

The Complexity of Counting Problems Over Incomplete Databases

MARCELO ARENAS and PABLO BARCELÓ, Universidad Católica & IMFD Chile
 MIKAËL MONET, Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRISTAL, F-59000 Lille, France

We study the complexity of various fundamental counting problems that arise in the context of incomplete databases, i.e., relational databases that can contain unknown values in the form of labeled nulls. Specifically, we assume that the domains of these unknown values are finite and, for a Boolean query q , we consider the following two problems: Given as input an incomplete database D , (a) return the number of completions of D that satisfy q ; or (b) return the number of valuations of the nulls of D yielding a completion that satisfies q . We obtain dichotomies between #P-hardness and polynomial-time computability for these problems when q is a self-join-free conjunctive query and study the impact on the complexity of the following two restrictions: (1) every null occurs at most once in D (what is called *Codd tables*); and (2) the domain of each null is the same. Roughly speaking, we show that counting completions is much harder than counting valuations: For instance, while the latter is always in #P, we prove that the former is not in #P under some widely believed theoretical complexity assumption. Moreover, we find that both (1) and (2) can reduce the complexity of our problems. We also study the approximability of these problems and show that, while counting valuations always has a fully polynomial-time randomized approximation scheme (FPRAS), in most cases counting completions does not. Finally, we consider more expressive query languages and situate our problems with respect to known complexity classes.

CCS Concepts: • **Theory of computation** → **Complexity theory and logic**; **Incomplete, inconsistent, and uncertain databases**; • **Mathematics of computing** → *Approximation algorithms*;

Additional Key Words and Phrases: Incomplete databases, closed-world assumption, counting complexity, Fully Polynomial-time Randomized Approximation Scheme (FPRAS)

ACM Reference format:

Marcelo Arenas, Pablo Barceló, and Mikaël Monet. 2021. The Complexity of Counting Problems Over Incomplete Databases. *ACM Trans. Comput. Logic* 22, 4, Article 21 (September 2021), 52 pages.
<https://doi.org/10.1145/3461642>

Marcelo Arenas also with Department of Computer Science and Institute for Mathematical and Computational Engineering, School of Engineering, Pontificia Universidad Católica de Chile.

Pablo Barceló also with Institute for Mathematical and Computational Engineering, School of Engineering and Faculty of Mathematics, Pontificia Universidad Católica de Chile.

This work was partially funded by ANID - Millennium Science Initiative Program - Code ICN17_002. Arenas is funded by Fondecyt grant 1191337 and Barceló by Fondecyt grant 1200967.

Authors' addresses: M. Arenas and P. Barceló, Vicuna Mackenna 4860, Hernán Briones building, Macul 7820436, Santiago, Chile; emails: marenas@ing.puc.cl, pbarcelo@uc.cl; M. Monet, Parc Scientifique de la Haute Borne, Park Plaza - INRIA Bâtiment B - avenue Halley 59650 Villeneuve d'Ascq, France; email: mikael.monet@inria.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1529-3785/2021/09-ART21 \$15.00

<https://doi.org/10.1145/3461642>

1 INTRODUCTION

Context. In the database literature, *incomplete databases* are often used to represent missing information in the data; see, e.g., References [1, 37, 52]. These are traditional relational databases whose active domain can contain both constants and *nulls*, the latter representing unknown values [30]. There are many ways in which one can define the semantics of such a database, each being equally meaningful depending on the intended application. Under the so-called *closed-world assumption* [1, 45], a standard, complete database $v(D)$ is obtained from an incomplete database D by applying a *valuation* v that replaces each null \perp in D with a constant $v(\perp)$. The goal is then to reason about the space formed by all valuations v and completions $v(D)$ of D .

Decision problems related to querying incomplete databases have been well studied already. Consider for instance the problem $\text{Certainty}(q(\bar{x}))$ that, for a fixed query $q(\bar{x})$, takes as input an incomplete database D and a tuple \bar{a} and asks whether \bar{a} is an answer to q for every possible completion of D . By now, we have a deep understanding of the complexity of these kind of decision problems for different choices of query languages, including **conjunctive queries (CQs)** and FO queries [2, 30]. However, having the answer to this question is sometimes of little help: What if it is not the case that q is certain on D ? Can we still infer some useful information? This calls for new notions that could be used to measure the certainty with which q holds, notions that should be finer than those previously considered. This is for instance what the recent work in Reference [38] does by introducing a notion of *best answer*, which are those tuples \bar{a} for which the set of completions of D over which $q(\bar{a})$ holds is maximal with respect to set inclusion.

A fundamental complementary approach to address this issue can be obtained by considering some *counting problems* related to incomplete databases; more specifically, determining the number of completions/valuations of an incomplete database that satisfy a query q . These problems are relevant, as they tell us, intuitively, how close is q from being certain over D , i.e., what is the level of *support* that q has over the set of completions/valuations of D . Surprisingly, such counting problems do not seem to have been studied for incomplete databases. A reason for this omission in the literature might be that, in general, it is assumed that the domain over which nulls can be interpreted is infinite, and thus incomplete databases might have an infinite number of completions/valuations. However, in many scenarios it is natural to assume that the domain over which nulls are interpreted is finite, in particular when dealing with uncertainty in practice [4, 6, 7, 11, 23, 46]. By assuming this, we can ensure that the number of completions and valuations are always finite, and thus that they can be counted. This is the setting that we study.

Problems studied. We focus on the problems $\#\text{Comp}(q)$ and $\#\text{Val}(q)$ for a Boolean query q , which take as input an incomplete database D together with a finite set $\text{dom}(\perp)$ of constants for every null \perp occurring in D , and ask the following: How many completions, respectively, valuations, of D satisfy q ? More formally, a *valuation* of D is a mapping v that associates to every null \perp of D a constant $v(\perp)$ in $\text{dom}(\perp)$. Then, given a valuation v of D , we denote by $v(D)$ the database that is obtained from D after replacing each null \perp with $v(\perp)$, and we call such a database a *completion*. Besides, in this article, we consider set semantics, that is, we remove repeated tuples from $v(D)$. For $\#\text{Comp}(q)$, we count all databases of the form $v(D)$ such that q holds in $v(D)$. Instead, for $\#\text{Val}(q)$, we count the number of valuations v such that q holds in $v(D)$. It is easy to see that these two values can differ, as a completion might be obtained from two different valuations, i.e., there might exist two distinct valuations v, v' such that $v(D) = v'(D)$. We think that both problems are meaningful: While $\#\text{Comp}(q)$ determines the support for q over the databases represented by D , we have that $\#\text{Val}(q)$ further refines this by incorporating the support for a particular completion that satisfies q over the set of valuations for D .

Table 1. Our Dichotomies for Counting Valuations and Completions of sjfBCQs

	Counting valuations		Counting completions	
	Non-uniform	Uniform	Non-uniform	Uniform
Naïve	$R(x, x)$ $R(x) \wedge S(x)$	$R(x, x)$ $R(x) \wedge S(x, y) \wedge T(y)$ $R(x, y) \wedge S(x, y)$	$R(x)$	$R(x, x)$ $R(x, y)$
Codd	$R(x) \wedge S(x)$	$R(x) \wedge S(x, y) \wedge T(y)$ $R(x, y) \wedge S(x, y)$	$R(x)$	$R(x, x)$ $R(x, y)$

For each of the eight cases, if an sjfBCQ q contains a pattern mentioned in that case, then the problem is #P-hard (and #P-complete for counting valuations, as well as for counting completions over Codd tables). In turn, for each case, if an sjfBCQ q does not have any of the patterns mentioned in that case, then the problem is in FP.

We deal with the *data complexity* of the problems #Comp(q) and #Val(q), focusing on obtaining dichotomy results for them in terms of counting complexity classes, as well as studying the existence of randomized algorithms that approximate their results under probabilistic guarantees. For the dichotomies, we concentrate on **self-join-free Boolean conjunctive queries** (sjfBCQs). This assumption simplifies the mathematical analysis, while at the same time defines a setting that is rich enough for many of the theoretical concepts behind these problems to appear in full force. Notice that a similar assumption is used in several works that study counting problems over probabilistic and inconsistent databases; see, e.g., References [18, 40]. To simplify further the presentation, in the bulk of the article, we mainly consider self-join-free Boolean conjunctive queries that do not contain constants; however, we explain later (in Section 7) how our results can be extended to queries that can contain constants and free variables.

To refine our analysis, we study two restrictions of the problems #Comp(q) and #Val(q) based on natural modifications of the semantics and analyze to what extent these restrictions simplify our problems. For the first restriction, we consider incomplete databases in which each null occurs exactly once, which corresponds to the well-studied setting of *Codd tables*—as opposed to *naive tables* where nulls are allowed to have multiple occurrences. We denote the corresponding problems by #Val_{Cd}(q) and #Comp_{Cd}(q). For the second restriction, we consider *uniform* incomplete databases in which all the nulls share the same domain—as opposed to the basic *non-uniform* setting in which all nulls come equipped with their own domain. We denote the corresponding problems by #Val^u(q) and #Comp^u(q). When both restrictions are in place, we denote the problems by #Val_{Cd}^u(q) and #Comp_{Cd}^u(q).

Our dichotomies for exact counting. We provide complete characterizations of the complexity of counting valuations and completions satisfying a given sjfBCQ q when the input is a Codd table or a naive table and is a non-uniform or a uniform incomplete database (hence, we have eight cases in total). Our eight dichotomies express that these problems are either tractable or #P-hard, and that the tractable cases can be fully characterized by the absence of certain forbidden *patterns* in q . In essence, a pattern is simply an sjfBCQ that can be obtained from q by deleting atoms and occurrences of variables (the exact definition of this notion is given in Section 3). Our characterizations are presented in Table 1. By analyzing this table, we can draw some important conclusions as explained next.

#Comp(q) and #Val(q) are computationally difficult: For very few sjfBCQs q , the aforementioned problems can be solved in polynomial time. Take as an example the uniform setting over naive tables. Then #Val^u(q) is #P-hard as long as q contains the pattern $R(x, x)$, or $R(x) \wedge S(x, y) \wedge T(y)$,

Table 2. Our Results on the Existence of FPRAS for Solving the Problems Studied in the Article (Assuming $NP \neq RP$)

	FPRAS for counting valuations		FPRAS for counting completions	
	Non-uniform	Uniform	Non-uniform	Uniform
Naïve	Always	Always	Never	Only when q has only unary atoms
Codd	Always	Always	Never	?

or $R(x, y) \wedge S(x, y)$. That is, as long as there is an atom in q that contains a repeated variable x , or a pair (x, y) of variables that appear in an atom and both x and y appear in some other atoms in q . By contrast, for this same setting, $\#Comp^u(q)$ is $\#P$ -hard as long as q contains the pattern $R(x, y)$ or $R(x, x)$, that is, as long as there is an atom in q that is not of arity one.

$\#Val(q)$ is always easier than $\#Comp(q)$: In all of the possible versions of our problem, the tractable cases for $\#Val(q)$ are a strict superset of the ones for $\#Comp(q)$. For instance, $\#Comp_{Cd}^u(\exists x \exists y R(x, y))$ is hard, while $\#Val_{Cd}^u(\exists x \exists y R(x, y))$ is tractable.

Even counting completions is hard: While counting the total number of valuations for an incomplete database can always be done in polynomial time, observe from Table 1 that the problem $\#Comp_{Cd}^u(\exists x \exists y R(x, y))$ is $\#P$ -hard, and thus that simply counting the completions of a uniform Codd table with a single binary relation R is $\#P$ -hard. Moreover, we show that in the non-uniform case a single unary relation suffices to obtain $\#P$ -hardness.

Codd tables help but not much: We show that counting valuations is easier for Codd tables than for naïve tables. In particular, there is always an sjfBCQ q such that counting the valuations that satisfy q is $\#P$ -hard, yet it becomes tractable when restricted to the case of Codd tables. However, for counting completions, both in the uniform and non-uniform setting, the sole restriction to Codd tables presents no benefits: For every sjfBCQ q , we have that $\#Comp(q)$ (respectively, $\#Comp^u(q)$) is $\#P$ -hard if and only if $\#Comp_{Cd}(q)$ (respectively, $\#Comp_{Cd}^u(q)$) is $\#P$ -hard.

Non-uniformity complicates things: All versions of our problems become harder in the non-uniform setting. This means that in all cases there is an sjfBCQ q for which counting valuations is tractable on uniform incomplete databases, but becomes $\#P$ -hard assuming non-uniformity, and the same holds for counting completions.

Our dichotomies for approximate counting. Although $\#Val(q)$ can be $\#P$ -hard, we prove that good randomized approximation algorithms can be designed for this problem. More precisely, we give a general condition under which $\#Val(q)$ admits a **fully polynomial-time randomized approximation scheme (FPRAS)** [33]. This condition applies in particular to all unions of Boolean conjunctive queries. Remarkably, we show that this no longer holds for $\#Comp(q)$; more precisely, there exists an sjfBCQ q such that $\#Comp(q)$ does not admit an FPRAS under a widely believed complexity theoretical assumption. More surprisingly, even counting the completions of a uniform incomplete database containing a single binary relation does not admit an FPRAS under such an assumption (and in the non-uniform case, a single unary relation suffices). Generally, for sjfBCQs, we obtain seven dichotomies for our problems between polynomial-time computability of exact counting and non-admissibility of an FPRAS. The only case that we did not completely solve is that of $\#Comp_{Cd}^u(q)$. Our dichotomies for approximate counting are illustrated in Table 2.

Beyond $\#P$. It is easy to see that the problem of counting valuations is always in $\#P$, provided that the model checking problem for q is in P . This is no longer the case for counting completions, and in fact, we show that, under a complexity theoretical assumption, there is an sjfBCQ q for

which $\#Comp^u(q)$ is not in $\#P$. This does not hold if restricted to Codd tables, however, as we prove that $\#Comp_{Cd}(q)$ is always in $\#P$ when the model checking problem for q is in P .

For reasons that we explain in the article, a suitable complexity class for the problem $\#Comp(q)$ is $SpanP$, which is defined as the class of counting problems that can be expressed as the number of different accepting outputs of a nondeterministic Turing machine running in polynomial time. While we have not managed to prove that there is an sjfBCQ q for which $\#Comp(q)$ is $SpanP$ -complete, we show that this is the case for the problem of counting completions for the negation of an sjfBCQ, even in the uniform setting; that is, we show that $\#Comp^u(-q)$ is $SpanP$ -complete for some sjfBCQ q . Finally, we also show that $SpanP$ is the right complexity class for counting valuations of queries for which model checking is in NP .

Extension to queries with constants and free variables. As we said already, for pedagogical reasons, we mostly present our results by considering queries that are Boolean and that do not have constants. In Section 7, however, we explain how to extend these results to the case of queries that have free variables and that can contain constants. For the case of a query $q(\bar{x})$ with free variables \bar{x} , our counting problems are defined in the expected way; for instance, the problem $\#Val(q(\bar{x}))$ takes as input an incomplete database D , a tuple of constants \bar{a} of same arity as \bar{x} , and it outputs the number of valuations ν of D such that \bar{a} is an answer to $q(\bar{x})$ on $\nu(D)$. We then extend our dichotomies and approximation results in this setting.

The current article extends the conference article cited in Reference [8] in the following ways:

- In Reference [8], we left open the dichotomy for $\#Val_{Cd}^u(q)$, i.e., for counting valuations of sjfBCQs for Codd tables under the uniform setting. We close this case here, by finding one more hard pattern (namely, the pattern $\exists x, y R(x, y) \wedge S(x, y)$) and showing that the problem can be solved in polynomial time for all other queries.
- We added Section 7, which explains how our framework can be extended to handle queries with constants and free variables;
- Proposition 6.3, which establishes the NP -completeness of checking if a set of facts is a possible completion of an incomplete database, is new;
- Finally, full proofs of most results are included in the body of the article.

Organization of the article. We start with the main terminology used in the article in Section 2, and then present in Section 3 our four dichotomies on $\#Val(q)$ when q is an sjfBCQ, and the input incomplete database can be Codd or not, and the domain can be uniform or not. We then establish the four dichotomies on $\#Comp(q)$ in Section 4. In Section 5, we study the approximability complexity of our problems. We then give in Section 6 some general considerations about the exact complexity of the problem $\#Comp(q)$ going beyond $\#P$. We explain in Section 7 how to extend our results to queries with constants and free variables. In Section 8, we discuss related work and explain the differences with the problems considered in this article. Last, we provide some conclusions and mention possible directions for future work in Section 9.

2 PRELIMINARIES

Relational databases and conjunctive queries. A *relational schema* σ is a finite non-empty set of relation symbols written R, S, T, \dots , each with its associated *arity*, which is denoted by $arity(R)$. Let $Consts$ be a countably infinite set of constants. A *database* D over σ is a set of *facts* of the form $R(a_1, \dots, a_{arity(R)})$ with $R \in \sigma$, and where each element $a_i \in Consts$. For $R \in \sigma$, we denote by $D(R)$ the subset of D consisting of facts over R . Such a set is usually called a *relation* of D .

A Boolean query q is a query that a database D can *satisfy* (written $D \models q$) or not (written $D \not\models q$). If q is a Boolean query, then $\neg q$ is the Boolean query such that $D \models \neg q$ if and only if $D \not\models q$. A **Boolean conjunctive query (BCQ)** over σ is an FO formula of the form

$$\exists \bar{x} \left(R_1(\bar{x}_1) \wedge \dots \wedge R_m(\bar{x}_m) \right), \quad (1)$$

where all variables are existentially quantified, and where for each $i \in [1, m]$, we have that R_i is a relation symbol in σ and \bar{x}_i is a tuple of variables with $|\bar{x}_i| = \text{arity}(R_i)$. To avoid trivialities, we will always assume that $m \geq 1$, i.e., the query has at least one atom, and also that $\text{arity}(R_i) \geq 1$ for all atoms. Observe that we do not allow constants to appear in the query (but we will come back to this issue in Section 7). For simplicity, we typically write a BCQ q of the form (1) as

$$R_1(\bar{x}_1) \wedge \dots \wedge R_m(\bar{x}_m),$$

and it will be implicitly understood that all variables in q are existentially quantified. As usual, we define the semantics of a BCQ in terms of *homomorphisms*. A homomorphism from q to a database D is a mapping from the variables in q to the constants used in D such that $\{R_1(h(\bar{x}_1)), \dots, R_m(h(\bar{x}_m))\} \subseteq D$. Then, we have $D \models q$ if there exists a homomorphism from q to D . A *self-join-free BCQ* (sjfBCQ) is a BCQ such that no two atoms use the same relation symbol.

Incomplete databases. Let Nulls be a countably infinite set of nulls (also called *labeled* or *marked* nulls in the literature), which is disjoint with Consts. An *incomplete database* over schema σ is a pair $D = (T, \text{dom})$, where T is a database over σ whose facts contain elements in $\text{Consts} \cup \text{Nulls}$, and where dom is a function that associates to every null \perp occurring in D a subset $\text{dom}(\perp)$ of Consts. Intuitively, T is a database that can mention both constants and nulls, while dom tells us where nulls are to be interpreted. Following the literature, we call T a *naive table* [30].

An incomplete database $D = (T, \text{dom})$ can represent potentially many complete databases, via what are called *valuations*. A valuation of D is simply a function ν that maps each null \perp occurring in T to a constant $\nu(\perp) \in \text{dom}(\perp)$. Such a valuation naturally defines a *completion* of D , denoted by $\nu(T)$, which is the complete database obtained from T by substituting each null \perp appearing in T by $\nu(\perp)$. It is understood, since a database is a *set* of facts, that $\nu(T)$ does not contain duplicate facts. By paying attention to completions of incomplete databases that are generated exclusively by applying valuations to them, we are sticking to the so-called *closed-world* semantics of incompleteness [1, 45]. This means that the databases represented by an incomplete database $D = (T, \text{dom})$ are not open to adding facts that are not “justified” by the facts in T .

Example 2.1. Let $D = (T, \text{dom})$ be the incomplete database consisting of the naive table $T = \{S(\perp_1, \perp_1), S(a, \perp_2)\}$, and where $\text{dom}(\perp_1) = \{a, b\}$ and $\text{dom}(\perp_2) = \{a, c\}$. Let ν_1 be the valuation mapping \perp_1 to b and \perp_2 to c . Then $\nu_1(T)$ is $\{S(b, b), S(a, c)\}$. Let ν_2 be the valuation mapping both \perp_1 and \perp_2 to a . Then $\nu_2(T)$ is $\{S(a, a)\}$. However, the function ν mapping \perp_1 and \perp_2 to b is not a valuation of D , because $b \notin \text{dom}(\perp_2)$.

When every null occurs at most once in T , then D is what is called a *Codd table* [16]; for instance, the incomplete database in Example 2.1 is not a Codd table, because \perp_1 occurs twice. We also consider *uniform* incomplete databases in which the domain of every null is the same. Formally, a uniform incomplete database is a pair $D = (T, \text{dom})$, where T is a database over σ and dom is a subset of Consts. The difference now is that a valuation ν of D must simply satisfy $\nu(\perp) \in \text{dom}$ for every null of D .

We will often abuse notation and use D instead of T ; for instance, we write $\nu(D)$ instead of $\nu(T)$, or $R(a, a) \in D$ instead of $R(a, a) \in T$, or again $D(R)$ instead of $T(R)$.

Counting problems on incomplete databases. We will study two kinds of counting problems for incomplete databases: problems of the form $\#\text{Val}(q)$, which count the number of *valuations* ν that yield a completion $\nu(D)$ satisfying a given BCQ q , and problems of the form $\#\text{Comp}(q)$, which count the number of *completions* that satisfy q . The query q is assumed to be fixed, so each query gives rise to different counting problems, and we are considering the *data complexity* [53] of these problems.

Before formally introducing our problems, let us observe that they are well defined if we assume that the set of constants to which a null can be mapped to is finite. Hence, for the (default) case of an incomplete database $D = (T, \text{dom})$, we assume that $\text{dom}(\perp)$ is always a finite subset of Consts . Similarly, for the case of a uniform incomplete database $D = (T, \text{dom})$, we assume that dom is a finite subset of Consts . Finally, given a Boolean query q , we use notation $\text{sig}(q)$ for the set of relation symbols occurring in q . With these ingredients, we can define our problems for the (default) case of incomplete naive tables and a Boolean query q .

PROBLEM :	$\#\text{Val}(q)$
INPUT :	An incomplete database D over $\text{sig}(q)$
OUTPUT :	Number of valuations ν of D with $\nu(D) \models q$

PROBLEM :	$\#\text{Comp}(q)$
INPUT :	An incomplete database D over $\text{sig}(q)$
OUTPUT :	Number of completions $\nu(D)$ of D with $\nu(D) \models q$

We also consider the uniform variants of these problems, in which the input D is a uniform incomplete database over $\text{sig}(q)$, and the restriction of these problems where the input is a Codd table instead of a naive table. We then use the terms $\#\text{Val}^u(q)$, $\#\text{Comp}^u(q)$ when restricted to the uniform case, $\#\text{Val}_{Cd}(q)$, $\#\text{Comp}_{Cd}(q)$ when restricted to Codd tables, and $\#\text{Val}_{Cd}^u(q)$, $\#\text{Comp}_{Cd}^u(q)$ when both restrictions are applied.

