\perp}$ are $\mathrm{P}^{\mathrm{NEXP}}$-hard in the ba-combined and combined complexity.

Proof sketch. Intuitively, the reduction for the $\mathrm{P}^{\mathrm{NEXP}}$-hardness proof in (Eiter, Lukasiewicz, and Predoiu 2016) for $\subseteq-A R-\mathrm{BCQ}$ answering for $\mathrm{A}_{\perp}$ in the $b a$-combined complexity is already a reduction for $\leq-A R-\mathrm{BCQ}$ answering, as maximal-cardinality consistent database subsets there coincide with maximal-subset consistent database subsets. Furthermore, since the reduction uses a ground atomic query, this also shows the $\mathrm{P}^{\text {NEXP }}$-hardness of $\leq-I C R$ - BCQ answering for $\mathrm{A}_{\perp}$, as $\leq-I C R-\mathrm{BCQ}$ answering coincides with $\leq-A R-\mathrm{BCQ}$ answering for ground queries. Finally, the reduction for the $\mathrm{P}^{\mathrm{NEXP}}$-hardness proof in (Eiter, Lukasiewicz, and Predoiu 2016) for $\subseteq-A R-\mathrm{BCQ}$ answering for $\mathrm{A}_{\perp}$ in the $b a$-combined complexity is turned into a $\mathrm{P}^{\mathrm{NEXP}}$-hardness proof for $\leq-I A R$-BCQ answering in this case. There, one encodes initial tiling assignments $v_{1}\left(X_{i}\right), \ldots, v_{n}\left(X_{n}\right)$ and has a ground atomic query $q$, which we now also include in the database along with a fresh ground atom $n q$ and the NC $v_{1}\left(X_{i}\right) \wedge \ldots \wedge v_{n}\left(X_{n}\right) \wedge q \wedge n q \rightarrow \perp$. This intuitively "forces" the atom $q$ into the database.

The $\Theta_{2}^{\mathrm{P}}$-hardness in the $f p$-combined complexity of $\leq-I A R$ - BCQ and $\leq-I C R-\mathrm{BCQ}$ answering over all the languages considered follows from the $\Theta_{2}^{\mathrm{P}}$-hardness in the data complexity of $\leq-I A R$-BCQ answering.
Theorem 12. $\leq-I A R$ and $\leq-I C R B C Q$ answering is $\Theta_{2}^{\mathrm{P}}-$ hard in the fp-combined complexity for $\mathrm{L}_{\perp}, \mathrm{LF}_{\perp}, \mathrm{A}_{\perp}, \mathrm{AF}_{\perp}$, $\mathrm{S}_{\perp}, \mathrm{SF}_{\perp}, \mathrm{G}_{\perp}, \mathrm{GF}_{\perp}, \mathrm{F}_{\perp}, \mathrm{WS}_{\perp}$, and $\mathrm{WA}_{\perp}$ knowledge bases.

## Summary and Outlook

We have given a precise picture of the complexity of BCQ answering under different cardinality-maximal incon-sistency-tolerant semantics, namely, the ABox repair, the intersection of repairs (IAR), and the intersection of closed repairs (ICR) semantics, for the most popular Datalog ${ }^{ \pm}$languages and complexity measures. Note that these complexity results can now also be used to derive further complexity
results for other Datalog ${ }^{ \pm}$languages. For example, for shy Datalog ${ }^{ \pm}$(Leone et al. 2012), by the results of this paper, it is immediate that BCQ answering is complete for $\Theta_{2}^{\mathrm{p}}$ and EXP in the data and combined complexity, respectively.

Future research lines include considering other classes of existential rules and defining other semantics for incon-sistency-tolerant ontological query answering. In particular, it would be interesting to explore whether there are datatractable and/or even first-order rewritable other such semantics. Furthermore, a more fine-grained way to analyze the complexity of query answering would be a non-uniform approach, looking at the complexity of a single ontology or a single ontology-mediated query (see, e.g., (Bienvenu et al. 2014; Koutris and Suciu 2014; Hernich et al. 2017)).
Acknowledgments. This work was supported by the Alan Turing Institute under the UK EPSRC grant EP/N510129/1, and by the EPSRC grants EP/R013667/1, EP/L012138/1, and EP/M025268/1.