As we will see, even though the problems $\#\text{Val}(q)$ and $\#\text{Comp}(q)$ look similar, they are of a different computational nature; this is because two distinct valuations can produce the same completion of an incomplete database. We illustrate this phenomenon in the following example:

Example 2.2. Let q be the Boolean conjunctive query $\exists x S(x, x)$, and D be the (non-uniform) incomplete database $D = (T, \text{dom})$, with $T = \{S(a, b), S(\perp_1, a), S(a, \perp_2)\}$, $\text{dom}(\perp_1) = \{a, b, c\}$ and $\text{dom}(\perp_2) = \{a, b\}$. We have depicted in Figure 1 the six valuations of D together with the completions that they define. Out of these six valuations ν , only four are such that $\nu(D) \models q$, so we have $\#\text{Val}(q)(D) = 4$. Moreover, there are only 3 distinct completions of D that satisfy q —because the first two are the same—so $\#\text{Comp}(q)(D) = 3$.

Counting complexity classes. Given two problems A, B , we write $A \leq_T^P B$ when A reduces to B under polynomial-time Turing reductions. When both A and B are counting problems, we write $A \leq_{\text{par}}^P B$ when A can be reduced to B under polynomial-time *parsimonious* reductions, i.e., when there exists a polynomial-time computable function f that transforms an input x of A to an input $f(x)$ of B such that $A(x) = B(f(x))$. We say that a counting problem is in FP when it can be solved in polynomial time. We will consider the counting complexity class #P [50] of problems that can be expressed as the number of accepting paths of a nondeterministic Turing machine running in polynomial time. Following Reference [50, 51], we define #P-hardness using Turing reductions. It is clear that $\text{FP} \subseteq \#\text{P}$. Moreover, this inclusion is widely believed to be strict. Therefore, proving

$(v(\perp_1), v(\perp_2))$	(a, a)	(a, b)	(b, a)	(b, b)	(c, a)	(c, b)
$v(D)$	S	S	S	S	S	S
	$a \ b$	$a \ b$	$a \ b$	$a \ b$	$a \ b$	$a \ b$
	$a \ a$	$a \ a$	$b \ a$	$b \ a$	$c \ a$	$c \ a$
			$a \ a$		$a \ a$	
$v(D) \models q?$	Yes	Yes	Yes	No	Yes	No

Fig. 1. The six valuations of the (non-uniform) incomplete database $D = (T, \text{dom})$ with $T = \{S(a, b), S(\perp_1, a), S(a, \perp_2)\}$ from Example 2.2, and their corresponding completions. The Boolean conjunctive query q is $\exists x S(x, x)$.

that a counting problem is #P-hard implies that it cannot be solved in polynomial time under such an assumption.

3 DICHOTOMIES FOR COUNTING VALUATIONS

In this section, for a fixed sjfBCQ q , we study the complexity of the problem of computing, given as input an incomplete database D , the number of valuations v of D such that $v(D)$ satisfies q . Recall that we have four cases to consider for this problem depending on whether we focus on naive or on Codd tables, where nulls are restricted to appear at most once, and whether we focus on non-uniform or uniform incomplete databases, where nulls are restricted to have the same domain. Our specific goal then is to understand whether the problem is tractable (in FP) or #P-hard in these scenarios, depending on the shape of q .

To this end, the shape of an sjfBCQ q will be characterized by the presence or absence of certain specific patterns. In the following definition, we introduce the necessary terminology to formally talk about the presence of a pattern in a query:

Definition 3.1. Let q, q' be sjfBCQs. We say that q' is a *pattern* of q if q' can be obtained from q by using an arbitrary number of times and in any order the following operations: deleting an atom, deleting an occurrence of a variable, renaming a relation to a fresh one, renaming a variable to a fresh one, and reordering the variables in an atom.¹

Example 3.2. Recall that we always omit existential quantifiers in Boolean queries. Then, we have that $q' = R'(u, u, y) \wedge S'(z)$ is a pattern of $q = R(u, x, u) \wedge S'(y, y) \wedge T(x, s, z, s)$. Indeed, q' can be obtained from q by deleting atom $T(x, s, z, s)$, renaming $R(u, x, u)$ as $R'(u, x, u)$ to obtain $R'(u, x, u) \wedge S'(y, y)$, reordering the variables in $R'(u, x, u)$ to obtain $R'(u, u, x) \wedge S'(y, y)$, renaming variable y into z to obtain $R'(u, u, x) \wedge S'(z, z)$, deleting the second variable occurrence in $S'(z, z)$ to obtain $R'(u, u, x) \wedge S'(z)$, and finally renaming variable x into y to obtain q' .

We point out that in Definition 3.1, the important parts are those about deleting atoms and variable occurrences. The parts about reordering variable occurrences inside an atom and about renaming relations and variables to fresh ones have obviously no effect on the complexity of the problem²; these are only here to allow us to formally say, for instance, that “ $R(x)$ is a pattern of $R(y)$,” or that “ $S(x, u, x)$ is a pattern of $T(w, z, z)$ ” (as these are, in essence, the same queries).

¹We remind the reader that we assume all sjfBCQs to contain at least one atom and that all atoms must contain at least one variable.

²This is, in particular, because the conjunctive queries we consider have no self-joins (otherwise, reordering variables inside an atom could change the complexity).

In the following general lemma, we show that if q' is a pattern of q , then each of the problems considered in this section is as hard for q as it is for q' . Recall in this result that unless stated otherwise, our problems are defined for naive tables under the non-uniform setting.

LEMMA 3.3. *Let q, q' be sjfBCQs such that q' is a pattern of q . Then, we have $\#\text{Val}(q') \leq_{\text{par}}^{\text{P}} \#\text{Val}(q)$. Moreover, the same results hold if we restrict to Codd tables, and/or to the uniform setting.*

PROOF. We first present the proof for $\#\text{Val}(q') \leq_{\text{par}}^{\text{P}} \#\text{Val}(q)$, that is, for naive tables in the non-uniform setting. First, observe that we can assume without loss of generality that we did not reorder the variables in the atoms nor renamed relation names or variables by fresh ones, because, as mentioned above, this does not change the complexity of the problem.³ We can then write q as $R_1(\bar{x}_1) \wedge \dots \wedge R_m(\bar{x}_m)$ and q' as $R_{j_1}(x'_{j_1}) \wedge \dots \wedge R_{j_p}(x'_{j_p})$, where $1 \leq j_1 < \dots < j_p \leq m$ and x'_{j_k} is obtained from \bar{x}_{j_k} by deleting some variable occurrences but not all,⁴ and the other atoms have been deleted. Let D' be an incomplete database input of $\#\text{Val}(q')$. Let A be the set of constants that are appearing in D' or are in a domain of some null occurring in D' . For $1 \leq k \leq p$, we construct the relation $D(R_{j_k})$ from the relation $D'(R_{j_k})$. Let us assume that \bar{x}_{j_k} is the tuple (x_1, \dots, x_r) (with some variables possibly being equal). We initialize $D(R_{j_k})$ to be empty, and then for every tuple \bar{t}' in $D'(R_{j_k})$, we add to $D(R_{j_k})$ all the tuples \bar{t} that can be obtained from \bar{t}' in the following way for $1 \leq i \leq r$:

- (a) If x_i is a variable occurrence that has not been deleted from \bar{x}_{j_k} , then copy the element (constant or null) of \bar{t}' corresponding to that variable occurrence to the i th position of \bar{t} ;
- (b) Otherwise, if x_i is a variable occurrence that has been deleted from \bar{x}_{j_k} , then fill the i th position of \bar{t} with every possible constant from A .

Then, we construct the relations $D(R_i)$ where R_i does not appear in q' (this can happen if we have deleted the atom $R_i(\bar{x}_i)$) by filling it with every possible R_i -fact over A . We leave the domains of all nulls unchanged. The whole construction can be performed in polynomial time (this uses the fact that q is assumed to be fixed, so the arities of the relations mentioned in q are fixed). Hence, it only remains to be checked that $\#\text{Val}(q')(D') = \#\text{Val}(q)(D)$, that is, that the reduction works and is indeed parsimonious. It is clear that the valuations of D' are exactly the same as the valuations of D (because they have the same sets of nulls). Hence, it is enough to verify that for every valuation ν , we have $\nu(D') \models q'$ if and only if $\nu(D) \models q$. Let h' be a homomorphism from q' to $\nu(D')$ witnessing that $\nu(D') \models q'$ (i.e., we have $h'(q) \subseteq \nu(D')$). Then h' can clearly be extended in the expected way into a homomorphism h from q to $\nu(D)$; this is, in particular, thanks to the fact that we filled the missing columns with every possible constant. Conversely, let h be a homomorphism from q to $\nu(D)$ witnessing that $\nu(D) \models q$. Then the restriction h' of h to the variables occurring in q' is such that $h(q') \subseteq \nu(D')$, hence, we have $\nu(D') \models q'$. This concludes the proof for the case of naive tables in the non-uniform setting. For the cases of Codd tables and/or for the uniform setting, the reduction is exactly the same. Indeed, the domains of the nulls are unchanged, and it is clear that the presented construction preserves the property of being a Codd table. \square

The idea is then to show the #P-hardness of our problems for some simple patterns, which we then combine with Lemma 3.3 and some tractability proofs to obtain the desired dichotomies. Our

³Formally, one can first check that we can assume without loss of generality that q' was obtained from q by first deleting some atoms and variable occurrences to obtain a query q'' , and then performing some renamings and variable reorderings to obtain q' (that is, we can always push the renaming and reordering parts at the end of the transformation). But then, since $\#\text{Val}(q'')$ and $\#\text{Val}(q')$ are obviously the same problem, we can assume $q' = q''$.

⁴See Footnote 1.

findings are summarized in the first two columns of Table 1 in the introduction. We first focus on the two dichotomies for the non-uniform setting in Section 3.1, and then we move to the case of uniform incomplete databases in Section 3.2. We explicitly state when a #P-hardness result holds even in the restricted setting in which there is a fixed domain over which nulls are interpreted. In other words, when there is a fixed domain A such that the incomplete databases used in the reductions are of the form $D = (T, \text{dom})$ and $\text{dom}(\perp) \subseteq A$, for each null \perp of T .

3.1 The Complexity on the Non-uniform Case

In this section, we study the complexity of the problems $\#\text{Val}(q)$ and $\#\text{Val}_{Cd}(q)$, providing dichotomy results in both cases. We start by proving the #P-hardness results needed for these dichotomies. We first show that $\#\text{Val}(R(x, x))$ is #P-hard by actually proving that hardness holds already in the uniform case.

PROPOSITION 3.4. $\#\text{Val}^u(R(x, x))$ is #P-hard. This holds even in the restricted setting in which all nulls are interpreted over the same fixed domain $\{1, 2, 3\}$.

PROOF. We reduce from the problem of counting the number of 3-colorings of a graph $G = (V, E)$, which is #P-hard [32]. For every node $v \in V$, we have a null \perp_v , and for every edge $\{u, v\} \in E$, we have the facts $R(\perp_v, \perp_u)$ and $R(\perp_u, \perp_v)$. The domain of the nulls is $\{1, 2, 3\}$. It is then clear that the number of valuations of the constructed database that do not satisfy $R(x, x)$ is exactly the number of 3-colorings of G . Since the total number of valuations can be computed in PTIME, this concludes the reduction. \square

The next pattern that we consider is $R(x) \wedge S(x)$. This time, we can show #P-hardness of the problem even for Codd databases.

PROPOSITION 3.5. $\#\text{Val}_{Cd}(R(x) \wedge S(x))$ is #P-hard.

PROOF. We start by recalling the setting of *consistent query answering under key constraints*. Intuitively, in this case, we are given a set Σ of keys and a database D that does not necessarily satisfy Σ . Then the task is to reason about the set of all *repairs* of D with respect to Σ [9]. In our context, this means that one wants to count the number of repairs of D with respect to Σ that satisfy a given CQ q . When q and Σ are fixed, we call this problem $\#\text{Repairs}(q, \Sigma)$; see, e.g., Reference [40]. We formalize these notions below.

Here, we focus on the case when Σ is a set of *primary* keys. Recall that this means that each relation name $R \in \sigma$ of arity n comes equipped with its own key, i.e., $\text{key}(R) = A$, where $A = \emptyset$ or $A = [1, \dots, p]$ for some $p \in \{1, \dots, n\}$. Henceforth, D is *inconsistent* with respect to Σ if there is a relation name $R \in \sigma$ and facts $R(\bar{a}), R(\bar{b}) \in D$ with $\bar{a} \neq \bar{b}$ such that

$$\text{key}(R) = A \quad \text{and} \quad \pi_A(\bar{a}) = \pi_A(\bar{b}).$$

In this case, we say that the pair $(R(\bar{a}), R(\bar{b}))$ is *key-violating*. Let us define a *block* in a database D with respect to a set Σ of primary keys to be any maximal set B of facts from D such that the facts in B are pairwise key-violating. A *repair* of D with respect to Σ is a subset D' of D that is obtained by choosing exactly one tuple from each block of D with respect to Σ .

Let us consider a schema σ with two binary relations R' and S' , such that $\text{key}(R') = \text{key}(S') = \{1\}$. That is, the first attribute of both R' and S' defines a key over such relations. We define this set of keys over σ to be Σ . Also, let $q = \exists x, y, z(R'(y, x) \wedge S'(z, x))$. For simplicity, we write the pair (q, Σ) as $R'(\underline{y}, x) \wedge S'(\underline{z}, x)$. The problem $\#\text{Repairs}(R'(\underline{y}, x) \wedge S'(\underline{z}, x))$, which given a database

D' over schema σ aims at computing the number of repairs of D' under Σ that satisfy q , is known to be #P-complete [40].⁵

Now, observe that the #P-hardness of $\#\text{Val}_{Cd}(R(x) \wedge S(x))$ easily follows from the hardness of the problem $\#\text{Repairs}(R'(\underline{y}, x) \wedge S'(\underline{z}, x))$. In fact, let D' be a database with binary relation R', S' . We construct an incomplete Codd database D with unary relations R, S as follows: For every constant a that appears in the first attribute of R' , we have a tuple $R(\perp)$ in D , where \perp is a fresh null, and we set $\text{dom}(\perp) = \{b \mid R'(a, b) \in D'\}$. For every constant a that appears in the first attribute of S' , we have a tuple $S(\perp)$ in D , where \perp is a fresh null, and we set $\text{dom}(\perp) = \{b \mid S'(a, b) \in D'\}$. It is then clear that the number of repairs of D' that satisfy $R'(\underline{y}, x) \wedge S'(\underline{z}, x)$ is equal to the number of valuations of D that satisfy $R(x) \wedge S(x)$, thus concluding the proof. We point out here that another proof of Proposition 3.5, which uses different techniques, can be found in the conference version of the article [8] (the proof that we presented here is shorter). \square

Already with Propositions 3.5 and 3.4, we have all the relevant hard patterns for the non-uniform setting. We start by proving our dichotomy result for naive tables, which is our default case.

THEOREM 3.6 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x) \wedge S(x)$ is a pattern of q , then $\#\text{Val}(q)$ is #P-complete. Otherwise, $\#\text{Val}(q)$ is in FP.*

PROOF. The #P-hardness part of the claim follows from the last two propositions and from Lemma 3.3. We explain why the problems are in #P right after this proof. When q does not have any of these two patterns, then all variables have exactly one occurrence in q . This implies that every valuation ν of D is such that $\nu(D)$ satisfies q (except when one relation is empty, in which case the result is simply zero). We can then compute the total number of valuations in FP by simply multiplying the sizes of the domains of every null in D . \square

Notice that in this theorem, the membership of $\#\text{Val}(q)$ in #P can be established by considering a nondeterministic Turing Machine M that, with input a non-uniform incomplete database D , guesses a valuation ν of D and verifies whether $\nu(D)$ satisfies q . This machine works in polynomial time, as we can verify whether $\nu(D)$ satisfies q in polynomial time (since q is a fixed FO query). Then, given that $\#\text{Val}(q)(D)$ is equal to the number of accepting runs of M with input D , we conclude that $\#\text{Val}(q)$ is in #P. Obviously, the same idea works for Codd tables, that is, $\#\text{Val}_{Cd}(q)$ is also in #P. But with this restriction, we obtain more tractable cases, as shown by the following dichotomy result:

THEOREM 3.7 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x) \wedge S(x)$ is a pattern of q , then $\#\text{Val}_{Cd}(q)$ is #P-complete. Otherwise, $\#\text{Val}_{Cd}(q)$ is in FP.*

PROOF. We only need to prove the tractability claim, since hardness follows from Proposition 3.5 and Lemma 3.3. We will assume without loss of generality that D contains no constants, as we can introduce a fresh null with domain $\{c\}$ for every constant c appearing in D , and the result is again a Codd table, and this does not change the output of the problem. Let q be $R_1(\bar{x}_1) \wedge \dots \wedge R_m(\bar{x}_m)$. Observe that, since q does not have $R(x) \wedge S(x)$ as a pattern, then any two atoms cannot have a variable in common. But then, since D is a Codd table, we have

$$\#\text{Val}_{Cd}(q)(D) = \prod_{i=1}^m \#\text{Val}_{Cd}(R_i(\bar{x}_i))(D(R_i)).$$

⁵To see that Reference [40] establishes the hardness of $q = R'(\underline{y}, x) \wedge S'(\underline{z}, x)$, first apply their rewrite rule R7 (from Figure 6) to obtain $q' = R'(\underline{y}, x) \wedge S'(\underline{x}, x)$, then apply rewrite rule R10 to obtain $q'' = R'(\underline{y}, x) \wedge S'(\underline{x}, a)$. Then, q'' is #P-hard by Lemma 19, and so is q by Lemma 7.

Hence, it is enough to show how to compute $\#\text{Val}_{Cd}(R_i(\bar{x}_i))(D(R_i))$ for every $1 \leq i \leq m$. Let $\bar{t}_1, \dots, \bar{t}_n$ be the tuples of $D(R_i)$. Let us write $\rho(\bar{t}_j)$ for the number of valuations of the nulls appearing in \bar{t}_j that do not match \bar{x}_i . Clearly, $\#\text{Val}_{Cd}(R_i(\bar{x}_i))(D(R_i)) = \prod_{\perp \text{ appears in } D(R_i)} |\text{dom}(\perp)| - \prod_{j=1}^n \rho(\bar{t}_j)$, so we only have to show how to compute $\rho(\bar{t}_j)$ for $1 \leq j \leq n$. Since we can easily compute the total number of valuations of \bar{t}_j , it is enough to show how to compute the number of valuations of \bar{t}_j that match \bar{x}_i . For every variable x that appears in \bar{x}_i , compute the size of the intersection of the domains of the corresponding nulls in \bar{t}_j , and denote it s_x . Then the number of valuations of \bar{t}_j that match \bar{x}_i is simply $\prod_{x \text{ appears in } \bar{x}_i} s_x$. This concludes the proof. \square

At this stage, we have completed the first column of Table 1, and we also know that $R(x, x)$ is a hard pattern in the uniform setting for naive tables (but not for Codd tables, by Theorem 3.7). In the next section, we treat the uniform setting.

3.2 The Complexity on the Uniform Case

In this section, we study the complexity of the problems $\#\text{Val}^u(q)$ and $\#\text{Val}_{Cd}^u(q)$, again providing dichotomy results in both cases.

3.2.1 Naïve Tables. We start our investigation with the case of naive tables. In Proposition 3.4, we already showed that $\#\text{Val}^u(R(x, x))$ is #P-hard. In the following proposition, we identify two other simple queries for which this problem is still intractable.

PROPOSITION 3.8. *$\#\text{Val}^u(R(x) \wedge S(x, y) \wedge T(y))$ and $\#\text{Val}^u(R(x, y) \wedge S(x, y))$ are both #P-hard. This holds even in the restricted setting in which all nulls are interpreted over the same fixed domain $\{0, 1\}$.*

PROOF. We reduce both problems from the problem of counting the number of independent sets in a graph (denoted by #IS), which is #P-complete [44]. We start with $\#\text{Val}^u(R(x) \wedge S(x, y) \wedge T(y))$. Let $q = R(x) \wedge S(x, y) \wedge T(y)$ and $G = (V, E)$ be a graph. Then, we define an incomplete database D as follows: For every node $v \in V$, we have a null \perp_v , and the uniform domain is $\{0, 1\}$. For every edge $\{u, v\} \in E$, we have facts $S(\perp_u, \perp_v)$ and $S(\perp_v, \perp_u)$ in D . Finally, we have facts $R(1)$ and $T(1)$ in D . For a valuation ν of the nulls, consider the corresponding subset S_ν of nodes of G , given by $S_\nu = \{t \in V \mid \nu(\perp_t) = 1\}$. This is a bijection between the valuations of the database and the node subsets of G . Moreover, we have that $\nu(D) \models q$ if and only if S_ν is an independent set of G . Since the total number of valuations of D is $2^{|V|}$, we have that the number of independent sets of G is equal to $2^{|V|} - \#\text{Val}^u(q)(D)$. Hence, we conclude that $\#\text{IS} \leq_{\text{T}}^{\text{P}} \#\text{Val}^u(q)$. The idea is similar for $\#\text{Val}^u(R(x, y) \wedge S(x, y))$: We encode the graph with the relation S in the same way, and this time, we add the fact $R(1, 1)$. \square

As shown in the following result, it turns out that the three aforementioned patterns are enough to fully characterize the complexity of counting valuations for naive tables in the uniform setting.

THEOREM 3.9 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x) \wedge S(x, y) \wedge T(y)$ or $R(x, y) \wedge S(x, y)$ is a pattern of q , then $\#\text{Val}^u(q)$ is #P-complete. Otherwise, $\#\text{Val}^u(q)$ is in FP.*

The #P-completeness part of the claim follows directly from what we have proved already. Here, the most challenging part of the proof is actually the tractability part. We only present a simple example to give an idea of the proof technique, and we defer the full proof to Appendix A.1. We will use the following definition: Given $n, m \in \mathbb{N}$, let us write $\text{surj}_{n \rightarrow m}$ for the number of surjective functions from $\{1, \dots, n\}$ to $\{1, \dots, m\}$. By an inclusion–exclusion argument, one can show that $\text{surj}_{n \rightarrow m} = \sum_{i=0}^{m-1} (-1)^i \binom{m}{i} (m-i)^n$ (for instance, see Reference [3]). It is clear that this can be computed in FP when n and m are given in unary.

Example 3.10. Let q be the sjfBCQ $R(x) \wedge S(x)$, and D be an incomplete database over relations R, S . Notice that q does not have any of the patterns mentioned in Theorem 3.9. We will show that $\#\text{Val}^u(q)$ is in FP. Since q contains only two unary atoms, we can also assume without loss of generality that the input D is a Codd table (otherwise all valuations are satisfying).

Since we can compute in FP the total number of valuations, it is enough to show how to compute the number of valuations of D that do not satisfy q . Let dom be the uniform domain, d be its size, n_R (respectively, n_S) be the number of nulls in $D(R)$ (respectively, in $D(S)$) and C_R (respectively, C_S) be the set of constants occurring in $D(R)$ (respectively, in $D(S)$), with c_R (respectively, c_S) its size. We can assume without loss of generality that $C_R \cap C_S = \emptyset$, as otherwise all the valuations are satisfying, and this is computable in PTIME. Furthermore, we can also assume that $C_R \cup C_S \subseteq \text{dom}$, since we can remove the constants that are not in dom , as these can never match.

Let $M := \text{dom} \setminus (C_R \cup C_S)$, and m its size (i.e., with our assumptions, we have $m = d - c_R - c_S$). Fix some subsets $M' \subseteq M$ and $R' \subseteq C_R$. The quantity $\text{surj}_{n_R \rightarrow |M'| + |R'|}$ then counts the number of valuations of the nulls of $D(R)$ that span exactly $M' \cup R'$. Moreover, letting ν_R be a valuation of the nulls of $D(R)$ that spans exactly $M' \cup R'$, the quantity $(d - c_R - |M'|)^{n_S}$ is the number of ways to extend ν_R into a valuation ν of all the nulls of D so $\nu(D) \not\models q$: Indeed, every null of $D(S)$ can take any value in $\text{dom} \setminus (C_R \cup M')$. The number of valuations of D that do not satisfy q is then (keeping in mind that a null in $D(R)$ cannot take a value in C_S):

$$\sum_{\substack{M' \subseteq M \\ R' \subseteq C_R}} \text{surj}_{n_R \rightarrow |M'| + |R'|} \times (d - c_R - |M'|)^{n_S}$$

and, since the summands only depends on the sizes of M' and R' , this is equal to

$$\sum_{\substack{0 \leq m' \leq m \\ 0 \leq r' \leq c_R}} \binom{m}{m'} \binom{c_R}{r'} \text{surj}_{n_R \rightarrow m' + r'} \times (d - c_R - m')^{n_S}.$$

This last expression can clearly be computed in PTIME.⁶

3.2.2 Codd Tables. We conclude this section by turning our attention to the case of Codd tables. Notice that none of the results proved so far provides a hard pattern in this case. We identify in the following proposition a simple query for which the problem is intractable.

PROPOSITION 3.11. $\#\text{Val}_{C_d}^u(R(x) \wedge S(x, y) \wedge T(y))$ is #P-hard.

PROOF. We reduce from the problem of counting the number of independent sets of a bipartite (simple) graph, written #BIS, which is #P-hard [44]. Let $G = (X \sqcup Y, E)$ be a bipartite graph. Without loss of generality, we can assume that $|X| = |Y| = n$; Indeed, if $|X| < |Y|$, then we could simply add $|Y| - |X|$ isolated nodes to complete the graph, which simply multiplies the number of independent sets by $2^{|Y| - |X|}$. Also, observe that counting the number of independent sets of G is the same as counting the number of pairs (S_1, S_2) with $S_1 \subseteq X, S_2 \subseteq Y$, such that $(S_1 \times S_2) \cap E = \emptyset$. We will call such a pair an *independent pair*. For $0 \leq i, j \leq n$, let $Z_{i,j}$ be the number of independent pairs (S_1, S_2) such that $|S_1| = i$ and $|S_2| = j$. It is clear that (\star) the number of independent sets of G is then $\#\text{BIS}(G) = \sum_{0 \leq i, j \leq n} Z_{i,j}$. The idea of the reduction is to construct in polynomial time $(n+1)^2$ incomplete databases $D_{a,b}$ for $0 \leq a, b \leq n$ such that, letting $C_{a,b}$ be the number of valuations ν of $D_{a,b}$ with $\nu(D_{a,b}) \not\models R(x) \wedge S(x, y) \wedge T(y)$, the values of the variables $Z_{i,j}$ and $C_{i,j}$ form a linear system of equations $\mathbf{AZ} = \mathbf{C}$, with \mathbf{A} an invertible matrix. This will allow us, using $(n+1)^2$ calls to an oracle for $\#\text{Val}_{C_d}^u(R(x) \wedge S(x, y) \wedge T(y))$, to recover the $Z_{i,j}$ values, and then to compute $\#\text{BIS}(G)$

⁶Note that, in the sum, we do not need to specify that $m' + r' \leq n_R$, as when $a < b$, we have $\text{surj}_{a \rightarrow b} = 0$.

using (\star) . We now explain how we construct $D_{a,b}$ from G for $0 \leq a, b \leq n$, and define \mathbf{A} . First, we fix an arbitrary linear order x_1, \dots, x_n of X , and similarly y_1, \dots, y_n for Y . The database $D_{a,b}$ has constants a_i for $1 \leq i \leq n$, and has a fact $S(a_i, a_j)$ whenever $(x_i, y_j) \in E$. It has nulls \perp_1, \dots, \perp_a and facts $R(\perp_i)$ for $1 \leq i \leq a$ (if $a = 0$ there are no such nulls and facts), and nulls $\perp'_1, \dots, \perp'_b$ and facts $T(\perp'_i)$ for $1 \leq i \leq b$; in particular, this is a Codd table. The uniform domain of the nulls is $\{a_i \mid 1 \leq i \leq n\}$. Given a valuation ν of $D_{a,b}$, let $P(\nu)$ be the pair of subsets of V defined by

$$P(\nu) \stackrel{\text{def}}{=} (\{x_i \mid \exists 1 \leq k \leq a \text{ s.t. } \nu(\perp_k) = a_i\}, \{y_i \mid \exists 1 \leq k \leq b \text{ s.t. } \nu(\perp'_k) = a_i\}).$$

One can then easily check that the following two claims hold:

- For every valuation ν of $D_{a,b}$, we have that $\nu(D_{a,b}) \not\models R(x) \wedge S(x, y) \wedge T(y)$ iff $P(\nu)$ is an independent pair of G^7 ;
- For every independent pair (S_1, S_2) of G , there are exactly $\text{surj}_{a \rightarrow |S_1|} \times \text{surj}_{b \rightarrow |S_2|}$ valuations ν such that $P(\nu) = (S_1, S_2)$.

But then, we have $C_{a,b} = \sum_{0 \leq i, j \leq n} (\text{surj}_{a \rightarrow i} \times \text{surj}_{b \rightarrow j}) Z_{i,j}$. In other words, we have the linear system of equations $\mathbf{AZ} = \mathbf{C}$, where \mathbf{A} is the $(n+1)^2 \times (n+1)^2$ matrix defined by $\mathbf{A}_{(a,b),(i,j)} \stackrel{\text{def}}{=} \text{surj}_{a \rightarrow i} \times \text{surj}_{b \rightarrow j}$. This matrix is the Kronecker product $\mathbf{A}' \otimes \mathbf{A}'$ of the $(n+1) \times (n+1)$ matrix with entries $\mathbf{A}'_{a,i} \stackrel{\text{def}}{=} \text{surj}_{a \rightarrow i}$. Since \mathbf{A}' is a triangular matrix with non-zero coefficients on the diagonal, it is invertible, hence, so is \mathbf{A} , which concludes the proof. \square

Note that in Proposition 3.8, we proved that $\#\text{Val}^u(R(x) \wedge S(x, y) \wedge T(y))$ is $\#\text{P}$ -hard in the general case where naive tables are allowed. Hence, the hardness of that query for naive tables was in fact a consequence of Proposition 3.11. However, we decided to provide a separate proof for Proposition 3.8, because in this case intractability holds already when nulls are interpreted over the fixed domain $\{0, 1\}$, whereas we do not know if this is true for Codd tables.

The second pattern that we show is hard for $\#\text{Val}^u$ for Codd tables is $R(x, y) \wedge S(x, y)$. Again, notice that we already showed this query to be hard in the case of naive tables (as Proposition 3.8), even for a fixed domain of $\{0, 1\}$. In the case of Codd tables, hardness still holds, but the proof is more complicated and uses domains of unbounded size (which is why we provide separate proofs). We show:

PROPOSITION 3.12. $\#\text{Val}_{\text{Codd}}^u(R(x, y) \wedge S(x, y))$ is $\#\text{P}$ -hard.

PROOF. Let q be the query $R(x, y) \wedge S(x, y)$. We reduce from the problem of counting the number of matchings of a 2–3–regular bipartite graph, which is $\#\text{P}$ -complete [14, 55]. Let $G = (A \sqcup B, E)$ be a 2–3–regular bipartite graph, with the nodes in A having degree 3 and those in B having degree 2, and let $n_A \stackrel{\text{def}}{=} |A|$, and $n_B \stackrel{\text{def}}{=} |B|$. Notice that, since G is 2–3–regular, we have that $n_B = \frac{3n_A}{2}$. We say that a set $S \subseteq E$ of edges of G is an *A-semimatching* if every node a of A is adjacent to at most 1 edge of S ; formally, for every $a \in A$, we have $|\{e \in S \mid a \in e\}| \leq 1$. The *type* of an *A-semimatching* S is the number c of nodes in B that are adjacent to exactly 2 edges of S ; formally, $c \stackrel{\text{def}}{=} |\{b \in B \mid |\{e \in S \mid b \in e\}| = 2\}|$. For $0 \leq c \leq n_B$, we write T_c for the number of *A-semimatchings* of G of type c . Observe then that T_0 is simply the number of matchings of G . The idea of the reduction is then as follows: We will construct databases D_k for $0 \leq k \leq n_B$ such that, letting C_k be the number of valuations ν of D_k such that $\nu(D_k) \not\models q$, the values of the variables T_k and C_k form a system of linear equations $\mathbf{AT} = \mathbf{C}$, with \mathbf{A} and invertible matrix. This will allow

⁷This observation, and in fact the idea of reducing from $\#\text{BIS}$, is due to Antoine Amarilli.

us, using $n_B + 1$ calls to an oracle for $\#\text{Val}_{Cd}^u(q)$, to recover the values T_k , and thus to obtain T_0 in polynomial time, that is, the number of matchings of G . We now explain how to construct the database D_k for $0 \leq k \leq n_B$. In what follows, we use the convention that $\{1, \dots, k\} = \emptyset$ when $k = 0$. The database D_k contains the following facts:

- (1) One fact $R(a, \perp_a)$ for every $a \in A$;
- (2) One fact $S(\perp_b, b)$ for every $b \in B$;
- (3) One fact $S(\perp_{a'}, a')$ for every $a \in A$ where a' is a fresh constant. In particular, observe that there are n_A such facts. In what follows, we will write $A' \stackrel{\text{def}}{=} \{a' \mid a \in A\}$;
- (4) One fact $S(a_1, a_2')$ for every $(a_1, a_2) \in A \times A$ with $a_1 \neq a_2$;
- (5) One fact $S(a, i)$ for every $(a, i) \in A \times \{1, \dots, k\}$;
- (6) One fact $S(u, v)$ for every $(u, v) \in (A \cup B)^2$ such that $\{u, v\}$ is not in E ;
- (7) One fact $R(u, v)$ for every $(u, v) \in (A' \cup B)^2$.

And finally, the (uniform) domain for all the nulls is $\text{dom} \stackrel{\text{def}}{=} A \cup B \cup A' \cup \{1, \dots, k\}$. Note that this is indeed a Codd database. Now, let us compute C_k , the number of valuations ν of D_k such that $\nu(D_k) \not\models q$. For such a valuation ν of D_k such that $\nu(D_k) \not\models q$, observe that (\star) for every $a \in A$ it holds that $\nu(\perp_a)$ is either a' or is one of the nodes in B that is a neighbor of a ; this is because otherwise, the facts from (4–6) would make the query be satisfied. Then, for a valuation ν of D_k such that $\nu(D_k) \not\models q$, let us define $\text{SM}(\nu) \stackrel{\text{def}}{=} \{\{a, b\} \mid (a, b) \in (A \times B) \cap R(\nu(D_k))\}$. Because of (\star) , observe that $\text{SM}(\nu)$ is a subset of E , and that it is in fact an A -semimatching. We can then partition the valuations ν of D_k with $\nu(D_k) \not\models q$ according to the type of $\text{SM}(\nu)$ as follows:

$$\{\nu \mid \nu \text{ is a valuation of } D_k \text{ with } \nu(D_k) \not\models q\} = \bigsqcup_{0 \leq c \leq n_B} \bigsqcup_{S: S \text{ is an } A\text{-semimatching of } G \text{ of type } c} \bigsqcup_{\substack{\nu: \nu \text{ valuation of } D_k \\ \nu(D_k) \not\models q \\ \text{SM}(\nu) = S}} \{\nu\}. \quad (2)$$

Fix an A -semimatching S of G of type $c \in \{0, \dots, n_B\}$, and let us count how many valuations ν of D_k there are such that $\nu(D_k) \not\models q$ and $\text{SM}(\nu) = S$. First, observe that for such a valuation ν , the value of $\nu(\perp_a)$ for every $a \in A$ is forced: It is a' if a is adjacent to no edge of S , and otherwise it is b for the unique $\{a, b\} \in S$. Therefore, we have to count how many possibilities there are for the remaining nulls, those of the form $\perp_{a'}$ and \perp_b from facts (2–3). We have:

- For a null \perp_b such that b is adjacent to two edges of S , there are $n_A + k - 2$ possible values to not satisfy the query. Note that there are c such nulls \perp_b .
- For a null $\perp_{a'}$ such that a is not adjacent to an edge in S (so we know that $\nu(\perp_a) = a'$), there are $n_A + k - 1$ possible values to not satisfy the query. For a null \perp_b such that b is adjacent to exactly one edge $\{a, b\}$ of S (i.e., we have $\nu(\perp_a) = b$) there are again $n_A + k - 1$ possible values to not satisfy the query. Observe that in total there are $n_A - 2c$ such nulls $\perp_{a'}$ or \perp_b , because S is an A -semimatching.
- For a null $\perp_{a'}$ such that a is adjacent to an edge in S (so we know that $\nu(\perp_a) \neq a'$) there are $n_A + k$ possibilities, and similarly for a null \perp_b such that b is not adjacent to an edge in S there are $n_A + k$ possible values. By the previous two items, in total there are $n_A + n_B - c - (n_A - 2c) = n_B + c$ such nulls $\perp_{a'}$ or \perp_b .

Therefore, there are exactly $(n_A + k - 2)^c (n_A + k - 1)^{n_A - 2c} (n_A + k)^{n_B + c}$ valuations ν of D_k such that $\nu(D_k) \not\models q$ and $\text{SM}(\nu) = S$. Since this depends only on the type of the A -semimatching S (and not on S itself), we obtain from Equation (2) that

$$\begin{aligned}
C_k &= |\{v \mid v \text{ is a valuation of } D_k \text{ with } v(D_k) \neq q\}| \\
&= \sum_{0 \leq c \leq n_B} T_c \times (n_A + k - 2)^c (n_A + k - 1)^{n_A - 2c} (n_A + k)^{n_B + c}.
\end{aligned}$$

That is, we have the linear system of equations $\mathbf{AT} = \mathbf{C}$, with $A_{k,c} = (n_A + k - 2)^c (n_A + k - 1)^{n_A - 2c} (n_A + k)^{n_B + c}$ for $0 \leq c, k \leq n_B$. But observe that we have $\mathbf{A} = \mathbf{DV}$, with \mathbf{D} being the diagonal $(n_B + 1) \times (n_B + 1)$ matrix with entries $(n_A + k - 2)^{n_A} (n_A + k)^{n_B}$, and \mathbf{V} being the $(n_B + 1) \times (n_B + 1)$ Vandermonde matrix with coefficients $\frac{(n_A + k - 2)(n_A + k)}{(n_A + k - 1)^2}$ for $0 \leq k \leq n_B$. Hence, to show that \mathbf{A} is invertible, we only need to argue that the coefficients of this Vandermonde matrix \mathbf{V} are all distinct. But for the function $f_{n_A}(x) \stackrel{\text{def}}{=} \frac{(n_A + x - 2)(n_A + x)}{(n_A + x - 1)^2}$, one can check that $f'_{n_A}(x) = \frac{2}{(n_A + x - 1)^3}$, so f_{n_A} is strictly increasing on $[0, n_B]$ (we assume $n_A \geq 2$ without loss of generality), so all the coefficients are indeed distinct. This concludes the proof. \square

As we show next, the patterns from Propositions 3.12 and 3.11 are the only hard patterns for $\#\text{Val}_{Cd}^u$. The following is then the last dichotomy of this section.

THEOREM 3.13 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x) \wedge S(x, y) \wedge T(y)$ or $R(x, y) \wedge S(x, y)$ is a pattern of q , then $\#\text{Val}_{Cd}^u(q)$ is $\#\text{P}$ -complete. Otherwise, $\#\text{Val}_{Cd}^u(q)$ is in FP.*

We only need to show the tractability part of that claim, as hardness follows from Propositions 3.12 and 3.11 and Lemma 3.3. First, observe that we can assume without loss of generality that the sjfBCQ q is connected. This is because q has no self-join and the database D is Codd, so, letting q_1, \dots, q_t be the connected components of q , and letting D_i be the database D restricted to the relations appearing in q_i , we have that

$$\#\text{Val}_{Cd}^u(q)(D) = \prod_{i=1}^t \#\text{Val}_{Cd}^u(q_i)(D_i).$$

Second, notice that, for a connected sjfBCQ q , not containing any of these two patterns is equivalent to the following: There exists a variable x such that all atoms of q contain variable x , and for any two atoms of q , the only variable that they have in common is x . In other words, x is in every atom and every other variable occurs in only one atom. For instance $R_1(x, y, y) \wedge R_2(x, x, z, z, z, u, u) \wedge R_3(x, x, v, t, t)$ is such a query. We now provide an example of a connected query with only two atoms.

Example 3.14. We consider the query $q = R(x, x, y, y) \wedge S(x, x)$. Let D be an incomplete Codd database over relations R, S , with uniform domain dom of size d . In this proof, we will use symbols a, a_1, a_2, \dots to denote a constant or a null, and symbols c, c_1, c_2, \dots to denote constants. Moreover, unless stated otherwise, these symbols can refer to the same constant (but not to the same null, because D is a Codd database). A fact that contains only constants is called a *ground fact*. Furthermore, since that database is Codd, we will always represent a null with \perp , being understood that they are all distinct.

We start with a few simplifications. First, we assume w.l.o.g. that D does not contain ground facts that already satisfy the query. Second, we assume w.l.o.g. that D does not contain facts of the form $R(a, a', c, c')$ or $R(c, c', a, a')$ or $S(c, c')$ where c and c' are distinct constants. Indeed, because D is Codd and such a fact f can never be part of a match, we could simply remove f from D and multiply the end result by the appropriate value (namely, d^t where t is the number of nulls of f). Last, we assume w.l.o.g. that for any fact of the form $R(a, a', c, \perp)$ or $R(a, a', \perp, c)$ or $R(\perp, c, a, a')$ or $R(c, \perp, a, a')$ or $S(c, \perp)$ or $S(\perp, c)$, we have $c \in \text{dom}$; indeed otherwise, those facts can never be part of a match and we could again remove them and multiply the result by the appropriate value.

Next, we need to introduce some notation. For a fact f of D , the *type* of f is the word in $\{0, 1\}^{\text{arity}(f)}$ that has a 1 in position i iff the i th element of f is a null. For instance the type of $R(\perp, \perp, c, \perp)$ is 1, 101 and that of $S(c, c)$ is 00. Observe that there are a fixed number of possible types, because the query is fixed. For a constant c and fact f of D , we say that f is *c-determined* if f contains the constant c at a position for variable x . For instance $R(\perp, c, c', \perp)$ is *c-determined* and so is $S(c, c)$. A fact f that contains only nulls on the positions for x is called *free*. With the simplifications of the last paragraph, a fact of D is either free or it is *c-determined* for a unique constant c . Let \mathbf{t} be a type of D (for R or for S). We say that \mathbf{t} is *free* if it has only 1s in the positions corresponding to variable x ; in other words, if it is the type of a free fact. Otherwise, we say that \mathbf{t} is *determined*. For a constant c and determined type \mathbf{t} , we write $n_{R,c,\mathbf{t}}$ (respectively, $n_{S,c,\mathbf{t}}$) the number of R -facts (respectively, of S -facts) of D that are *c-determined* of type \mathbf{t} . For a free type \mathbf{t} of R (respectively, of S), we write $F_{R,\mathbf{t}}$ (respectively, $F_{S,\mathbf{t}}$) for the set of free R -facts (respectively, of free S -facts) and $f_{R,\mathbf{t}}$ (respectively, $f_{S,\mathbf{t}}$) for its size. Let f be a fact that is *c-determined*. We write α_f for the number of valuations ν of the nulls in f such that f matches the corresponding atom in q . For instance if f is $R(\perp, c, c', \perp)$ or again $R(c, c, c', c')$, then $\alpha_f = 1$, while if f is $R(c, c, \perp, \perp)$, then $\alpha_f = d$.⁸ Similarly, we let β_f denote the number of valuations ν of the nulls in f such that f does not match the corresponding atom in q ; which is then equal to $d^t - \alpha_f$, for t the number of nulls in f . Observe that α_f and β_f depend only on the type \mathbf{t} of f . Hence, we will write $\alpha_{\mathbf{t}}$ and $\beta_{\mathbf{t}}$ instead. Let f be a free fact. We let α_f be the number of valuations of the nulls in f that are not on a position for variable x that match the corresponding part of the atom in q . For instance if f is $R(\perp, \perp, c', \perp)$, then $\alpha_f = 1$ and if f is $R(\perp, \perp, \perp, \perp)$, then $\alpha_f = d$. We also define $\beta_f \stackrel{\text{def}}{=} d^t - d\alpha_f$ where t is the number of nulls, which correspond to the number of valuations of the nulls in f such that the ground fact obtained does not match its corresponding atom. Again, since α_f and β_f depend only on the free type \mathbf{t} of f , we will write $\alpha_{\mathbf{t}}$ and $\beta_{\mathbf{t}}$ instead. It is clear that we can compute all values $\alpha_{f,c}$ and $\beta_{f,c}$, for every determined fact f and constant c of D , and values α_f for every free fact in polynomial time. Last, we fix a linear order c_1, \dots, c_n on the constants c_i such that there exists a fact of D that is c_i -determined.

Next, we explain how we can compute in FP the number of valuations of D that *do not* satisfy q . To do so, we will first define some quantities and then show that we can compute these quantities in polynomial time using a dynamic programming approach. One of these quantities will be the number of valuations of D that do not satisfy q , which is what we want to compute. The quantities that we will define are of the form $V(\text{params})$, where *params* consist of the following parameters (\dagger):

- (a) one parameter v whose range is $0 \dots n$;
- (b) one parameter $q_{R,\mathbf{t}}$ for every free type $\mathbf{t} \in \{0, 1\}^4$ of R , with range $0 \dots f_{R,\mathbf{t}}$;
- (c) one parameter r_R with range $0, \dots, n$;
- (d) one parameter r_S with range $0, \dots, r_R$.

We now explain what the quantity $V(\text{params})$ with these parameters represent. To that end, we define the incomplete database $D(\text{params})$ to be the database that contains only

- all of the facts that are c_i -determined for $1 \leq i \leq v$ (if $v = 0$, then D_v contains no determined facts);
- for every free type \mathbf{t} of R , $D(\text{params})$ contains $q_{R,\mathbf{t}}$ free R -facts of D of type \mathbf{t} (it does not matter which ones, say the first $q_{R,\mathbf{t}}$ ones);
- $D(\text{params})$ contains all the free S -facts of D .