## References

Arenas, M.; Bertossi, L. E.; and Chomicki, J. 1999. Consistent query answers in inconsistent databases. In Proc. PODS, 68-79.
Bienvenu, M., and Bourgaux, C. 2016. Inconsistencytolerant querying of description logic knowledge bases. In Reasoning Web, 156-202.
Bienvenu, M., and Rosati, R. 2013. Tractable approximations of consistent query answering for robust ontologybased data access. In Proc. IJCAI, 775-781.
Bienvenu, M.; ten Cate, B.; Lutz, C.; and Wolter, F. 2014. Ontology-based data access: A study through disjunctive Datalog, CSP, and MMSNP. ACM TODS 39(4):33:1-33:44.
Bienvenu, M.; Bourgaux, C.; and Goasdoué, F. 2014. Querying inconsistent description logic knowledge bases under preferred repair semantics. In Proc. AAAI, 996-1002.
Bienvenu, M.; Bourgaux, C.; and Goasdoué, F. 2016. Explaining inconsistency-tolerant query answering over description logic knowledge bases. In Proc. AAAI, 900-906.
Bienvenu, M. 2012. On the complexity of consistent query answering in the presence of simple ontologies. In Proc. AAAI, 705-711.
Calì, A.; Gottlob, G.; and Kifer, M. 2013. Taming the infinite chase: Query answering under expressive relational constraints. J. Artif. Intell. Res. 48:115-174.
Calì, A.; Gottlob, G.; and Lukasiewicz, T. 2012. A general Datalog-based framework for tractable query answering over ontologies. J. Web Sem. 14:57-83.
Calì, A.; Gottlob, G.; and Pieris, A. 2012. Towards more expressive ontology languages: The query answering problem. Artif. Intell. 193:87-128.
Eiter, T.; Lukasiewicz, T.; and Predoiu, L. 2016. Generalized consistent query answering under existential rules. In Proc. KR, 359-368.
Fagin, R.; Kolaitis, P. G.; Miller, R. J.; and Popa, L. 2005. Data exchange: semantics and query answering. Theor. Comput. Sci. 336(1):89-124.

Greco, G.; Malizia, E.; Palopoli, L.; and Scarcello, F. 2009. On the complexity of compact coalitional games. In Proc. IJCAI, 147-152.
Greco, G.; Malizia, E.; Palopoli, L.; and Scarcello, F. 2011. On the complexity of core, kernel, and bargaining set. Artif. Intell. 175(12-13):1877-1910.
Hernich, A.; Lutz, C.; Papacchini, F.; and Wolter, F. 2017. Dichotomies in ontology-mediated querying with the guarded fragment. In Proc. PODS, 185-199.
Koutris, P., and Suciu, D. 2014. A dichotomy on the complexity of consistent query answering for atoms with simple keys. In Proc. ICDT, 165-176.
Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2010. Inconsistency-tolerant semantics for description logics. In Proc. RR, 103-117.
Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2015. Inconsistency-tolerant query answering in ontology-based data access. J. Web Sem. 33:3-29.
Leone, N.; Manna, M.; Terracina, G.; and Veltri, P. 2012. Efficiently computable Datalog ${ }^{\exists}$ programs. In Proc. KR.
Lopatenko, A., and Bertossi, L. E. 2007. Complexity of consistent query answering in databases under cardinality-based and incremental repair semantics. In Proc. ICDT, 179-193.
Lopatenko, A., and Bertossi, L. 2016. Complexity of consistent query answering in databases under cardinality-based and incremental repair semantics (extended version). Technical Report 1605.07159, arXiv.
Lukasiewicz, T., and Malizia, E. 2016. On the complexity of $m$ CP-nets. In Proc. AAAI, 558-564.
Lukasiewicz, T., and Malizia, E. 2017. A novel characterization of the complexity class $\Theta_{k}^{P}$ based on counting and comparison. Theor. Comput. Sci. 694:21-33.
Lukasiewicz, T.; Martinez, M. V.; Pieris, A.; and Simari, G. I. 2015. From classical to consistent query answering under existential rules. In Proc. AAAI, 1546-1552.
Lukasiewicz, T.; Malizia, E.; and Molinaro, C. 2018a. Complexity of approximate query answering under inconsistency in Datalog ${ }^{ \pm}$. In Proc. IJCAI, 1921-1927.
Lukasiewicz, T.; Malizia, E.; and Molinaro, C. 2018b. Complexity of approximate query answering under inconsistency in Datalog ${ }^{ \pm}$. In Proc. SEBD, 22:1-8.
Lukasiewicz, T.; Martinez, M. V.; and Simari, G. I. 2012. Inconsistency handling in Datalog+/- ontologies. In Proc. ECAI, 558-563.
Lukasiewicz, T.; Martinez, M. V.; and Simari, G. I. 2013. Complexity of inconsistency-tolerant query answering in Datalog+/-. In Proc. OTM, 488-500.
Rosati, R. 2011. On the complexity of dealing with inconsistency in description logic ontologies. In Proc. IJCAI.
Schaefer, M. 2001. Graph Ramsey theory and the polynomial hierarchy. J. Comput. Syst. Sci. 62(2):290-322.
Vardi, M. Y. 1982. The complexity of relational query languages (extended abstract). In Proc. STOC, 137-146.


[^0]:    Copyright (C) 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