⁸When f is a ground fact—such as $R(c, c, c', c')$ —we recall the mathematical convention that there exists a unique function with empty domain, hence, a unique valuation of the nulls of f .

(Note that parameters (c) and (d) are not used to define the database but we still write $D(\text{params})$.)

The quantity $V(\text{params})$ is then defined to be the number of valuations ν of database $D(\text{params})$ such that $\nu(D(\text{params})) \not\models q$ and such that the following holds: (1) for every $1 \leq i \leq r_S$ and free fact f of R or of S , the ground fact $\nu(f)$ does not match the corresponding atom in q with variable x mapped to c_i (this condition is empty if $r_S = 0$); and (2) for every $r_S + 1 \leq i \leq r_R$ and every free fact f of R , the ground fact $\nu(f)$ does not match $R(x, x, y, y)$ with variable x mapped to c_i (this condition is empty if $r_R = r_S$). By definition, we then have that, when $\nu = n$, $q_{R,t} = f_{R,t}$ for every free type t of R —so $D(\text{params})$ is equal to D^- , $r_R = 0$ and $r_S = 0$ then $V(\text{params})$ is then equal to $\#\text{Val}_{Cd}^u(q)(D)$. Observe that there are a fixed number of parameters in params (because the query is fixed), and that the possible values of these parameters are polynomial in the size of the input. We explain how compute the values $V(\text{params})$ by dynamic programming.

Base case. Our base case will be when ν is equal to zero (and the other parameters are arbitrary as in (\dagger)); in other words, the database $D(\text{params})$ does not contain any determined facts, it contains only free facts. We have to compute the number of valuations of $D(\text{params})$ that do not satisfy the query and such that (1) and (2) hold. We do so as follows: We guess a subset S_R of $\text{dom} \setminus \{c_1, \dots, c_{r_R}\}$; this will be the set of constants c such that $R(c, c, c', c') \in \nu(D(\text{params}))$ for some c' . We also guess a subset S_S of $\text{dom} \setminus (\{c_1, \dots, c_{r_S}\} \cup S_R)$; this will be the set of constants c such that $S(c, c) \in \nu(D(\text{params}))$ (observe that S_R and S_S are disjoint; this is to avoid satisfying the query). To “achieve” these subsets, we for each free type t of R guess a subset $W_{R,t}$ of the free facts of R of type t ; these will be the R -facts f of type t such that $\nu(f)$ matches $R(x, x, y, y)$ with a constant in S_R for variable x . Similarly, we for each free type t of S guess a subset $W_{S,t}$ of the free facts of S of type t ; these will be the S -facts f of type t such that $\nu(f)$ matches $S(x, x)$ with a constant in S_S for variable x . To ensure that these subsets are indeed enough to cover S_R and S_S , we surjectively assign each of the corresponding facts a constant in S_R (for R -facts) or S_S (for S -facts). For such facts f , we then choose one of the α_t valuations of f that ensure that the ground fact obtained satisfies the corresponding atom of the query; crucially, this quantity depends only on the type of f . For the remaining (free) facts f , we choose one of the β_t valuations of f that ensure that the ground facts obtained do not satisfy the corresponding atom of the query. In the end, we obtain that $V(\text{params})$ is equal to the following expression:

$$\sum_{S_R \subseteq \text{dom} \setminus \{c_1, \dots, c_{r_R}\}} \sum_{S_S \subseteq \text{dom} \setminus (\{c_1, \dots, c_{r_S}\} \cup S_R)} \sum_{W_{R,t_1} \subseteq Q_{R,t_1}} \cdots \sum_{W_{R,t_j} \subseteq Q_{R,t_j}} \sum_{W_{S,t'_1} \subseteq F_{S,t'_1}} \cdots \sum_{W_{S,t'_k} \subseteq F_{S,t'_k}} \gamma, \quad (3)$$

where $t_1 \dots t_j$ are all the free types of R , $t'_1 \dots t'_k$ are all the free types of S , Q_{R,t_i} is the set of the first q_{R,t_i} free R -facts of type t_i , and where, letting $z \stackrel{\text{def}}{=} \sum_{i=1}^j |W_{R,t_i}|$ and $z' \stackrel{\text{def}}{=} \sum_{i=1}^k |W_{S,t'_i}|$, we have

$$\gamma = \text{surj}_{z \rightarrow |S_R|} \times \text{surj}_{z' \rightarrow |S_S|} \times \left(\prod_{i=1}^j \alpha_{t_i}^{|W_{R,t_i}|} \beta_{t_i}^{|Q_{R,t_i}| - |W_{R,t_i}|} \right) \times \left(\prod_{i=1}^k \alpha_{t'_i}^{|W_{S,t'_i}|} \beta_{t'_i}^{|F_{S,t'_i}| - |W_{S,t'_i}|} \right).$$

Obviously, we cannot compute the expression in Equation (3) in polynomial time, since we are summing over subsets of the input facts. However, because γ depends only on the sizes of all these sets, we can express $V(\text{params})$ as

$$\sum_{0 \leq s_R \leq d - r_R} \sum_{0 \leq s_S \leq d - r_S - s_R} \sum_{0 \leq w_{R,t_1} \leq q_{R,t_1}} \cdots \sum_{0 \leq w_{R,t_j} \leq q_{R,t_j}} \sum_{0 \leq w_{R,t'_1} \leq f_{R,t'_1}} \cdots \sum_{0 \leq w_{R,t'_k} \leq f_{R,t'_k}} \delta, \quad (4)$$

where, letting again $z \stackrel{\text{def}}{=} \sum_{i=1}^J w_{R,t_i}$ and $z' \stackrel{\text{def}}{=} \sum_{i=1}^K w_{S,t'_i}$, we have

$$\begin{aligned} \delta = & \text{surj}_{z \rightarrow s_R} \times \text{surj}_{z' \rightarrow s_S} \times \left(\prod_{i=1}^J \alpha_{t_i}^{w_{R,t_i}} \beta_{t_i}^{q_{R,t_i} - w_{R,t_i}} \right) \times \left(\prod_{i=1}^K \alpha_{t'_i}^{w_{S,t'_i}} \beta_{t'_i}^{f_{S,t'_i} - w_{S,t'_i}} \right) \\ & \times \binom{d - d_R}{s_R} \binom{d - r_S - s_R}{s_S} \times \left(\prod_{i=1}^J \binom{q_{R,t_i}}{w_{R,t_i}} \right) \times \left(\prod_{i=1}^K \binom{f_{S,t'_i}}{w_{S,t'_i}} \right). \end{aligned}$$

But then, because the expression (4) contains a fixed number of nested sums (this number depends only on the query q), and because indices and summands are polynomial, this quantity can be computed in polynomial time. This concludes the base case (i.e., when $v = 0$).

Inductive case. Next, we explain how we can compute the quantities $V(\text{params})$ (with the parameters params as in (†)) in polynomial time from quantities $V(\text{params}')$ with a strictly smaller value for parameter (a). Hence, we assume that parameter v is ≥ 1 . The idea is to get rid of the c_v -determined facts of $D(\text{params})$ by partitioning the valuations of $D(\text{params})$ that do not satisfy q and satisfy (1) and (2) into

- (A) those valuations ν that do not satisfy the query and satisfy (1) and (2) and such that no c_v -determined R -fact or free R -fact f of $D(\text{params})$ is such that $\nu(f)$ matches $R(x, x, y, y)$ with variable x mapped to c_v ; and
- (B) those valuations ν that do not satisfy the query and satisfy (1) and (2) and such that at least one c_v -determined R -fact or free R -fact f of $D(\text{params})$ is such that $\nu(f)$ matches $R(x, x, y, y)$ with variable x mapped to c_v (and thus, no c_v -determined S -fact or free S -fact f of $D(\text{params})$ is such that $\nu(f)$ matches $S(x, x)$ with variable x mapped to c_v).

To compute valuations in (A), we choose for each R -fact f that is c_v -determined one of the β_{f,c_v} valuations of its nulls that is such that $\nu(f)$ does not match $R(x, x, y, y)$, we disallow the free facts of R to have valuations that would match on c_v by increasing r^R by one, and we choose any valuation for the c_v -determined facts of S (since we know that they will not be part of a match). Formally, letting t_f be the number of nulls of a fact f , we have the expression

$$A = \left(\prod_{\substack{c_v\text{-determined} \\ R\text{-fact } f}} \beta_f \right) \times \left(\prod_{\substack{c_v\text{-determined} \\ S\text{-fact } f}} d^{t_f} \right) \times V(\text{params}'),$$

where params' is equal to params except that $v' = v - 1$ and $r'_R = r_R + 1$. This can be computed in polynomial time if we know the value $V(\text{params}')$. For valuations in (B), we do as follows: We choose exactly which subset of the c_v -determined and free facts of R will match $R(x, x, y, y)$ with $x = c_v$ by partitioning according to the types; for the remaining c_v -determined R -facts, we choose one of the β valuations that do not match on c_v ; for the remaining free facts of R , we disallow to match on c_v by increasing r^R by one; for each c_v -determined fact f of S , we choose one of the β_{f,c_v} valuations of f that are such that $\nu(f)$ is not $S(c_v, c_v)$; and for the free facts of S , we disallow matching on c_v by increasing r_S by one. Formally, letting $\mathbf{t}_1, \dots, \mathbf{t}_j$ be the types of the c_v -determined R -facts and $\mathbf{t}'_1, \dots, \mathbf{t}'_k$ be the free types of R , we obtain an expression of the form

$$B = \sum_{0 \leq h_1 \leq n_{R,c_v,t_1}} \cdots \sum_{0 \leq h_j \leq n_{R,c_v,t_j}} \sum_{0 \leq h'_1 \leq q_{R,t'_1}} \cdots \sum_{0 \leq h'_k \leq q_{R,t'_k}} \gamma \times \delta \times V(\text{params}'), \quad (5)$$

where γ is equal to 1 if $\sum_{i=1}^J h_{t_i} + \sum_{i=1}^K h_{t'_i} \geq 1$ and to 0 otherwise (to match on c_v in R), δ is

$$\left(\prod_{i=1}^J \binom{n_{R,c_v,t_i}}{h_{t_i}} \alpha_{t_i}^{h_{t_i}} \beta_{t_i}^{n_{R,c_v,t_i} - h_{t_i}} \right) \times \left(\prod_{i=1}^J \binom{q_{R,t'_i}}{h'_{t'_i}} \alpha_{t'_i}^{h'_{t'_i}} \right) \times \left(\prod_{\substack{c_v\text{-determined} \\ S\text{-fact } f}} \beta_f \right),$$

and where params' is equal to params except that:

- we have $v' = v - 1$;
- for every free type t of R , we have $q'_{R,t} = q_{R,t} - h_{t_i}$;
- we have $r'_R = r_R + 1$;
- we have $r'_S = r_S + 1$.

The crucial point to see that expressions A and B indeed compute what we want is that we do not need to remember exactly which subsets of constants are disallowed to match for the free facts, but we only need to remember their numbers (this is what allows us to use a dynamic programming approach); and similarly for the free facts of R , we only need to remember how many we have left at each stage of each type, but not their precise subsets. Again, if we know the values of $V(\text{params}')$ where in params' parameter v' is equal to $v - 1$, then we can compute in polynomial time the value $V(\text{params}) = A + B$. Thus, we can compute all values $V(\text{params})$ in polynomial time, and this concludes this example.

The proof in this example can be extended as-is to any connected sjfBCQ containing only two atoms and having only one variable (x) that joins. We claim that the same idea works for an arbitrary number of atoms, which concludes Theorem 3.13. Since a full proof is technically tedious and does not provide new insights, we omit it here.

4 DICHOTOMIES FOR COUNTING COMPLETIONS

In this section, we study the complexity of the problems of counting completions satisfying an sjfBCQ q , in the four cases that can be obtained by considering naive or Codd tables and non-uniform or uniform domains. We will again use the notion of pattern as introduced in Definition 3.1. Our first step is to observe that Lemma 3.3, which we used in the last section for the problems or counting valuations, extends to the problems of counting completions.

LEMMA 4.1. *Let q, q' be sjfBCQs such that q' is a pattern of q . Then, we have that $\#\text{Comp}(q') \leq_{\text{par}}^P \#\text{Comp}(q)$. Moreover, the same results hold if we restrict to the case of Codd tables, and/or to the uniform setting.*

PROOF. The reduction is exactly the same as the one of Lemma 3.3. To show that this reduction works properly for counting completions, it is enough to observe that for every pair of valuations v_1, v_2 of D' (or of D , since D and D' have exactly the same set of nulls), we have that $v_1(D') = v_2(D')$ iff $v_1(D) = v_2(D)$. \square

We will then follow the same general strategy as in the last section, i.e., prove hardness for some simple patterns and combine these with Lemma 4.1 and tractability proofs to obtain dichotomies. Our findings are summarized in the last two columns of Table 1 in the introduction. We start in Section 4.1 with the non-uniform cases and continue in Section 4.2 with the uniform cases. Again, we explicitly state when a #P-hardness result holds even in the restricted setting in which there is a fixed domain over which nulls are interpreted.

4.1 The Complexity on the Non-uniform Case

Here, we study the complexity of the problems $\#\text{Comp}(q)$ and $\#\text{Comp}_{Cd}(q)$, providing dichotomy results in both cases. In fact, it turns out that these problems are $\#\text{P}$ -hard for all sjfBCQs. To prove this, it is enough to show that the problem $\#\text{Comp}_{Cd}(R(x))$ is hard, that is, even counting the completions of a single unary table is $\#\text{P}$ -hard in the non-uniform setting.

PROPOSITION 4.2. $\#\text{Comp}_{Cd}(R(x))$ is $\#\text{P}$ -hard.

PROOF. We provide a polynomial-time parsimonious reduction from the problem of counting the vertex covers of a graph, which we denote by $\#\text{VC}$. Let $G = (V, E)$ be a graph. We construct a Codd table D using a single unary relation R such that the number of completions of D equals the number of vertex covers of G . For every edge $e = \{u, v\}$ of G , we have one null \perp_e with $\text{dom}(\perp_e) = \{u, v\}$ and the fact $R(\perp_e)$. Let a be a fresh constant. For every node $u \in V$, we have a null \perp_u with $\text{dom}(\perp_u) = \{u, a\}$ and the fact $R(\perp_u)$. Last, we add the fact $R(a)$. We now show that the number of completions of D equals the number of vertex covers of G .

Let $\text{VC}(G)$ be the set of vertex covers of G . For a valuation ν of D , define the set $S_\nu := \{u \in V \mid R(u) \in D\}$. Since the fact $R(a)$ is in every completion of D , it is clear that the number of completions of D is equal to $|\{S_\nu \mid \nu \text{ is a valuation of } D\}|$. We claim that $\text{VC}(G) = \{S_\nu \mid \nu \text{ is a valuation of } D\}$, which shows that the reduction works. (\subseteq) Let $C \in \text{VC}(G)$, and let us show that there exists a valuation ν of D such that $S_\nu = C$. For a null of the form \perp_e with $e = \{u, v\} \in E$, assuming w.l.o.g. that $u \in C$, we define $\nu(\perp_e)$ to be u . For a null of the form \perp_u with $u \in V$, we define $\nu(\perp_u)$ to be u if $u \in C$ and a otherwise. It is then clear that $S_\nu = C$. (\supseteq) Let ν be a valuation of D , and let us show that S_ν is a vertex cover. Assume by contradiction that there is an edge $e = \{u, v\}$ such that $e \cap S_\nu = \emptyset$. By definition of D , we must have $\nu(\perp_e) \in \{u, v\}$, so one of u or v must be in S_ν , hence, a contradiction. Therefore, we conclude that $\#\text{VC} \leq_{\text{par}}^{\text{P}} \#\text{Comp}_{Cd}(R(x))$. \square

Recall from Section 2 that, to avoid trivialities, we assume all sjfBCQs to contain at least one atom and that all atoms have at least one variable. Using Lemma 4.1, this allows us to obtain the following dichotomy result:

THEOREM 4.3 (DICHOTOMY). *For every sjfBCQ q , it holds that $\#\text{Comp}(q)$ and $\#\text{Comp}_{Cd}(q)$ are $\#\text{P}$ -hard.*

Notice here that we do not claim membership in $\#\text{P}$; in fact, we will come back to this issue in Section 6 to show that this is unlikely to be true for naive tables. However, we can still show that membership in $\#\text{P}$ holds for Codd tables. We then obtain:

THEOREM 4.4 (DICHOTOMY). *For every sjfBCQ q , the problem $\#\text{Comp}_{Cd}(q)$ is $\#\text{P}$ -complete.*

PROOF. Hardness is from Theorem 4.3. To show membership in $\#\text{P}$ we will actually prove a more general result in Section 6.1. There, we show that for every Boolean query q such that q has model checking in P the problem $\#\text{Comp}_{Cd}(q)$ is in $\#\text{P}$. This in particular applies to all sjfBCQs. \square

4.2 The Complexity on the Uniform Case

We now investigate the complexity of $\#\text{Comp}^u(q)$ and $\#\text{Comp}_{Cd}^u(q)$. Recall that in the non-uniform case, even counting the completions of a single unary table is a $\#\text{P}$ -hard problem. This no longer holds in the uniform case, as we will show that $\#\text{Comp}^u(q)$ is in FP for every sjfBCQ that is defined over a schema consisting exclusively of unary relation symbols.

Such a positive result, however, cannot be extended much further. In fact, we show next that $R(x, x)$ and $R(x, y)$ are hard patterns, both for naive and Codd tables (and, thus, we also conclude that the problem of counting the completions of a single binary Codd table is a $\#\text{P}$ -hard

problem). We start with the case of naive tables, for which hardness even holds when nulls are interpreted over the fixed domain $\{0, 1\}$.

PROPOSITION 4.5. *The problems $\#\text{Comp}^u(R(x, x))$ and $\#\text{Comp}^u(R(x, y))$ are both $\#\text{P}$ -hard, even when nulls are interpreted over the same fixed domain $\{0, 1\}$.*

PROOF. We reduce from $\#\text{IS}$, the problem of counting the number of independent sets of a graph. Let $G = (V, E)$ be a graph. We will construct an incomplete database D containing a single binary predicate R such that each completion of D satisfies $R(x, x)$ and the number of completions of D is $2^{|V|} + \#\text{IS}(G)$, thus establishing hardness for the two queries. For every node $u \in V$, we have a null \perp_u with $\text{dom}(\perp_u) = \{0, 1\}$. We then construct the naive table D as follows:

- for every node $u \in V$, we add to D the fact $R(u, \perp_u)$;
- then for every edge $\{u, v\} \in E$, we add the facts $R(\perp_u, \perp_v)$ and $R(\perp_v, \perp_u)$ to D ; and
- last, we add the facts $R(0, 0)$, $R(0, 1)$, $R(1, 0)$, and $R(\perp, \perp)$, where \perp is a fresh null.

It is clear that every completion of D satisfies $R(x, x)$ (thanks to the fact $R(\perp, \perp)$). Let us now count the number of completions of D . First, we observe that, thanks to the facts of the form $R(u, \perp_u)$, for $u \in V$, for every two valuations ν, ν' that do not assign the same value to the nulls of the form \perp_u , it is the case that $\nu(D) \neq \nu'(D)$. We then partition the completions of D into those that contain the fact $R(1, 1)$ and those that do not contain $R(1, 1)$. Because of the facts of the form $R(u, \perp_u)$, for $u \in V$, and thanks to the fact $R(\perp, \perp)$, which becomes $R(1, 1)$ when we assign 1 to \perp , there are exactly $2^{|V|}$ completions of D that contain $R(1, 1)$. Moreover, it is easy to see that there are $\#\text{IS}(G)$ valuations ν of D that assign 0 to \perp and that yield a completion not containing $R(1, 1)$. Indeed, one can check that a valuation of D that assigns 0 to \perp yields a completion not containing $R(1, 1)$ if and only if the set $\{u \in V \mid \nu(\perp_u) = 1\}$ is an independent set of G . Therefore, we conclude that the number of completions of D is indeed $2^{|V|} + \#\text{IS}(G)$, and therefore that $\#\text{IS} \leq_{\text{T}}^{\text{P}} \#\text{Comp}^u(q)$, where q can be $R(x, x)$ or $R(x, y)$. \square

Next, we prove hardness of the same queries for Codd tables (but in this case we do not know if hardness holds when nulls are interpreted over a fixed domain, as our proof will use domains of unbounded size). We will reduce from the problem of counting the number of *induced pseudoforests* of a graph, as defined next.

Definition 4.6. A graph G is a *pseudoforest* if every connected component of G contains at most one cycle. Let $G = (V, E)$ be a graph. For $S \subseteq E$, let us denote by $G[S]$ the graph (V', S) , where V' is the set of nodes of G that appear in some edge of S . The problem $\#\text{PF}$ is the problem that takes as input a graph $G = (V, E)$ and outputs the number of edge sets $S \subseteq E$ such that $G[S]$ is a pseudoforest.

Using techniques from matroid theory, the authors of Reference [26] have shown that $\#\text{PF}$ is $\#\text{P}$ -hard on graphs. We explain in Appendix B.1 how their proof actually shows hardness of this problem for *bipartite graphs* (which we need); formally, we have:

PROPOSITION 4.7 (IMPLIED BY [26]). *The problem $\#\text{PF}$ restricted to bipartite graphs is $\#\text{P}$ -hard.*

To prove that the reduction that we will present is correct, we will also need the following folklore lemma about pseudoforests. We recall that an *orientation* of an undirected graph $G = (V, E)$ is a directed graph that can be obtained from G by orienting every edge of G . Equivalently, one can see such an orientation as a function $f : E \rightarrow V$ that assigns to every edge in G a node to which it is incident. We then have:

LEMMA 4.8. *A graph G is a pseudoforest if and only if there exists an orientation of G such that every node has outdegree at most 1.*

PROOF. Folklore, see, e.g., References [22, 27, 36]. \square

Using the hardness of #PF on bipartite graphs, we are able to show hardness of $\#\text{Comp}_{Cd}^u(R(x, x))$ and $\#\text{Comp}_{Cd}^u(R(x, y))$ as follows:

PROPOSITION 4.9. *The problems $\#\text{Comp}_{Cd}^u(R(x, x))$ and $\#\text{Comp}_{Cd}^u(R(x, y))$ are both #P-hard.*

PROOF. We reduce both problems from #PF on bipartite graphs. Let $G = (U \sqcup V, E)$ be a bipartite graph. We will construct a uniform Codd table D over binary relation R such that (1) all the completions of D satisfy both queries; and (2) the number of completions of D is equal to #PF(G), thus establishing hardness. For every $(t, t') \in (U \cup V)^2 \setminus E$, we add to D the fact $R(t, t')$; we call these the *complementary facts*. For every $u \in U$, we add to D the fact $R(u, \perp_u)$ and for every $v \in V$ the fact $R(\perp_v, v)$. Finally, we add to D a fact $R(f, f)$ where f is a fresh constant. The uniform domain of the nulls is $\text{dom} = U \cup V$. It is clear that D is a Codd table and that every completion of D satisfies both queries (thanks to the fact $R(f, f)$), so (1) holds. We now prove that (2) holds. First, observe that a completion $\nu(D)$ of D is uniquely determined by the set of edges $\{(u, v) \in E \mid R(u, v) \in \nu(D)\}$; this is because $\nu(D)$ already contains all the complementary facts. For a set $S \subseteq E$ of edges, let us define D_S to be the complete database that contains all the complementary facts and all the facts $R(u, v)$ for $(u, v) \in S$ (note that D_S is not necessarily a completion of D). We now argue that for every set $S \subseteq E$, we have that D_S is a completion of D if and only if $G[S]$ is a pseudoforest, which would conclude the proof. By Lemma 4.8, we only need to show that D_S is a completion of D if and only if $G[S]$ admits an orientation with maximum outdegree 1. We show each direction in turn. (\Rightarrow) Assume D_S is a completion of D , and let ν be a valuation witnessing this fact, i.e., such that $\nu(D) = D_S$. First, observe that we can assume without loss of generality that (\star) for every $e = (u, v) \in S$, we have either $\nu(\perp_u) = v$ or $\nu(\perp_v) = u$ but not both. Indeed, if we had both, then we could modify ν into ν' by redefining, say, $\nu'(\perp_u)$ to be u , and we would still have that $\nu'(D) = D_S$ (because $R(u, u)$ is already present in D : It is a complementary fact). We now define an orientation $f_\nu : S \rightarrow U \cup V$ of $G[S]$ from ν as follows: Let $e = (u, v) \in S$. Then: if we have $\nu(\perp_u) = v$, we define $f_\nu((u, v))$ to be v , i.e., we orient the (undirected) edge (u, v) from u to v . Else, if we have $\nu(\perp_v) = u$, then we define $f_\nu((u, v))$ to be u , i.e., we orient the (undirected) edge (u, v) from v to u . Observe that by (\star) f_ν is well defined. It is then easy to check that the maximal outdegree of the directed graph defined by f_ν is 1; this is because for every $u \in U$ (respectively, $v \in V$), there is only one fact in D of the form $R(u, \text{null})$ (respectively, $R(\text{null}, v)$), namely, the fact $R(u, \perp_u)$ (respectively, $R(\perp_v, v)$). (\Leftarrow) Let $f : S \rightarrow U \cup V$ be an orientation of $G[S]$ with maximum outdegree 1. Let ν_f be the valuation of D defined from f as follows: For every $u \in U$ (respectively, $v \in V$), if there is an edge $(u, v) \in S$ such that $f((u, v)) = v$ (respectively, such that $f((u, v)) = u$), then define $\nu_f(\perp_u)$ to be v (respectively, define $\nu_f(\perp_v)$ to be u). Observe that there can be at most one such edge, because f has maximum outdegree 1, so this is well defined. If there is no such edge, define $\nu_f(\perp_u)$ to be u (respectively, define $\nu_f(\perp_v)$ to be v). Since all edges in S are given an orientation by f , it is clear that for every $(u, v) \in S$, we have $R(u, v) \in \nu_f(D)$. Moreover, since $\nu_f(D)$ contains all the complementary facts, we have that $\nu_f(D) = D_S$, which shows that D_S is a completion of D and concludes this proof. \square

As we show next, these two patterns suffice to characterize the complexity of $\#\text{Comp}^u(q)$ and $\#\text{Comp}_{Cd}^u(q)$.

THEOREM 4.10 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x, y)$ is a pattern of q , then $\#\text{Comp}^u(q)$ and $\#\text{Comp}_{Cd}^u(q)$ are #P-hard. Otherwise, these problems are in FP.*

From what precedes, we only have to prove the tractability part of that claim, and this for naive tables. To this end, we define a *conjunction of basic singletons* sjfBCQ to be an sjfBCQ of

the form $C_1(x_1) \wedge \dots \wedge C_m(x_m)$, where each $C_i(x_i)$ is a conjunction of unary atoms over the same variable x_i (here, the x_i are pairwise distinct). Since q does not contain the pattern $R(x, x)$ nor the pattern $R(x, y)$, observe that q must in fact be a conjunction of basic singletons sjfBCQ. The main difficulty is to decompose the computation in such a way that we do not count the same completion twice. Moreover, the fact that the database is naive and not Codd, and the fact that constants can appear everywhere, complicate a lot of the description of the algorithm. For these reasons, we provide in the next section an example that gives the intuition of the general proof, and defer the general proof to Appendix B.2.

Note that, as in the last section, we do not claim membership in #P in the statement of Theorem 4.10. However, and also as in the last section, we can show that these problems are in #P for Codd tables, which allows us to obtain our last dichotomy for exact counting.

THEOREM 4.11 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x, y)$ is a pattern of q , then $\#Comp_{Cd}^u(q)$ is #P-complete. Otherwise, this problem is in FP.*

PROOF. Hardness follows from Theorem 4.10, while membership in #P follows from the result proven in Section 6.1. \square

4.3 Example of Tractability for Theorem 4.10

In this section, we prove that for the query $q \stackrel{\text{def}}{=} R(x) \wedge S(y)$, the problem $\#Comp_{Cd}^u(q)$ is in FP. Observe that q is a conjunction of basic singletons query and that it is always satisfied (as long as the database is not empty). We will show that $\#Comp^u(q)$ is in FP. Let C_{RS}, C_R, C_S (respectively, N_{RS}, N_R, N_S) be the sets of constants (respectively, nulls) that occur respectively: in R and in S , only in R , only in S , and denote c_{RS}, c_R, c_S (respectively, n_{RS}, n_R, n_S) their sizes. Let dom be the uniform domain of the nulls, and let d be its size. First, observe that we can assume without loss of generality that $C \stackrel{\text{def}}{=} C_{RS} \cup C_R \cup C_S \subseteq \text{dom}$, as, otherwise, we could simply remove the facts of D that are over constants that are not in dom and this would not change the result. Let $c = c_{RS} + c_R + c_S$. The next two claims explain how we can decompose the computation in a way that we do not count a same completion twice (Claim 4.12) and that we count all the possible completions (Claim 4.13).

CLAIM 4.12. *For a triplet (I_R, I_S, I_{RS}) of subsets of dom satisfying the conditions (\star) $I_R \subseteq \text{dom} \setminus C$, $I_S \subseteq \text{dom} \setminus (C \cup I_R)$, and $I_{RS} \subseteq \text{dom} \setminus (C_{RS} \cup I_R \cup I_S)$, let us define $P(I_R, I_S, I_{RS})$ to be the complete database consisting of the following facts:*

- (1) $R(a)$ and $S(a)$ for $a \in C_{RS} \cup I_{RS}$;
- (2) $R(a)$ for $a \in I_R \cup (C_R \setminus I_{RS})$;
- (3) $S(a)$ for $a \in I_S \cup (C_S \setminus I_{RS})$.

Then, for any two such triplets of sets (I_R, I_S, I_{RS}) and (I'_R, I'_S, I'_{RS}) that are different, the complete databases $P(I_R, I_S, I_{RS})$ and $P(I'_R, I'_S, I'_{RS})$ are distinct.

PROOF. To help the reader, we have drawn in Figure 2 how the sets can intersect. If we have $I_{RS} \neq I'_{RS}$ with $a \in I_{RS}$ and $a \notin I'_{RS}$, then $P(I_R, I_S, I_{RS})$ contains both facts $R(a)$ and $S(a)$, while $P(I'_R, I'_S, I'_{RS})$ does not. So let us assume now that $I_{RS} = I'_{RS}$. If we have $I_R \neq I'_R$ with $a \in I_R$ and $a \notin I'_R$, then one can check that $P(I_R, I_S, I_{RS})$ contains the fact $R(a)$ while $P(I'_R, I'_S, I'_{RS})$ does not. Hence, let us assume that $I_R = I'_R$. Using the same reasoning, we obtain that $I_S = I'_S$, thus completing the proof. \square

Our next step is to show that every completion of D is of the form $P(I_R, I_S, I_{RS})$ for some triplet (I_R, I_S, I_{RS}) satisfying (\star) :




Fig. 2. How the sets dom , I_R , I_S , I_{RS} , C_{RS} , C_R and C_S from Claim 4.12 are allowed to intersect when they satisfy (\star) . The sets themselves and the intersections can be empty.

CLAIM 4.13. *For every completion D' of D , there exists a triplet (I_R, I_S, I_{RS}) satisfying (\star) such that $D' = P(I_R, I_S, I_{RS})$.*

PROOF. We define:

- $I_R \stackrel{\text{def}}{=} D'(R) \setminus (C_R \cup D'(S))$, where we see $D'(R)$ as the set of constants occurring in relation R of D' ;
- $I_S \stackrel{\text{def}}{=} D'(S) \setminus (C_S \cup D'(R))$;
- $I_{RS} \stackrel{\text{def}}{=} (D'(R) \cap D'(S)) \setminus C_{RS}$.

Then one can easily check that (I_R, I_S, I_{RS}) satisfies (\star) and that $D' = P(I_R, I_S, I_{RS})$. \square

By combining these two claims, we have that the result that we wish to compute (which, we recall, is simply the total number of completions of D , because any valuation satisfies the query) is equal to

$$\sum_{I_R \subseteq \text{dom} \setminus C} \sum_{I_S \subseteq \text{dom} \setminus (C \cup I_R)} \sum_{I_{RS} \subseteq \text{dom} \setminus (C_{RS} \cup I_R \cup I_S)} \text{check}(I_R, I_S, I_{RS}),$$

where $\text{check}(I_R, I_S, I_{RS}) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } P(I_R, I_S, I_{RS}) \text{ is a possible completion of } D \\ 0 & \text{otherwise} \end{cases}$.

Next, we show that the value of $\text{check}(I_R, I_S, I_{RS})$ can be computed in polynomial time and actually only depends on the sizes of these sets. To show this, we will use the following:

CLAIM 4.14. *We have $\text{check}(I_R, I_S, I_{RS}) = 1$ if and only if the following conditions hold:*

- (1) *If $n_R \geq 1$ and $|C_R \cup C_{RS} \cup I_{RS}| = 0$, then we have $|I_R| \neq 0$. Intuitively, this means that the value of a null in N_R cannot be “absorbed” by $C_R \cup C_{RS} \cup I_{RS}$.*
- (2) *If $n_S \geq 1$ and $|C_S \cup C_{RS} \cup I_{RS}| = 0$, then we have $|I_S| \neq 0$. (Same reasoning for nulls in N_S .)*
- (3) *If $n_{RS} \geq 1$ and $|C_{RS}| = 0$, then we have $|I_{RS}| \neq 0$. (Same reasoning for nulls in N_{RS} .)*

(4) The following system of equations, whose variables are natural numbers between 0 and d , has a solution:

$$\begin{aligned} z_{N_R}^{\{N_R\}} + z_{N_R}^{\{N_R, C_S\}} + z_{N_R}^{\{N_R, N_S\}} &\leq n_R \\ z_{N_S}^{\{N_S\}} + z_{N_S}^{\{N_S, C_R\}} + z_{N_S}^{\{N_S, N_R\}} &\leq n_S \\ z_{C_R}^{\{C_R, N_S\}} &\leq c_R \\ z_{C_S}^{\{C_S, N_R\}} &\leq c_S \\ z_{N_R}^{\{N_R\}} &\geq |I_R| \\ z_{N_S}^{\{N_S\}} &\geq |I_S| \end{aligned}$$

$$n_{RS} + \min(z_{N_R}^{\{N_R, C_S\}}, z_{C_S}^{\{C_S, N_R\}}) + \min(z_{N_R}^{\{N_R, N_S\}}, z_{N_S}^{\{N_S, N_R\}}) + \min(z_{N_S}^{\{N_S, C_R\}}, z_{C_R}^{\{C_R, N_S\}}) \geq |I_{RS}|.$$

PROOF. We prove the claim informally by explaining the main ideas, because a formal proof would be too long and not that interesting. Conditions (1–3) are easily checked to be necessary. We now explain why condition (4) is also necessary. Suppose that $P(I_R, I_S, I_{RS})$ is a completion of D . Observe that (\dagger) to obtain the constants in I_{RS} , we had to use some or all of the following:

- the nulls in N_{RS} ; or
- the nulls in N_R together with those in N_S ; or
- the nulls in N_R together with the constants in C_S ; or
- the nulls in N_S together with the constants in C_R .

But then, to obtain $P(I_R, I_S, I_{RS})$ as a completion, we must have used three disjoint (possibly empty) sets $Z_{N_R}^{\{N_R\}}, Z_{N_R}^{\{N_R, C_S\}}, Z_{N_R}^{\{N_R, N_S\}}$ of the nulls in N_R of sizes $0 \leq z_{N_R}^{\{N_R\}}, z_{N_R}^{\{N_R, C_S\}}, z_{N_R}^{\{N_R, N_S\}} \leq d$, we have done the same for the nulls in N_S and we also used a subset of the constants of C_R (and C_S) in such a way that, according to (\dagger) :

- The nulls in $Z_{N_R}^{\{N_R\}}$ have been used to obtain the set I_R (which, we recall, is the set of constants that occur only in R and that are not in C_R). Note that only the nulls in N_R could have been used to obtain constants in I_R . This is what the fifth equation expresses.
- The nulls in $Z_{N_R}^{\{N_R, C_S\}}$ have values in $Z_{C_S}^{\{C_S, N_R\}}$, which gives us constants in I_{RS} . Observe that at maximum, we could obtain $\min(z_{N_R}^{\{N_R, C_S\}}, z_{C_S}^{\{C_S, N_R\}})$ constants in this manner.
- The nulls in $Z_{N_R}^{\{N_R, N_S\}}$ and those in $Z_{N_S}^{\{N_S, N_R\}}$ have common values, which gives us constants in I_{RS} . Again, observe that we can get at most $\min(z_{N_R}^{\{N_R, N_S\}}, z_{N_S}^{\{N_S, N_R\}})$ constants using these.
- The nulls in $Z_{N_S}^{\{N_S, C_R\}}$ have values in $Z_{C_R}^{\{C_R, N_S\}}$, which gives us constants in I_{RS} . Observe that at maximum, we could obtain $\min(z_{N_S}^{\{N_S, C_R\}}, z_{C_R}^{\{C_R, N_S\}})$ constants in this manner.

The first four equations express the partitioning process, and the last equation then expresses that by combining all these constants, we indeed obtained the whole set I_{RS} .

We now explain why conditions (2–4) are sufficient. If $|I_R|, |I_S|$ and I_{RS} are all ≥ 1 , then condition (4) is sufficient, because we can use the nulls and constants as explained above, and we have enough of them to obtain the sets I_R, I_S, I_{RS} . We explain what happens when $I_R = \emptyset$ for instance. In that case, condition (1) ensures us that we have either $n_R = 0$ or $C_R \cup C_{RS} \cup I_{RS} \neq \emptyset$. If we have $n_R = 0$, then it is fine, since the only nulls that could be used to fill I_R are those in N_R . If we have $n_R \geq 1$ and $C_R \cup C_{RS} \cup I_{RS} \neq \emptyset$, then we can use these to absorb the values of the nulls in N_R , and we are fine (i.e., we will be able to obtain $I_R = \emptyset$). We leave it to the reader to complete the small gaps in this proof. \square

Using this, we have that the value of $\text{check}(I_R, I_S, I_{RS})$ only depends on the sizes of I_R, I_S, I_{RS} , and moreover can be computed in polynomial time

CLAIM 4.15. *The value of $\text{check}(I_R, I_S, I_{RS})$ only depends on $|I_R|, |I_S|, |I_{RS}|, n_R, n_S, n_{RS}, c_R, c_S, c_{RS}$, and can be computed in FP.*

PROOF. The fact that this value only depends on the sizes of these sets is simply by inspection of the conditions in Claim 4.14. Conditions (1–3) can obviously be checked in PTIME. The fact that condition (4) can be checked in PTIME is because we can test all possible assignments between 0 and d for all these variables and see if there is one assignment that satisfies the equations; indeed, observe that the number of variables is fixed because the query is fixed. \square

But then, we can express the result as follows:

$$\sum_{0 \leq i_R, i_S, i_{RS} \leq d} \binom{d-c}{i_R} \binom{d-c-i_R}{i_S} \binom{d-c_{RS}-i_R-i_S}{i_{RS}} \times \text{check}(i_R, i_S, i_{RS}),$$

and we can evaluate this expression as-is in FP, because computing $\text{check}(I_R, I_S, I_{RS}) \in \{0, 1\}$ is in PTIME by the last claim. This concludes this (long) example.

5 APPROXIMATING THE NUMBERS OF VALUATIONS AND COMPLETIONS

As we saw in the previous sections, counting valuations and completions of an incomplete database are usually intractable problems. However, this does not necessarily rule out the existence of efficient approximation algorithms for such counting problems, in particular if some source of randomization is allowed. In this section, we investigate this question by focusing on the well-known notion of **Fully Polynomial-time Randomized Approximation Scheme (FPRAS)** for counting problems [33]. Formally, let Σ be a finite alphabet and $f : \Sigma^* \rightarrow \mathbb{N}$ be a counting problem. Then f is said to have an FPRAS if there is a randomized algorithm $\mathcal{A} : \Sigma^* \times (0, 1) \rightarrow \mathbb{N}$ and a polynomial $p(u, v)$ such that, given $x \in \Sigma^*$ and $\varepsilon \in (0, 1)$, algorithm \mathcal{A} runs in time $p(|x|, 1/\varepsilon)$ and satisfies the following condition:

$$\Pr\left(|f(x) - \mathcal{A}(x, \varepsilon)| \leq \varepsilon f(x)\right) \geq \frac{3}{4}.$$

Observe that the property of having an FPRAS is closed under polynomial-time parsimonious reductions, that is, if we have an FPRAS for a counting problem A and for counting problem B , we have that $B \leq_{\text{par}}^p A$, then we also have an FPRAS for B .

In the following sections, we investigate the existence of FPRAS for the problems of counting valuations and completions of an incomplete database. The overall picture that we obtain is shown in Table 2. We first deal with counting valuations in Section 5.1, where we show a general condition under which this problem has an FPRAS (which will apply, in particular, to all Boolean conjunctive queries). Then, in Section 5.2, we show that the situation is quite different for counting completions, as in most cases this problem does not admit an FPRAS.

5.1 Approximating the Number of Valuations

To prove the main result of this section, we need to consider the counting complexity class SpanL [5]. Given a finite alphabet Σ , an NL-transducer M over Σ is a nondeterministic Turing Machine with input and output alphabet Σ , a read-only input tape, a write-only output tape (where the head is always moved to the right once a symbol in Σ is written on it, so the output cannot be read by M), and a work-tape of which, on input x , only the first $c \cdot \log(|x|)$ cells can be used for a fixed constant $c > 0$ (so the space used by M is logarithmic). Moreover, $y \in \Sigma^*$ is said to be an output of M with input x if there exists an accepting run of M with input x such that y is the string

in the output tape when M halts. Then a function $f : \Sigma^* \rightarrow \mathbb{N}$ is said to be in SpanL if there exists an NL-transducer M over Σ such that for every $x \in \Sigma^*$, the value $f(x)$ is equal to the number of distinct outputs of M with input x . In Reference [5], it was proved that $\text{SpanL} \subseteq \#\text{P}$, and also this inclusion is strict unless $\text{NL} = \text{NP}$.

Very recently, the authors of Reference [10] have shown that every problem in SpanL has an FPRAS.

THEOREM 5.1 ([10, COROLLARY 3]). *Every problem in SpanL has an FPRAS.*

By using this result, we can give a general condition on a Boolean query q under which $\#\text{Val}(q)$ has an FPRAS, as this condition ensures that $\#\text{Val}(q)$ is in SpanL. More precisely, a Boolean query q is said to be *monotone* if for every pair of (complete) databases D, D' such that $D \subseteq D'$, if $D \models q$, then $D' \models q$. Moreover, q is said to have *bounded minimal models* if there exists a constant C_q (that depends only on q) satisfying that for every (complete) database D , if $D \models q$, then there exists $D' \subseteq D$ such that $D' \models q$ and the number of facts in D' is at most C_q . Finally, the *model checking problem* for q , denoted by $\text{MC}(q)$, is the problem of deciding, given a (complete) database D , whether $D \models q$. Then q is said to have a model checking in a complexity class C if $\text{MC}(q) \in C$. With this terminology, we can state the main result of this section.

PROPOSITION 5.2. *Assume that a Boolean query q is monotone, has model checking in nondeterministic linear space, and has bounded minimal models. Then $\#\text{Val}(q)$ is in SpanL.*

PROOF. Let D be the input incomplete database, with the domains for each null. First, the machine guesses a subset $D' \subseteq D$ of size $\leq C_q$, such that each fact of D' is over a relation symbol that appears in q . Observe that D' contains at most $|D'| \times \text{arity}(q) \leq C_q \times \text{arity}(q)$ distinct nulls and that this is a constant. The machine then guesses and remembers a valuation ν of D' and computes $\nu(D')$. The encoding size $\|\nu(D')\|$ of $\nu(D')$ is $O(\log |D|)$, so the machine can check in nondeterministic linear space whether $\nu(D') \models q$ and stops and rejects in the branches that fail the test. Then, the machine reads the input tape left to right and for every occurrence of a null \perp (appearing in D) that it finds, it does the following:

- It checks whether \perp appears before on the input tape and if so, then it simply continues;
- Else, if \perp does not appear before on the input tape but appears in D' , then the machine writes $\nu(\perp)$ on its output tape;
- Else, if \perp does not appear before on the input tape and does not appear in D' , then it guesses a value for it and writes that value on the output tape (but it does not remember that value).

It is easy to see that this procedure can be carried out by a logspace nondeterministic transducer, so we only need to show that the distinct outputs of the machine correspond exactly to the distinct valuations ν of D such that $\nu(D) \models q$. Since the machine writes values for nulls in order of first appearance on the input tape, it is clear that every valuation is outputted exactly once. Let ν be a valuation that is outputted, and let D' be the subdatabase such that $\nu(D') \models q$. Since $\nu(D') \subseteq \nu(D)$ and q is monotone, we have $\nu(D) \models q$. Inversely, let ν be a valuation of D such that $\nu(D) \models q$, and let us show that it must be outputted. Since $\nu(D) \models q$ and q has bounded minimal models, there exists $D_\nu \subseteq \nu(D)$ of size $\leq C_q$ such that $D_\nu \models q$. But D_ν is $\nu(D')$ for some $D' \subseteq D$ of size $\leq C_q$. Then it is clear that one of the branches of the machine has guessed D' and then $\nu|_{D'}$ and then has written ν on the output tape. \square

In particular, given that a union of Boolean of conjunctive queries satisfies the three properties of the previous proposition, we conclude from Theorem 5.1 that $\#Val(q)$ can be efficiently approximated by using a randomized algorithm if q is a union of BCQs.⁹

COROLLARY 5.3. *If q is a union of BCQs, then $\#Val(q)$ has an FPRAS (and the same holds if restricted to the uniform setting and/or to Codd tables).*

We prove in the next section that the good properties stated in Proposition 5.2 do not hold for counting completions.

5.2 Approximating the Number of Completions

In this section, we prove that the problem of counting completions of an incomplete database is much harder in terms of approximation than the problem of counting valuations. In this investigation, two randomized complexity classes play a fundamental role. Recall that RP is the class of decision problems L for which there exists a polynomial-time probabilistic Turing Machine M such that: (a) if $x \in L$, then M accepts with probability at least $3/4$; and (b) if $x \notin L$, then M does not accept x . Moreover, BPP is defined exactly as RP but with condition (b) replaced by: (b') if $x \notin L$, then M accepts with probability at most $1/4$. Thus, BPP is defined as RP but allowing errors for both the elements that are and are not in L . It is easy to see that $RP \subseteq BPP$. Besides, it is known that $RP \subseteq NP$, and this inclusion is widely believed to be strict. Finally, it is not known whether $BPP \subseteq NP$ or $NP \subseteq BPP$, but it is widely believed that NP is not included in BPP.

The non-uniform case. Recall that #IS is the problem of counting the number of independent sets of a graph. This problem will play a fundamental role when showing non-approximability of counting completions in the non-uniform case. More precisely, the following is known about the approximability of #IS:

THEOREM 5.4 ([20, THEOREM 3.1]). *The problem #IS does not admit an FPRAS unless $NP = RP$.*

In the proof of Proposition 4.2, we considered the problem #VC of counting the number of vertex covers of a graph $G = (V, E)$ and showed that $\#VC \leq_{\text{par}}^P \#\text{Comp}_{Cd}(R(x))$. By observing that $S \subseteq V$ is an independent set of G if and only if $V \setminus S$ is a vertex cover of G , we can conclude that $\#IS(G) = \#VC(G)$ and, thus, the same reduction from the proof of Proposition 4.2 establishes that $\#IS \leq_{\text{par}}^P \#\text{Comp}_{Cd}(R(x))$. Therefore, from the fact that the reduction in Lemma 4.1 is also parsimonious and preserves the property of being a Codd table, and the fact that the existence of an FPRAS is closed under polynomial-time parsimonious reductions, we obtain the following result from Theorem 5.4:

THEOREM 5.5 (DICHOTOMY). *For every sjfBCQ q , it holds that $\#\text{Comp}_{Cd}(q)$ does not admit an FPRAS unless $NP = RP$ (and, hence, the same holds for $\#\text{Comp}(q)$).¹⁰*

The uniform case. Recall that from Theorem 4.10, we know that if an sjfBCQ q contains neither $R(x, x)$ nor $R(x, y)$ as a pattern, then $\#\text{Comp}^u(q)$ is in FP. Thus, the question to answer in this section is whether $\#\text{Comp}^u(q)$ and $\#\text{Comp}_{Cd}^u(q)$ can be efficiently approximated if q contains any of these two patterns. For the case of naive tables, we will give a negative answer to this question. Notice that, this time, our reduction from #IS in Proposition 4.5 is not parsimonious, so we cannot

⁹As a matter of fact, this holds even for the larger class of unions of BCQs with *inequalities* (that is, atoms of the form $x \neq y$), as such queries also satisfy the aforementioned three properties.

¹⁰Again, we remind the reader that, to avoid trivialities, we assume all sjfBCQs to contain at least one atom and all atoms to have at least one variable.

use Theorem 5.4 as we did for the non-uniform case. Instead, we will rely on the following well-known fact: If there exists a BPP algorithm for a problem that is NP-complete, then $\text{NP} \subseteq \text{BPP}$, which in turn implies that $\text{NP} = \text{RP}$ [34].

PROPOSITION 5.6. *Neither $\#\text{Comp}^u(R(x, x))$ nor $\#\text{Comp}^u(R(x, y))$ admits an FPRAS unless $\text{NP} = \text{RP}$. This holds even in the restricted setting in which all nulls are interpreted over the same fixed domain $\{1, 2, 3\}$.*

PROOF. Let $G = (V, E)$ be a graph. First, we explain how to construct an incomplete database D containing a single binary relation R , with uniform domain $\{1, 2, 3\}$, and such that (a) all completions of D satisfy both queries; (b) if G is 3-colorable, then D has eight completions; and (c) if G is not 3-colorable, then D has seven completions. For every node $u \in V$, we have a null \perp_u . The database D consists of the following three disjoint sets of facts:

- For every edge $\{u, v\} \in E$, we have the two facts $R(\perp_u, \perp_v)$ and $R(\perp_v, \perp_u)$; we call these the *coding facts*.
- We have the facts $R(1, 2), R(2, 1), R(2, 3), R(3, 2), R(1, 3)$, and $R(3, 1)$; we call these the *triangle facts*;
- We have six fresh nulls $\perp_1, \perp'_1, \perp_2, \perp'_2, \perp_3, \perp'_3$ and the facts $R(\perp_i, \perp'_i)$ and $R(\perp'_i, \perp_i)$ for $1 \leq i \leq 3$; we call these the *auxiliary facts*;
- Last, we have a fact $R(c, c)$, where c is a fresh constant.

It is clear that all the completions of D satisfy both queries (thanks to the fact $R(c, c)$), so we only need to prove (b) and (c). Observe that a candidate completion of D can be equivalently seen as an undirected graph, possibly with self-loops, over the nodes $\{1, 2, 3\}$ (we omit the fact $R(c, c)$, since it is in every completion) and that contains the triangle. Thanks to the auxiliary facts, it is easy to show that all such graphs with at least one self-loop can be obtained as a completion of D . For instance, the completion that is triangle with a self-loop only on 1 can be obtained by assigning 1 to all the nulls in the coding facts, assigning 1 to $\perp_1, \perp'_1, \perp_2$ and \perp_3 and assigning 2 to \perp'_2 and \perp'_3 . There are seven such completions in total. Then, the completion whose graph is the triangle with no self-loops is obtainable if and only if G is 3-colorable (we assign a 3-coloring to the nulls in the coding facts, and assign 1 to \perp_i and 2 to \perp'_i for every $i \in \{1, 2, 3\}$). This indeed proves (b) and (c). Next, we show that any FRPAS with $\varepsilon = 1/16$ for counting the number of completions of D would yield a BPP algorithm to solve 3-colorability, thus implying $\text{NP} = \text{RP}$, since 3-colorability is an NP-complete problem.

Let \mathcal{A} be an FPRAS for $\#\text{Comp}^u(q)$, with q being $R(x, x)$ or $R(x, y)$, and let us define a BPP algorithm \mathcal{B} for 3-colorability using \mathcal{A} . On input graph G , algorithm \mathcal{B} does the following. First, it computes in polynomial time the naive table D as described above. Then \mathcal{B} calls \mathcal{A} with input $(D, 1/16)$, and if $\mathcal{A}(D, 1/16) \geq 7.5$, then \mathcal{B} accepts, otherwise \mathcal{B} rejects. We now prove that \mathcal{B} is indeed a BPP algorithm for 3-colorability. Assume first that G is 3-colorable. Then by (b) and by definition of what is an FPRAS, we have that $\Pr(|8 - \mathcal{A}(D, 1/16)| \leq 8/16) \geq \frac{3}{4}$. This implies in particular that $\Pr(\mathcal{A}(D, 1/16) \geq 8 - 8/16) \geq \frac{3}{4}$. Since $8 - 8/16 = 7.5$, we conclude that if G is 3-colorable, then \mathcal{B} accepts with probability at least $3/4$. Next, assume that G is not 3-colorable. Then by (c), we have that $\Pr(|7 - \mathcal{A}(D, 1/16)| \leq 7/16) \geq \frac{3}{4}$. This implies in particular that $\Pr(\mathcal{A}(D, 1/16) \leq 7 + 7/16) \geq \frac{3}{4}$. Since $7 + 7/16 < 7.5$, this implies in particular that $\Pr(\mathcal{A}(D, 1/16) < 7.5) \geq \frac{3}{4}$. From this, we conclude that if G is not 3-colorable, then \mathcal{B} rejects with probability at least $3/4$. This concludes the proof of the proposition. \square

By observing again that the reduction in Lemma 4.1 is parsimonious, and that the existence of an FPRAS is closed under parsimonious reductions, we obtain that $\#\text{Comp}^u(q)$ cannot be efficiently approximated if q contains $R(x, x)$ or $R(x, y)$ as a pattern.

THEOREM 5.7 (DICHOTOMY). *Let q be an sjfBCQ. If q has $R(x, x)$ or $R(x, y)$ as a pattern, then $\#\text{Comp}^u(q)$ does not admit an FPRAS unless $\text{NP} = \text{RP}$. Otherwise, this problem is in FP (by Theorem 4.10).*

We do not know if this result still holds for Codd tables, or if it is possible to design an FPRAS in this setting. We leave this question open for future research.

6 ON THE GENERAL LANDSCAPE: BEYOND #P

Recall that, when studying the complexity of counting completions for sjfBCQs in Section 4, we did claim that these problems are in #P for Codd tables, but that we did not claim so for naive tables. The goal of this section is then threefold. First, we want to prove that the problem of counting completions is indeed in #P for Codd tables. Second, we want to give formal evidence that we indeed could not show membership in #P for naive tables. Third, we want to identify a counting complexity class that is more appropriate to describe the complexity of $\#\text{Comp}(q)$. We deal with these three objectives in the next three sections.

6.1 Membership in #P of $\#\text{Comp}_{Cd}(q)$

In this section, and as promised in the proofs of Propositions 4.4 and 4.11, we show that for any Boolean query q , if the model checking problem for q (denoted $\text{MC}(q)$; recall the definition from Section 5.1) is in P, then the problem of counting completions for q under Codd tables is in #P.

PROPOSITION 6.1. *If a Boolean query q has the property that model checking is in P, then we have that $\#\text{Comp}_{Cd}(q)$ is in #P.*

We recall that a fact that contains only constants is a *ground fact*. To show Proposition 6.1, we first prove that we can check in polynomial time if a given set of ground facts is a possible completion of an incomplete database:

LEMMA 6.2. *Given as input an incomplete Codd table D and a set S of ground facts, we can decide in polynomial time whether there exists a valuation v of D such that $v(D) = S$.*

PROOF. For every fact f of D , let us denote by $P(f)$ the set of ground facts that can be obtained from f via a valuation ($P(f)$ can be $\{f\}$ if f is already a ground fact). The first step is to check that for every fact f of D , it holds that $(\star) P(f) \cap S \neq \emptyset$. If this is not the case, then we know for sure that for every valuation v of D , we will have $v(D) \not\subseteq S$, so we can safely reject. Next, we build the bipartite graph $G_{D,S}$ defined as follows: The nodes in the left partition of $G_{D,S}$ are the facts of D , the nodes in the right partition are the facts in S , and we connect a fact f of D with all the ground facts in the right partition that are in $S \cap P(f)$. It is clear that we can construct $G_{D,S}$ in polynomial time. We then compute in polynomial time the size m of a maximum-cardinality matching of $G_{D,S}$, for instance using Reference [21]. It is clear that we have $m \leq |S|$. At this stage, we claim that there exists a valuation v of D such that $v(D) = S$ if and only if $m = |S|$. We prove this by analyzing the two possible cases:

- If $m < |S|$, then let us show that there is no such valuation. Indeed, assume by way of contradiction that such a valuation v exists. Let B be a subset of D of minimal size such that $v(B) = S$. It is clear that such a subset exists, and moreover that its size is exactly $|S|$. But then, consider the set M of edges of $G_{D,S}$ defined by $M \stackrel{\text{def}}{=} \{(f, v(f)) \mid f \in B\}$. Then M is a matching of $G_{D,S}$ of size $|S|$, contradicting the fact that m is the size of a maximum-cardinality matching.
- If $m = |S|$, then let us show that such a valuation exists. Let M be a matching of $G_{D,S}$ of size $|S|$. It is clear that every node corresponding to a ground fact $f \in S$ is incident to

(exactly) one edge of M ; let us denote that edge by e_f . Moreover, since M is a matching, the mapping that associates to a ground fact $f \in S$ the fact f'_f at the other end of e_f is injective. Hence, we can define $v(\perp)$ of every null \perp occurring in such a fact $f'_f \in D$ to be the unique constant such that $v(f'_f) = f$ holds, and for every other fact f' in D not incident to an edge in M , we choose a value for its nulls so $v(f') \in S$, which we can do thanks to (\star) . It is then clear that we have $v(D) = S$.

But then, we can simply accept if $m = |S|$ and reject otherwise. \square

We can now prove Proposition 6.1:

PROOF OF PROPOSITION 6.1. We define a non-deterministic Turing machine M_q such that, given as input an incomplete Codd table D , its number of accepting computation paths is exactly the number of completions of D that satisfy q . First, compute in polynomial time the set $A \stackrel{\text{def}}{=} \bigcup_{f \in D} P(f)$, where $P(f)$ is defined just as in Lemma 6.2. Then, the machine M_q guesses a subset S of A . It then checks in polynomial time if S , when seen as a database, satisfies q , and rejects if it is not the case. Then, using Lemma 6.2, it checks in polynomial time whether there exists a valuation v of D such that $v(D) = S$, and accepts iff this is the case. It is then clear that M_q satisfies the conditions, which shows that $\#\text{Comp}_{Cd}(q)$ is in $\#\text{P}$. \square

In the next two sections, all upper bounds will be proved for the most general scenario of non-uniform naive tables, while all lower bounds will be proved for the most restricted scenario of uniform naive tables with a fixed domain.

6.2 Non-membership in $\#\text{P}$ of $\#\text{Comp}(q)$

We now want to show that $\#\text{P}$ is not the right complexity class for problems of the form $\#\text{Comp}(q)$, over naive tables. One could try to show membership in $\#\text{P}$ of $\#\text{Comp}(q)$ as we did in the proof of Proposition 6.1; that is, guess a set of ground facts, then check in polynomial time that it satisfies the query and that it is a possible completion of the incomplete database. However, this proof strategy does not work, as we show next that checking if a set of ground fact is a completion of an incomplete database is an NP-complete problem. Moreover, this holds already for a fixed schema containing only a binary relation and for a fixed set of ground facts.

PROPOSITION 6.3. *There exists a set S of ground facts over binary relation R such that the following is NP-complete: Given as input an incomplete database D over R , decide if S is a completion of D .*

PROOF. We reduce from 3-colorability. Given a graph G , we build the same incomplete database D as in the proof of Proposition 5.6, and the (fixed) set of ground facts is the triangle, that is, $S = \{R(1, 2), R(2, 1), R(2, 3), R(3, 2), R(1, 3), R(3, 1)\}$. Then, as in that proof, we have that S is a completion of D if and only if G is 3-colorable. \square

This does not, however, constitute a proof that $\#\text{Comp}(q)$ is not in $\#\text{P}$ (but it is a good hint). To prove that $\#\text{Comp}(q)$ is unlikely to be in $\#\text{P}$, we need to define the complexity class SPP introduced in References [24, 28, 42]. Given a nondeterministic Turing Machine M and a string x , let $\text{accept}_M(x)$ (respectively, $\text{reject}_M(x)$) be the number of accepting (respectively, rejecting) runs of M with input x , and let $\text{gap}_M(x) = \text{accept}_M(x) - \text{reject}_M(x)$. Then a language L is said to be in SPP [24] if there exists a polynomial-time nondeterministic Turing Machine M such that: (a) if $x \in L$, then $\text{gap}_M(x) = 1$; and (b) if $x \notin L$, then $\text{gap}_M(x) = 0$. It is conjectured that $\text{NP} \not\subseteq \text{SPP}$ as, for example, for every known polynomial-time nondeterministic Turing Machine M accepting an NP-complete problem, the function gap_M is not bounded. In the following proposition, we show how this conjecture helps us to reach our second goal.

PROPOSITION 6.4. *There exists an sjfBCQ q such that $\#\text{Comp}^u(q)$ is not in $\#P$ unless $\text{NP} \subseteq \text{SPP}$.*

The proof of this result relies on the proof of Theorem 6.6, in the next section (we presented the results in this order for narrative purposes). We will then defer its presentation until the proof of Theorem 6.6 is given.

6.3 An Appropriate Counting Complexity Class for $\#\text{Comp}(q)$: SpanP

To meet our third goal, we need to introduce one last counting complexity class. The class SpanP [35] is defined exactly as the class SpanL introduced in Section 5.1, but considering *polynomial-time* nondeterministic Turing machines with output, instead of *logarithmic-space* nondeterministic Turing machines with output. It is straightforward to prove that $\#P \subseteq \text{SpanP}$. Besides, it is known that $\#P = \text{SpanP}$ if and only if $\text{NP} = \text{UP}$ [35].¹¹ Therefore, it is widely believed that $\#P$ is properly included in SpanP. The following easy observation can be seen as a first hint that SpanP is a good alternative to describe the complexity of counting completions:

OBSERVATION 6.5. *If q is a Boolean query such that $\text{MC}(q)$ is in P , then $\#\text{Comp}(q)$ is in SpanP.*

Notice that this result applies to all sjfBCQs and, more generally, to all FO Boolean queries. In fact, this result applies to even more expressive query languages such as Datalog [1]. More surprisingly, in the following theorem, we show that $\#\text{Comp}^u(q)$ can be SpanP-complete for an FO query q and, in fact, already for the negation of an sjfBCQ.

THEOREM 6.6. *There exists an sjfBCQ q such that $\#\text{Comp}^u(\neg q)$ is SpanP-complete under polynomial-time parsimonious reductions.*

To prove this result, we will use the problem of counting the number of satisfying assignments of a 3-CNF formula that are distinct in the first k variables, that we denote by $\#\text{k3SAT}$. Formally, the problem $\#\text{k3SAT}$ takes as input a 3-CNF formula F on variables $\{x_1, \dots, x_n\}$ and an integer $1 \leq k \leq n$, and outputs the number of assignments of the first k variables that can be extended to a satisfying assignment of F . This problem is shown to be SpanP-complete in Reference [35]:

PROPOSITION 6.7 ([35, SECTION 6]). *$\#\text{k3SAT}$ is SpanP complete (under polynomial-time parsimonious reductions).*

We are ready to prove Theorem 6.6.

PROOF OF THEOREM 6.6. Notice that we only need to show hardness for a fixed sjfBCQ q . We reduce from $\#\text{k3SAT}$ to $\#\text{Comp}^u(\neg q)$, for a fixed sjfBCQ q to be defined. Let F be a 3-CNF on variables $\{x_1, \dots, x_n\}$, and $1 \leq k \leq n$. We first explain how we build the incomplete database D , and we will define the sjfBCQ q after. For every variable x_i , $1 \leq i \leq n$, we have a null \perp_{x_i} , and the (uniform) domain is $\{0, 1\}$. For $(a, b, c) \in \{0, 1\}^3$, we have a relation C_{abc} of arity 3, and we fill it with every tuple of the form $C_{abc}(a', b', c')$ with $(a', b', c') \in \{0, 1\}^3$ such that $a = a' \vee b = b' \vee c = c'$ holds; hence, for every $(a, b, c) \in \{0, 1\}^3$ there are exactly seven facts of this form. For every clause $K = l_1 \vee l_2 \vee l_3$ of F with l_1, l_2, l_3 being literals over variables y_1, y_2, y_3 , letting $(a_1, a_2, a_3) \in \{0, 1\}^3$ be the unique tuple such that $a_i = 1$ iff l_i is a positive literal, we add to $C_{a_1 a_2 a_3}$ the fact $C_{a_1 a_2 a_3}(\perp_{y_1}, \perp_{y_2}, \perp_{y_3})$. Last, we have a binary relation S that we fill with the tuples $S(i, \perp_{x_i})$ for $1 \leq i \leq k$. The sjfBCQ q then simply says that there exists a tuple that appears

¹¹Recall that UP is the class Unambiguous Polynomial-Time introduced in Reference [49], and that $L \in \text{UP}$ if and only if there exists a polynomial-time nondeterministic Turing Machine M such that if $x \in L$, then $\text{accept}_M(x) = 1$, and if $x \notin L$, then $\text{accept}_M(x) = 0$.

in all the relations C_{abc} :

$$q = \exists x \exists y S(x, y) \wedge \exists x \exists y \exists z \left(\bigwedge_{(a,b,c) \in \{0,1\}^3} C_{abc}(x, y, z) \right). \quad (6)$$

Note that we added the seemingly useless query $\exists x \exists y S(x, y)$ to q , because the set of relations in D has to be a subset of the set of relations occurring in q (indeed, this is how we defined our problems in Section 2). We now show that the number of completions of D that do not satisfy q is equal to the number of assignments of the first k variables that can be extended to a satisfying assignment of F , thus establishing that $\#\text{Comp}^u(\neg q)$ is SpanP-hard (under polynomial-time parsimonious reductions). First, observe that the assignments of the variables are in bijection with the valuations of the nulls of D . One can then readily observe the following:

- If q is falsified in a completion of D , then it can only be because there does not exist a tuple that occurs in all the relations; this is because the query $\exists x \exists y S(x, y)$ is always satisfied by any completion of D .
- For every assignment of the variables, letting ν be the corresponding valuation of the nulls, there exists a tuple that is in all relations C_{abc} of $\nu(D)$ if and only if that assignment is not satisfying for F . Indeed, this happens if and only if there exists a relation C_{abc} such that $\nu(D)(C_{abc})$ contains exactly eight facts.
- For every two valuations ν, ν' such that the corresponding assignments are not satisfying the formula, we have that $\nu(D) \neq \nu'(D)$ if and only if ν and ν' differ on the first k variables. This is because, by the previous item, each relation C_{abc} contains exactly the seven ground tuples that we initially put in D .

By putting it all together, we obtain that the reduction works as expected. \square

This theorem gives evidence that SpanP is the right class to describe the complexity of counting completions for FO queries (and even for queries with model checking in polynomial time). It is important to notice that SpanP-hardness is proved in Theorem 6.6 by considering parsimonious reductions. This is a delicate issue, because from the main result in Reference [48], it is possible to conclude that every counting problem that is #P-hard (even under polynomial-time parsimonious reductions) is also SpanP-hard under polynomial-time Turing reductions, so a more restrictive notion of reduction has to be used when proving that a counting problem is SpanP-hard [35].

Before continuing, we prove Proposition 6.4.

PROOF OF PROPOSITION 6.4. Let q be the sjfBCQ defined in Equation (6) in the proof of Theorem 6.6. Its schema $\sigma = \{S\} \cup \{C_{abc} \mid (a, b, c) \in \{0, 1\}^3\}$ consists of 10 relation symbols, with S being binary and each C_{abc} being ternary. Let us denote by $\#\text{Comp}^u(\sigma)$ the problem that takes as input an incomplete database over schema σ and outputs its number of completions. The first part of our proof is to reduce $\#\text{Comp}^u(\sigma)$ to $\#\text{Comp}^u(q)$; formally, we claim that $\#\text{Comp}^u(\sigma) \leq_{\text{par}}^p \#\text{Comp}^u(q)$. Indeed, let D be an incomplete database over schema σ , that is an input of $\#\text{Comp}^u(\sigma)$. We construct in polynomial time an incomplete database D' over the same schema such that $\#\text{Comp}^u(\sigma)(D) = \#\text{Comp}^u(q)(D')$, thus establishing the parsimonious reduction. Let f be a fresh constant that does not occur in D nor in the domain of some null. Then the relation $D'(S)$ is the same as the relation $D(S)$, plus a fact $S(f, f)$. Moreover, for every $(a, b, c) \in \{0, 1\}^3$, the relation $D'(C_{abc})$ consists of all the facts in $D(C_{abc})$, plus a fact $C_{abc}(f, f, f)$. It is easy to see that D and D' have the same number of completions. Moreover, thanks to the facts that use the constant f , we have that every completion of D' satisfies q . Therefore, we indeed have that $\#\text{Comp}^u(\sigma)(D) = \#\text{Comp}^u(q)(D')$. This proves that $\#\text{Comp}^u(\sigma) \leq_{\text{par}}^p \#\text{Comp}^u(q)$.

For the second part of the proof, we need to introduce the complexity class GapP. This class consists of function problems that can be expressed as the difference of two functions in #P [24, 29]. It is known that if the inclusion $\text{SpanP} \subseteq \text{GapP}$ holds, then we have that $\text{NP} \subseteq \text{SPP}$ [39].¹² With this, we are able to prove the proposition. Assume that $\#\text{Comp}^u(q)$ is in #P. Then, by the first part of the proof, we have that $\#\text{Comp}^u(\sigma) \in \#P$ as well (because #P is closed under polynomial-time parsimonious reductions). Now, observe that for every incomplete database D over σ , the following holds:

$$\#\text{Comp}^u(\neg q)(D) = \#\text{Comp}^u(\sigma)(D) - \#\text{Comp}^u(q)(D).$$

But then this means that $\#\text{Comp}^u(\neg q)$ is in GapP (since both problems in the right-hand side are in #P). Since $\#\text{Comp}^u(\neg q)$ is SpanP-complete by Theorem 6.6 under polynomial-time parsimonious reductions and, since GapP is closed under polynomial-time parsimonious reductions, this would indeed imply that $\text{SpanP} \subseteq \text{GapP}$ and, hence, that $\text{NP} \subseteq \text{SPP}$. \square

We conclude this section by considering an even more general scenario where queries have model checking in NP. Interestingly, in this case, SpanP is again the right class to describe the complexity not only of counting completions, but also of counting valuations.

THEOREM 6.8. *If q is a Boolean query with $\text{MC}(q) \in \text{NP}$, then both $\#\text{Val}(q)$ and $\#\text{Comp}(q)$ are in SpanP. Moreover, there exists such a Boolean query q for which $\#\text{Val}^u(q)$ is SpanP-complete under polynomial-time parsimonious reductions (and for $\#\text{Comp}^u(q)$, we can even take q to be the negation of an sjfBCQ, hence, with model checking in P, as given by Theorem 6.6).*

PROOF. It is straightforward to prove that these problems are in SpanP. The part in between parentheses has been shown in Theorem 6.6. Thus, we need to prove that $\#\text{Val}^u(q)$ is SpanP-hard for a fixed Boolean query q such that $\text{MC}(q) \in \text{NP}$, under polynomial-time parsimonious reductions. To do this, we will reduce from the SpanP-complete problem #HamSubgraphs, defined as follows:

Let $G = (V, E)$ be a undirected graph, and let $S \subseteq V$. The *subgraph of G induced by S* , denoted by $G[S]$, is the graph with set of nodes S and set of edges $\{\{u, v\} \in E \mid u, v \in S\}$. We recall that a graph G is *Hamiltonian* when there exists a cycle in G that visits every node of G exactly once. The problem #HamSubgraphs takes as input a simple graph $G = (V, E)$ and an integer k , and outputs the number of induced subgraphs $G[S]$ with $|S| = k$ such that $G[S]$ is Hamiltonian.

PROPOSITION 6.9 ([35, SECTION 6]). *#HamSubgraphs is SpanP-complete (under polynomial-time parsimonious reductions).*

Next, we show that $\#\text{HamSubgraphs} \leq_{\text{par}}^{\text{P}} \#\text{Val}^u(q)$, for a fixed Boolean query q (to be defined). Let $G = (V, E)$ be an undirected graph. We first explain how we construct the incomplete database D , and we will then define the query q . The schema contains two binary relation symbols R, T and one unary relation symbol K . Fix a linear order a_1, \dots, a_n of the nodes of G . For every edge $\{u, v\} \in E$, we have the facts $R(u, v)$ and $R(v, u)$. For $1 \leq i \leq n$, we have a fact $T(a_i, \perp_i)$, and the domain of the nulls is $\{0, 1\}$. For $1 \leq j \leq k$, we have a fact $K(j)$. Observe that D is a Codd table. We now define the Boolean query q , which will be a sentence in existential second-order logic ($\exists\text{SO}$) over relational signature R, T, K . Before doing so, we explain the main idea: Intuitively, q will check that there are exactly k facts of the form $T(a_i, 1)$ in the relation T and that, letting S

¹²In fact, the class GapSpanP is defined in Reference [39], where it is proved that a function f is in GapSpanP if and only if $f = g - h$, where h, g are functions in SpanP. Then it is shown in Reference [39, Corollary 3.5] that the inclusion $\text{GapSpanP} \subseteq \text{GapP}$ implies that $\text{NP} \subseteq \text{SPP}$. But if we have that $\text{SpanP} \subseteq \text{GapP}$, then we also have that $\text{GapSpanP} \subseteq \text{GapP}$ as GapP is closed under subtraction and, therefore, we conclude that $\text{NP} \subseteq \text{SPP}$, as desired.

be the set of nodes v such that $T(v, 1)$ is in relation T , the induced subgraph $G[S]$ is Hamiltonian. This will indeed ensure that we have $\#\text{Val}^u(q)(D) = \#\text{HamSubgraphs}(G, k)$, thus completing this reduction, which is parsimonious and can be performed in polynomial-time. The query is

$$q = \exists S \psi_1(S) \wedge \psi_2(S),$$

where S is a unary second order variable and the formula $\psi_1(S)$ states that (a) the elements s of S are exactly all the elements such that $T(s, 1)$ holds, and that (b) there are exactly the same number of elements in S as there are elements j for which $K(j)$ holds. It is clear that (a) can be expressed in FO. Moreover, (b) can be expressed in $\exists\text{SO}$ by asserting the existence of a binary second-order relation U that represents a bijective function from S to the elements in K . Then $\psi_2(S)$ is a formula that asserts that $G[S]$ is Hamiltonian. Since this is a property in NP, $\psi_2(S)$ can be expressed in $\exists\text{SO}$ by Fagin's theorem (see, e.g., Reference [31]). This shows that the reduction is correct. Finally, the fact that $\text{MC}(q)$ is in NP again follows from Fagin's theorem. This concludes the proof. \square

7 EXTENSIONS TO QUERIES WITH CONSTANTS AND FREE VARIABLES

So far, we have only considered our counting problems for queries that are Boolean and that do not contain constants. In this section, we explain how our framework can be adapted to queries with constants and with free variables. Specifically, we will explain how one can obtain dichotomies for self-join-free conjunctive queries with constants and free variables.

Before that, we have to formally define our counting problems for a query with free variables. Let $q(\bar{x})$ be a query with free variables \bar{x} . For a tuple of constants \bar{t} of appropriate arity, we write $q(\bar{t})$ the Boolean query obtained by substituting the variables \bar{x} with the constants \bar{t} . The problem $\#\text{Val}(q(\bar{x}))$ then takes as input an incomplete database D over relations $\text{sig}(q(\bar{x}))$, a tuple of constants \bar{t} , and returns the number of valuations ν of D such that $\nu(D) \models q(\bar{t})$. We write this output $\#\text{Val}(q(\bar{x}))(D, \bar{t})$. The problem $\#\text{Comp}(q(\bar{x}))$ is defined similarly.

We first explain in Section 7.1 how to obtain dichotomies for self-join-free conjunctive query with free variables and constants, assuming we have dichotomies for self-join-free Boolean conjunctive queries with constants. In Section 7.2, we then explain how to obtain dichotomies for the later case.

7.1 Dealing with Free Variables

Suppose in this section that we have a dichotomy between $\#\text{P}$ -hardness and FP of our counting problems for every sjfBCQ that is allowed to contain constants. We then show how to obtain dichotomies for every self-join-free CQ that is allowed to have constants *and* free variables. To this end, we will need the following definition: Let $q(\bar{x})$ be a self-join-free CQ with free variables \bar{x} , and let \bar{t} and \bar{t}' be two tuples of constants with appropriate arity. We say that \bar{t} and \bar{t}' are *equivalent with respect to $q(\bar{x})$* if one can go from $q(\bar{t})$ to $q(\bar{t}')$ by iteratively renaming constants into fresh constants. For instance, if $q(x_1, x_2)$ is $\exists y R(y, c, x_1, x_2) \wedge S(c', x)$, then (c, c') is equivalent to (c, c''') , and (c'', c''') and (c''', c'') are also equivalent. It is clear that, if $q(\bar{x})$ is fixed, then this defines an equivalence relation that has only finitely many classes. Furthermore, it is also clear that if \bar{t} and \bar{t}' are equivalent with respect to $q(\bar{x})$, then the problems $\#\text{Val}(q(\bar{t}))$ and $\#\text{Val}(q(\bar{t}'))$ (respectively, $\#\text{Comp}(q(\bar{t}))$ and $\#\text{Comp}(q(\bar{t}'))$) have the same complexity. We can now show what we wanted. In what follows, hardness refers to $\#\text{P}$ -hardness (but it is not important for the proof):

LEMMA 7.1. *Assume that the following is true: for every sjfBCQ q that is allowed to have constants, the problem $\#\text{Val}(q)$ is either hard or is tractable. Then the following is also true: For every self-join-free conjunctive query $q(\bar{x})$, the problem $\#\text{Val}(q(\bar{x}))$ is either hard or is tractable. This holds also for counting completions, and when restricted to Codd tables and/or to the uniform setting.*

PROOF. We only deal with counting valuations for the naive and non-uniform case, as the other cases are similar. We prove the following for any self-join-free CQ $q(\bar{x})$, which implies the claim: If there exists a tuple of constants \bar{t} such that $\#Val(q(\bar{t}))$ is hard, then $\#Val(q(\bar{x}))$ is hard as well, otherwise $\#Val(q(\bar{x}))$ is tractable. We start with the “if” direction. Let \bar{t} be a tuple of constants such that $\#Val(q(\bar{t}))$ is hard. By definition, it is clear that, for any incomplete database D , we have that $\#Val(q(\bar{t}))(D) = \#Val(q(\bar{x}))(D, \bar{t})$; this shows hardness of $\#Val(q(\bar{x}))$. Now for the “otherwise” direction: Let $\bar{t}_1, \dots, \bar{t}_k$ be representatives of the finitely many equivalence classes of the equivalence relation defined above. We have access to oracles for $\#Val(q(\bar{t}_1)), \dots, \#Val(q(\bar{t}_k))$. Let D, \bar{t} be an input of $\#Val(q(\bar{x}))$. We then simply recognize (in constant time, since the query is fixed) to which \bar{t}_i the tuple \bar{t} is equivalent with respect to $q(\bar{x})$, and call the appropriate oracle. \square

(Notice that this idea actually works for conjunctive queries (with self-joins), or even unions of conjunctive queries.) Hence, the problem becomes that of obtaining dichotomies for self-join-free Boolean conjunctive queries that can contain constants. We explain in the next section how this can be done.

7.2 Dealing with Constants

In this section, we simply write “an sjfBCQ” to mean an sjfBCQ that can contain some constants. To the best of our knowledge, there is no general reduction that allows us to easily obtain dichotomies for the case with constants from the case where queries do not have constants. Hence, the strategy to obtain dichotomies will be the same as in Section 3; namely, we will again use the notion of pattern to find a set of hard patterns, and then show that when a query does not have any of the hard patterns, then the problem is tractable. The notion of pattern for an sjfBCQ that can contain constants is the same as the one we used in Definition 3.1, but we simply add the possibility of deleting an occurrence of a constant. For instance, the query $R(x, c)$ is a pattern of the query $R(x, x, c, c, c')$ (in this section all variables are existentially quantified, since we consider only Boolean queries). The main property of patterns that we used then extends in this setting, as shown next.

LEMMA 7.2. *Let q, q' be sjfBCQs such that q' is a pattern of q . Then, we have $\#Val(q') \leq_{\text{par}}^p \#Val(q)$. Moreover, the same results hold if we restrict to Codd tables, and/or to the uniform setting, and/or to counting completions.*

PROOF. The reduction is exactly the same as that of Lemma 3.3, the only difference being that, when we delete an occurrence of a constant c , we simply fill the corresponding columns of every tuple with this constant. \square

Next, we prove dichotomies for counting valuations for the non-uniform setting for naive and Codd databases. For Codd databases, it turns out that there is no new hard pattern that involves constants, so the dichotomy is the same as for sjfBCQs without constants. Formally:

THEOREM 7.3 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x) \wedge S(x)$ is a pattern of q , then $\#Val_{Cd}(q)$ is #P-complete. Otherwise, $\#Val_{Cd}(q)$ is in FP.*

PROOF. The proof is exactly as the proof of Theorem 3.7, with the following modification: We let $\rho(\bar{t}_j)$ be the number of valuations of the nulls appearing in \bar{t}_j that do not match the corresponding atom of q ; and clearly, this can again be computed in polynomial time. \square

For naive tables, however, we find two new hard patterns that involve constants, as shown next.

PROPOSITION 7.4. *Let c, c' be two distinct constants. The problems $\#Val(R(c, c))$ and $\#Val(R(c, c'))$ are both #P-hard.*

PROOF. We only explain for $\#\text{Val}(R(c, c'))$, as the other case is analogous. The reduction is similar to that used in Proposition 3.8, but we reduce from counting the number of independent sets in *bipartite* graphs (in Proposition 3.8, we did not need the graphs to be bipartite). Let $G = (U \sqcup V, E)$ be a bipartite graph. We have one null \perp_u for every node $u \in U$ with domain $\text{dom}(\perp_u) = \{0, c\}$, and one null \perp_v for every node $u \in V$ with domain $\text{dom}(\perp_v) = \{0, c'\}$. For every edge $(u, v) \in E$, we have a fact $R(\perp_u, \perp_v)$ in D . Then it is clear that the number of valuations of D that do not contain $R(c, c')$ is equal to the number of independent sets of G , thus establishing hardness. (Note that for $R(c, c)$, we do not need the graph to be bipartite for the reduction to work.) \square

We then claim that these are the only additional patterns that are necessary to obtain a dichotomy for $\#\text{Val}(q)$ for sjfBCQs with constants.

THEOREM 7.5 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x) \wedge S(x)$ or $R(c, c)$ or $R(c, c')$ for $c \neq c'$ is a pattern of q , then $\#\text{Val}(q)$ is $\#\text{P}$ -complete. Otherwise, $\#\text{Val}(q)$ is in FP.*

The $\#\text{P}$ -hardness part of the claim follows from Lemma 7.2 and Propositions 3.4, 3.5, and 7.4. We now show the tractability claim. First, observe that not having any of these patterns means the following: every variable in q has exactly one occurrence and every atom of q contains at most one constant (but notice that the same constant can appear in multiple atoms). But then, because the database is Codd and because q has no self-joins, by multiplying by the appropriate factor, the problem $\#\text{Val}(q)$ reduces to the problem $\#\text{Val}(q')$ where q' is an sjfBCQ of the form $R_1(c_1) \wedge R_2(c_2) \wedge \dots \wedge R_k(c_k)$, where the constants c_1, \dots, c_k are not necessarily distinct. We give next an example proof that the problem is tractable for such a simple query.

Example 7.6. Let q be the query $R(c) \wedge S(c')$ with $c \neq c'$ and D be an incomplete naive database. We explain how to compute $\#\text{Val}(\neg q)(D)$ (the number of valuations that do not satisfy the query) in FP. This is enough, since the total number of valuations can clearly be computed in FP. First, we can assume without loss of generality that D does not contain ground atoms that already satisfy the query. Then, let $B^{RS} \stackrel{\text{def}}{=} \{\perp_1^{RS}, \dots, \perp_{n_{RS}}^{RS}\}$ be the set of nulls that occur in both $D(R)$ and $D(S)$, $B^R \stackrel{\text{def}}{=} \{\perp_1^R, \dots, \perp_{n_R}^R\}$ be the set of nulls that occur in $D(R)$ but not in $D(S)$, and $B^S \stackrel{\text{def}}{=} \{\perp_1^S, \dots, \perp_{n_S}^S\}$ be the set of nulls that occur in $D(S)$ but not in $D(R)$. Let D' be the database that contains only the facts $R(\perp)$ and $S(\perp)$ for $\perp \in B^{RS}$. Notice that the number of valuations of the nulls in B^R such that no null has value c is $\prod_{i=1}^{n_R} |\text{dom}(\perp_i^R) \setminus \{c\}|$ and that the number of valuations of the nulls in B^R such that some null has value c is then $\prod_{i=1}^{n_R} |\text{dom}(\perp_i^R)| - \prod_{i=1}^{n_R} |\text{dom}(\perp_i^R) \setminus \{c\}|$, and a similar expression can be obtained for the valuations of the nulls in B^S and constant c' . Then, by case analysis of whether some null in B^R has value c or not, and whether some null in B^S has value c' or not, we obtain that $\#\text{Val}(\neg q)(D) = A + B + C$, where

$$A = \left(\prod_{i=1}^{n_R} |\text{dom}(\perp_i^R)| - \prod_{i=1}^{n_R} |\text{dom}(\perp_i^R) \setminus \{c\}| \right) \times \prod_{i=1}^{n_S} |\text{dom}(\perp_i^S) \setminus \{c'\}| \times \prod_{i=1}^{n_{RS}} |\text{dom}(\perp_i^{RS}) \setminus \{c'\}|$$

is the number of valuations of D that do not satisfy q and such that some null in B^R has value c ,

$$B = \prod_{i=1}^{n_R} |\text{dom}(\perp_i^R) \setminus \{c\}| \times \left(\prod_{i=1}^{n_S} |\text{dom}(\perp_i^S)| - \prod_{i=1}^{n_S} |\text{dom}(\perp_i^S) \setminus \{c'\}| \right) \times \prod_{i=1}^{n_{RS}} |\text{dom}(\perp_i^{RS}) \setminus \{c\}|$$

is the number of valuations of D that do not satisfy q and such that no null in B^R has value c and some null in B^S has value c' , and

$$C = \prod_{i=1}^{n_R} |\text{dom}(\perp_i^R) \setminus \{c\}| \times \prod_{i=1}^{n_S} |\text{dom}(\perp_i^S) \setminus \{c'\}| \times \#\text{Val}(\neg q)(D')$$

is the number of valuations of D that do not satisfy q and such that no null in B^R has value c and no null in B^S has value c' . Hence, we only have to explain how to compute $\#\text{Val}(\neg q)(D')$ in polynomial time. For $k \in \{0, \dots, n_{RS}\}$, let D'_k be the database containing only the facts of D' over the nulls $\perp_1^{RS}, \dots, \perp_k^{RS}$. We then define the quantities $V(S, k)$, for $S \in \{\{c\}, \{c'\}, \{c, c'\}, \emptyset\}$ and $k \in \{0, \dots, n_{RS}\}$ to be the number of valuations of D_k that do not satisfy q and such that no null has a value that is in S . Notice then that $V(\emptyset, n_{RS}) = \#\text{Val}(\neg q)(D')$. For an element a and set A , we write $[a \in A]$ to mean 1 if $a \in A$ and 0 otherwise. But then we can easily compute the quantities $V(S, k)$ by dynamic programming using the following relations:

$$\begin{aligned} V(\emptyset, k) = & [c \in \text{dom}(\perp_k^{RS})] \times V(\{c\}, k - 1) \\ & + [c' \in \text{dom}(\perp_k^{RS})] \times V(\{c'\}, k - 1) \\ & + |\text{dom}(\perp_k^{RS}) \setminus \{c, c'\}| \times V(\emptyset, k - 1) \end{aligned}$$

and

$$V(\{c, c'\}, k) = |\text{dom}(\perp_k^{RS}) \setminus \{c, c'\}| \times V(\{c, c'\}, k - 1)$$

and

$$\begin{aligned} V(\{c\}, k) = & [c' \in \text{dom}(\perp_k^{RS})] \times V(\{c, c'\}, k - 1) \\ & + |\text{dom}(\perp_k^{RS}) \setminus \{c, c'\}| \times V(\{c, c'\}, k - 1) \end{aligned}$$

and a similar relation for $V(\{c'\}, k)$.

The proof for the general case is simply a tedious generalization of the proof for this example and is not very interesting, so we omit it.

Therefore, by combining Theorems 7.3 and 7.5 with Lemma 7.1, we obtain dichotomies for counting valuations of self-join-free conjunctive queries that can contain constants and free variable for the non-uniform setting for both naive and Codd tables. It is likely that dichotomies can be obtained for counting valuations in the uniform setting by using the same methodology, but we do not pursue this further, as our goal in this section is not to be exhaustive but rather to give the main ideas to be able to handle free variables and constants.

For counting completions, the reader can check that the pattern $R(c)$ is hard for $\#\text{Comp}_{Cd}$ (so in this case again all queries are hard), and that for $\#\text{Comp}_{Cd}^u$ any query that contains an atom that is not unary is again hard (it is evident from the proof of Proposition 4.9). By combining these observations with Lemma 7.1 and with the tractability proof for $\#\text{Comp}^u$ (which can be shown to extend in this case), we obtain four dichotomies for counting completions for self-join-free conjunctive queries that can contain constants and free variables.

Last, we point out that by using the same methodology, one also can extend our results on approximations to the case of queries containing constants and free variables (since the relevant reductions are parsimonious).

8 RELATED WORK

There are two main lines of work that must be compared to what we do in this article. In both cases the goal is to go beyond the traditional notion of *certain answers* that so far had been used almost exclusively to deal with query answering over uncertain data. We discuss them here, explain how they relate to our problems, and what are the fundamental differences.

Best answers and 0-1 laws for incomplete databases. Libkin has recently introduced a framework that can be used to measure the certainty with which a Boolean query holds on an incomplete database, and also to compare query answers (for a non-Boolean query) [38].

For a Boolean query q , incomplete database D , and integer k , he defines the quantity $\mu^k(q, D)$ as $\frac{|\text{Supp}^k(q, D)|}{|V^k(D)|}$, where $V^k(D)$ denotes the set of valuations of D with domain $\{1, \dots, k\}$, and $\text{Supp}^k(q, D)$ denotes the set of valuations $v \in V^k(D)$ such that $v(D) \models q$; hence, $\mu^k(q, D)$ represents the relative frequency of valuations v in $\{1, \dots, k\}$ for which the query is satisfied. He then shows that, for a very large class of queries (namely, *generic queries*), the value $\mu^k(q, d)$ always tends to 0 or 1, as k tends to infinity (and the same results hold when considering completions instead of valuations). This means that, intuitively, over an infinite domain, the query q is either almost certainly true or almost certainly false.

He also studies the complexity of finding best answers for a non-Boolean query q . As mentioned in the introduction, a tuple \bar{a} is a better answer than another tuple \bar{b} when for every valuation v of D , if we have $\bar{b} \in q(v(d))$, then we also have $\bar{a} \in q(v(d))$. A best answer is then an answer such that there is no other answer strictly better than it (under inclusion of the sets of satisfying valuations). He studies the complexity of comparing answers under this semantics and that of computing the set of best answers (see also Reference [25]).

There are several crucial differences between this previous work and ours. First, Libkin does not study the complexity of computing $\mu^k(q, d)$. We do this under the name $\#\text{Val}^u(q)$; moreover, we also study the setting in which the domains are not uniform. Second, knowing that a tuple is the best answer might not tell us anything about the size of its “support,” i.e., the number of valuations that support it. In particular, a best answer is not necessarily an answer that has the biggest support. Finally, under the semantics of better answers it does not matter if we look at the completions or at the valuations (i.e., a tuple is a best answer with respect to inclusion of valuations iff it is the best answer with respect to completions); while we have shown that it does matter for counting problems.

Counting problems for probabilistic databases and consistent query answering. Remarkably, counting problems have received considerable attention in other database scenarios where uncertainty issues appear. As mentioned in the introduction, this includes the settings of probabilistic databases and inconsistent databases. In the former case, uncertainty is represented as a probability distribution on the possible states of the data [19, 47]. There, query answering amounts to computing a weighted sum of the probabilities of the possible states of the data that satisfy a query q . We call this problem $\text{Prob}(q)$. In the case of inconsistent databases, we are given a set Σ of constraints and a database D that does not necessarily satisfy Σ ; cf. References [9, 12, 13]. Then the task is to reason about the set of all *repairs* of D with respect to Σ [9]. In our context, this means that one wants to count the number of repairs of D with respect to Σ that satisfy a given query q . When q and Σ are fixed, we call this problem $\#\text{Repairs}(q, \Sigma)$.

Both $\text{Prob}(q)$ and $\#\text{Repairs}(q, \Sigma)$ have been intensively studied already. To start with, counting complexity dichotomies have been obtained for the problem $\#\text{Repairs}(q, \Sigma)$; e.g., Reference [40] gives a dichotomy for this problem when q is an sjfBCQ and σ consists of primary keys, and Reference [41] extends this result to CQs with self-joins but only for unary keys constraints. We also mention Reference [15], where the problem of counting repairs such that a particular input tuple is in the result of the query on the repair is studied. A seemingly close counting problem for probabilistic databases is the problem $\text{Prob}(q)$ over **block independent disjoint (BIDs)** databases. We do not define it formally here, but counting repairs under primary keys can be seen as a special case of this problem, where the tuples in a “block” all have the same probability, and where the sum of the probabilities sum to 1 (and in BIDs this sum is allowed to be < 1 , meaning that a block can be completely erased). Dichotomies for this problem have been obtained in Reference [18] for sjfBCQs. Counting complexity dichotomies for other models of probabilistic databases also exist;

e.g., for *tuple-independent* probabilistic databases in which each fact is assigned an independent probability of being part of the actual dataset. Interestingly, dichotomies in this case hold for arbitrary unions of BCQs, and thus not just for sjfBCQs [19].

In some cases, one can use a problem of the form $\#\text{Repairs}(q, \Sigma)$ (or $\text{Prob}(q)$) to show the hardness of a problem of the form $\#\text{Val}(q')$. For instance, in Section 3.1, we used the $\#\text{P}$ -hardness of $\#\text{Repairs}(R'(\underline{y}, x) \wedge S'(\underline{z}, x))$ to prove that of $\#\text{Val}_{Cd}(R(x) \wedge S(x))$. In general, however, the problems $\#\text{Repairs}(q, \Sigma)$ and $\text{Prob}(q)$ seem to be unrelated to our problems for the following reasons: First, in our setting, the nulls can appear anywhere, so there is no notion of primary keys here; hence, it seems unlikely that one can design a generic reduction from the problem of counting valuations/completions to the problem of counting repairs. In fact, it would perfectly make sense to study our counting problems where we add constraints such as functional dependencies. Second, in the BID and counting repairs problems, each “valuation” (repair) gives a different complete database, while in our case, we have seen that this is not necessarily the case. In particular, problems of the form $\#\text{Comp}(q)$ have no analogues in these settings, whereas we have seen that they behave very differently in our setting.

Concerning approximation results, it is known that the problems $\#\text{Repairs}(q, \Sigma)$ and $\text{Prob}(q)$ admit an FPRAS in some important settings. In particular, when q is a union of BCQs, this holds for $\#\text{Repairs}(q, \Sigma)$ when Σ is a set of primary keys [15], and for $\text{Prob}(q)$ over BID and tuple-independent probabilistic databases [18]. We observe here that this is reminiscent of our Corollary 5.3, which shows that problems of the form $\#\text{Val}(q)$ have an FPRAS for every union of BCQs.

9 FINAL REMARKS

Our work aims to be a first step in the study of counting problems over incomplete databases. The main conclusion behind our results is that the counting problems studied in this article are particularly hard from a computational point of view, especially when compared to more positive results obtained in other uncertainty scenarios; e.g., over probabilistic and inconsistent databases. As we have shown, a particularly difficult problem in our context is that of counting completions, even in the uniform setting where all nulls have the same domain. In fact, Proposition 4.9 shows that this problem is $\#\text{P}$ -hard even in very restricted scenarios, and Proposition 5.6 that it cannot be approximated by an FPRAS. It seems, then, that the only way in which one could try to tackle this problem is by developing suitable tractable heuristics, without provable quantitative guarantees, but that works sufficiently well in practical scenarios. An example of this could be developing algorithms that compute “under-approximations” for the number of completions of a naive table satisfying a certain sjfBCQ q . Notice that a related approach has been proposed by Console et al. for constructing under-approximations of the set of certain answers by applying methods based on many-valued logics [17].

We plan to continue working on several interesting problems that are left open in this article. First, we would like to pinpoint the complexity of $\#\text{Comp}(q)$ when q is an sjfBCQ; in particular, whether this problem is SpanP -complete for at least one such a query. We also want to study whether the non-existence of FPRAS for $\#\text{Comp}(q)$ established in Proposition 5.6 continues to hold over Codd tables. We would also like to develop a more thorough understanding of the role of fixed domains in our dichotomies. In several cases, that we have explicitly stated, our lower bounds hold even if nulls in tables are interpreted over a fixed domain. Still, in some cases, we do not know whether this holds. These include, e.g., Proposition 3.11, Proposition 4.2, and Proposition 4.9. Finally, it would also be interesting to study these counting problems under bag semantics (instead of the set semantics used in this article), or consider arbitrary conjunctive queries as opposed to only self-join-free ones.

APPENDICES

A PROOFS FOR SECTION 3 (DICHOTOMIES FOR COUNTING VALUATIONS)

A.1 Proof of Theorem 3.9

In this section, we prove the tractability claim of the following dichotomy theorem.

THEOREM 3.9 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x) \wedge S(x, y) \wedge T(y)$ or $R(x, y) \wedge S(x, y)$ is a pattern of q , then $\#\text{Val}^u(q)$ is $\#\text{P}$ -complete. Otherwise, $\#\text{Val}^u(q)$ is in FP.*

First, to characterize the queries that do not have these patterns, we will use the notion of *connectivity graph* of an sjfBCQ q :

Definition A.1. Let q be an sjfBCQ. The *connectivity graph* of q is the graph $G_q = (V, E)$ with labeled edges, where V is the set of atoms of q , and for every two atoms $R(\bar{x}_i), S(\bar{y}_i)$ of q , if they share a variable then we have an edge between the corresponding nodes of G_q , that edge being labeled with the variables in $\bar{x}_i \cap \bar{y}_i$.

Example A.2. Figure 3 shows the connectivity graph of the query

$$R_1(x_1, x_1, y_1, t_1), R_2(x_1, y_1, t_2), S_1(x_2, t_3), S_2(x_2, t_4), S_3(x_2), T_1(x_3), T_2(x_3), T_3(x_3), T_4(x_3, t_5).$$




Fig. 3. The connectivity graph G_q of the sjfBCQ q from Example A.2.

The following is then readily observed:

LEMMA A.3. *Let q be an sjfBCQ that does not contain any of the patterns mentioned in Theorem 3.9. Then for every connected component C of G_q , C is a clique and there exists a variable such that all edges of C are labeled by exactly that variable.*

PROOF. First, observe that every edge of G_q must be labeled by exactly one variable, as otherwise the query q would contain the pattern $R(x, y) \wedge S(x, y)$. Let C be a connected component of G_q . Then, we have:

- C is a clique. Indeed, assume by contradiction that C is not a clique. Then, since C is connected and is not a clique, we can find 3 nodes $A_1(\bar{x}), A_2(\bar{x}'), A_3(\bar{x}'')$ such that $A_1(\bar{x})$ is adjacent to $A_2(\bar{x}')$, $A_2(\bar{x}')$ is adjacent to $A_3(\bar{x}'')$, and $A_1(\bar{x})$ is not adjacent to $A_3(\bar{x}'')$. Let X be $\bar{x} \cap \bar{x}'$ and Y be $\bar{x}' \cap \bar{x}''$, i.e., the labels on the two corresponding edges of C . By definition of G_q and, since $A_1(\bar{x})$ is not adjacent to $A_3(\bar{x}'')$, we must have $X \cap Y = \emptyset$. But X and Y are not empty (again by definition of G_q), so by picking x in X and y in Y , we see that q contains the pattern $R(x) \wedge S(x, y) \wedge T(y)$, a contradiction.
- There exists a variable that labels every edge of C . Indeed, since every edge of G_q is labeled by exactly one variable and, since C is a clique, if it was not the case, then again we could find the pattern $R(x) \wedge S(x, y) \wedge T(y)$ in q .

This concludes the proof. \square

For instance, the query from Example A.2 does not satisfy this criterion, since the edge in the first connected component of G_q is labeled by two variables. However, if we consider the query $S_1(x_2, t_3), S_2(x_2, t_4), S_3(x_2), T_1(x_3), T_2(x_3), T_3(x_3), T_4(x_3, t_5)$ (i.e., we remove the first connected component), then it satisfies the criterion.

We will also use the general fact that for an sjfBCQ q , we can assume wlog that q does not contain variables that occur only once:

LEMMA A.4. *Let q be an sjfBCQ, and let q' be the sjfBCQ obtained from q by deleting all the variables that have only one occurrence in q . Then $\#\text{Val}^u(q) \leq_1^p \#\text{Val}^u(q')$.*

PROOF. Let D be an incomplete database input of $\#\text{Val}^u(q)$. Let S be set of nulls \perp such that:

- \perp occurs in a column corresponding to a variable that has been deleted; and
- \perp does not occur in a column corresponding to a variable that has not been deleted.

Then, letting D' be the database obtained from D by projecting out the columns corresponding to the deleted variables, it is clear that we have $\#\text{Val}^u(q)(D) = \#\text{Val}^u(q')(D') \times \prod_{\perp \in S} |\text{dom}(\perp)|$, where dom is the uniform domain of the nulls. We note here that this lemma is also true in the non-uniform setting. \square

By Lemma A.3 and Lemma A.4, it is enough to show the tractability of $\#\text{Val}^u(q)$ when q is of the form $C_1(x_1) \wedge \dots \wedge C_m(x_m)$, where each $C_i(x_i)$ is what we call a *basic singleton query*, i.e., is a conjunction of unary atoms over the same variable x_i . We call such an sjfBCQ a *conjunction of basic singletons*. For instance,

$$S_1(x_2), S_2(x_2), S_3(x_2), T_1(x_3), T_2(x_3), T_3(x_3), T_4(x_3)$$

is such a query, with $m = 2$. We will use the following:

LEMMA A.5. *Let $q = C_1(x_1) \wedge \dots \wedge C_m(x_m)$ be a conjunction of basic singletons sjfBCQ, and let D be an incomplete database. For $S \subseteq [m]$, we define $N_S(D) \stackrel{\text{def}}{=} |\{v \text{ valuation of } D \mid v(D) \not\models \bigvee_{i \in S} C_i(x_i)\}|$. Then, we have $\#\text{Val}^u(q)(D) = \sum_{S \subseteq [m]} (-1)^{|S|} N_S(D)$.*

PROOF. Direct, by inclusion–exclusion. \square

Hence, and remembering that we consider data complexity, it is enough to show how to compute $N_S(D)$ for every $S \subseteq [m]$. The main difficulties in computing $N_S(D)$ is that the relations can have nulls in common (since we consider naive tables), and that they may also have constants; this makes it technically painful to express a closed-form expression for $N_S(D)$. We explain how to do it next, thus finishing the proof of Theorem 3.9.

PROPOSITION A.6. *Let $q = C_1(x_1) \wedge \dots \wedge C_m(x_m)$ be a conjunction of basic singletons sjfBCQ and $S \subseteq [m]$. Then, given an incomplete database D as input, we can compute $N_S(D)$ in polynomial time.*

PROOF. First, observe that to compute $N_S(D)$, we can assume without loss of generality that the input database D only contains facts over relation names that occur in some $C_i(x_i)$, for $i \in S$. Indeed, $N_S(D)$ counts the valuations v of D that do not satisfy any of the $C_i(x_i)$ for $i \in S$, so for any $j \notin S$, we do not care if v satisfies $C_j(x_j)$ or not; hence, we could simply multiply the result by the appropriate factor. Therefore, we can assume that S is $[m]$. We now need to fix some notation. Let us write the conjunction of basic singleton sjfBCQ q as

$$R_1(x_1) \wedge \dots \wedge R_{m_1}(x_1) \wedge R_{m_1+1}(x_2) \wedge \dots \wedge R_{m_1+m_2}(x_2) \wedge \dots \wedge R_{\sum_{i=1}^{m-1} m_i}(x_m) \wedge \dots \wedge R_{\sum_{i=1}^m m_i}(x_m)$$

and let K be the number of atoms in q , that is, $K \stackrel{\text{def}}{=} \sum_{i=1}^m m_i$. Let dom be the uniform domain of the nulls occurring in D and d its size. For $\mathbf{s} \subseteq [K]$, we write $C_{\mathbf{s}}$ the set of constants that occur in each of the relations $D(R_i)$ for $i \in \mathbf{s}$ but in none of the others, and write $c_{\mathbf{s}}$ the size of that set. We call such a set a *block* of constants. Similarly for the nulls, we write $N_{\mathbf{s}}$ the set of nulls that occur in each of the relations $D(R_i)$ for $i \in \mathbf{s}$ but in none of the others (and we call this a block of nulls), and $n_{\mathbf{s}}$ for its size. We can assume wlog that:

- (a) For every $1 \leq i \leq m$, there is no constant that occurs in every $D(R)$ for R a relation name in $C_i(x_i)$. Indeed otherwise any valuation would satisfy $C_i(x_i)$, thus $N_{[m]}(D)$ would simply be 0.
- (b) Every constant c appearing in D is in dom . Indeed otherwise, with the last item, this constant would have no chance to be part of a match, so we could simply remove it (i.e., remove all tuples of the form $R(c)$ from D).

For a subset $A \subseteq \text{dom}$, let us write $A^C \stackrel{\text{def}}{=} \text{dom} \setminus A$. Finally, for a set $Z = \{A_1, \dots, A_l\}$ of subsets of dom , we denote by $\mathcal{I}(Z)$ the set

$$\mathcal{I}(Z) \stackrel{\text{def}}{=} \left\{ \bigcap_{i=1}^l B_i \mid (B_1, \dots, B_l) \in \{A_1, A_1^C\} \times \dots \times \{A_l, A_l^C\} \right\}$$

We now explain informally how we can compute $N_{[m]}(D)$. Let $L = \mathbf{s}_1, \dots, \mathbf{s}_{2^K}$ be an arbitrary linear order of the set of subsets of $[K]$. We will define by induction on $i \in [2^K]$ an expression computing $N_{[m]}(D)$, which will be a nested sum of the form

$$\sum_{\text{sum}_{\mathbf{s}_1}} f_{\mathbf{s}_1} \times \left(\sum_{\text{sum}_{\mathbf{s}_2}} f_{\mathbf{s}_2} \times \left(\dots \left(\sum_{\text{sum}_{\mathbf{s}_{2^K}}} f_{\mathbf{s}_{2^K}} \right) \dots \right) \right) \quad (7)$$

where each $\text{sum}_{\mathbf{s}_i}$ sums over the possible images $A_{\mathbf{s}_i}$ of the nulls in $N_{\mathbf{s}_i}$ by a valuation, and $f_{\mathbf{s}_i}$ will simply be $\text{surj}_{N_{\mathbf{s}_i} \rightarrow A_{\mathbf{s}_i}}$, where $a_{\mathbf{s}_i} \stackrel{\text{def}}{=} |A_{\mathbf{s}_i}|$, i.e., the number of valuations ν of $N_{\mathbf{s}_i}$ with image exactly $A_{\mathbf{s}_i}$. But there are two technicalities:

- First, we need to ensure that each basic singleton query $C_i(x_i)$ of q will not be satisfied. In order to do that, $\text{sum}_{\mathbf{s}_i}$ will actually sum over all the possible partitions $(B_{\mathbf{s}_1}^1, \dots, B_{\mathbf{s}_1}^{\mathcal{I}(Z_{i-1})})$ of $A_{\mathbf{s}_i}$, where each of the $B_{\mathbf{s}_i}^j$ is included in one of the sets in $\mathcal{I}(Z_{i-1})$, where Z_{i-1} contains all the blocks of constants and all the other $B_{\mathbf{s}_j}^k$ for $j < i$. We iteratively build that sum from the outside to the inside, starting with $Z_0 \stackrel{\text{def}}{=} \{\text{dom}\} \cup \{C_{\mathbf{s}} \mid \mathbf{s} \subseteq [K]\}$. This will allow us to avoid summing over the $B_{\mathbf{s}_1}^j$ that would render a basic singleton query true.
- Second, as is, such a sum is obviously not going to be computable in PTIME, as we are summing over subsets of dom . To fix this, observe that the value of the subsum for \mathbf{s}_i actually only depends on the *sizes* of the sets in Z_{i-1} . Hence, iterating from the outside to the inside, whenever $\text{sum}_{\mathbf{s}_i}$ contains a sum of the form, say, $B_{\mathbf{s}_i}^k \subseteq B_{\mathbf{s}_j}^{k'}$ for $j < i$, we can replace this with a sum over $0 \leq b_{\mathbf{s}_i}^k \leq b_{\mathbf{s}_j}^{k'}$, and add to $f_{\mathbf{s}_i}$ a factor of $\binom{b_{\mathbf{s}_j}^{k'}}{b_{\mathbf{s}_i}^k}$. Now, because of how Z_0 is defined, and because of how \mathcal{I} works, all the initial numbers in the first sum are either $|\text{dom} \setminus \bigcup_{i=1}^K C_{\{i\}}|$ or one of the numbers $c_{\mathbf{s}}$ for $\mathbf{s} \subseteq [K]$. These can all be computed in polynomial time.

The resulting expression then indeed evaluates to $N_{[m]}(D)$, and is in a form that allows us to directly compute it in polynomial time (but non-elementary in the query). This concludes the proof of Proposition A.6. \square

B PROOFS FOR SECTION 4 (DICHOTOMIES FOR COUNTING COMPLETIONS)

B.1 Proof for Proposition 4.7

In this section, we explain how to obtain the following hardness result.

PROPOSITION 4.7 (IMPLIED BY [26]). *The problem #PF restricted to bipartite graphs is #P-hard.*

This result is proven for (non-necessarily bipartite) graphs in [26] using techniques from matroid theory, in particular using the notions of *bicircular matroid* of a graph and of *Tutte polynomial* of a matroid. We did not find a way to show that the result holds on bipartite graphs without explaining their proof for general graphs, and we did not find a way to explain the proof for general graphs without introducing these concepts. Therefore, we need to define these concepts here. We have tried to keep this exposition as brief as possible, but more detailed introductions to matroid theory and to the Tutte polynomial can be found in [43, 54]. First, we define what is a matroid.

Definition B.1. A matroid $M = (E, \mathcal{I})$ is a pair where E is a finite set (called the *ground set*) and \mathcal{I} is a set of subsets of E whose elements are called *independent sets* and that satisfies the following properties:

Non emptiness. $\mathcal{I} \neq \emptyset$;

Heritage. For every $A' \subseteq A \subseteq E$, if $A \in \mathcal{I}$ then $A' \in \mathcal{I}$;

Independent set exchange. For every $A, B \in \mathcal{I}$, if $|A| > |B|$ then there exists $x \in A \setminus B$ such that $B \cup \{x\} \in \mathcal{I}$.

In a matroid $M = (E, \mathcal{I})$, an independent set $A \in \mathcal{I}$ is called a *basis* if every strict superset $A \subsetneq A' \subseteq E$ is not in \mathcal{I} . Notice that, thanks to the independent set exchange property, all bases of M have the same number of elements. The *rank* of M is defined as the number of elements in any basis of M . Given a matroid $M = (E, \mathcal{I})$ and $A \subseteq E$, we can define the *submatroid* of M generated by A to be $M_A = (A, \mathcal{I}')$, where for $A' \subseteq A$, we have $A' \in \mathcal{I}'$ iff $A' \in \mathcal{I}$ (one should check that this is indeed a matroid). The *rank function* $\text{rk}_M : \{A \mid A \subseteq E\} \rightarrow \mathbb{N}$ of M is then defined with $\text{rk}_M(A)$ being the rank of the matroid M_A . We will now omit the subscript in rk_M as this will not cause confusion. We are ready to define the Tutte polynomial of a matroid.

Definition B.2. Let $M = (E, \mathcal{I})$ be a matroid. The *Tutte polynomial* of M , denoted $T(M; x, y)$, is the two-variables polynomial defined by

$$T(M; x, y) = \sum_{A \subseteq E} (x-1)^{\text{rk}(M)-\text{rk}(A)} (y-1)^{|A|-\text{rk}(A)}$$

We will use the following observation:

OBSERVATION B.3. *Let $M = (E, \mathcal{I})$ be a matroid. Then $T(M; 2, 1) = |\mathcal{I}|$, i.e., evaluating the Tutte polynomial of a matroid at point $(2, 1)$ simply counts its number of independent sets.*

PROOF. We have $T(M; 2, 1) = \sum_{A \subseteq E} 0^{|A|-\text{rk}(A)}$. We recall the convention that $0^0 = 1$, and the fact that $0^k = 0$ for $k > 0$. Observe then that we always have $\text{rk}(A) \leq |A|$, and that we have $\text{rk}(A) = |A|$ if and only if $A \in \mathcal{I}$, which proves the claim. \square

Next, we define what is called the *bicircular matroid* of a graph $G = (V, E)$. Recall from Definition 4.6 the definition of the induced subgraph $G[S]$ for $S \subseteq E$.

Definition B.4. Let $G = (V, E)$ be a graph and $\mathcal{I} = \{S \subseteq E \mid G[S] \text{ is a pseudoforest}\}$. Then one can check that (E, \mathcal{I}) is a matroid [56]. This matroid is called the *bicircular matroid* of G , and is denoted by $B(G)$.

Notice then that the problem #PF is exactly the same as the problem of computing, given as input a graph G , the quantity $T(B(G); 2, 1)$. We now explain the steps used in [26] to prove that computing $T(B(G); 2, 1)$ is #P-hard for graphs. The starting point of our explanation is that computing $T(B(G); 1, 1)$ is #P-hard.

PROPOSITION B.5 ([26, COROLLARY 4.3]). *The problem of computing, given a graph G , the quantity $T(B(G); 1, 1)$ is #P-hard.*

Second, let us define the following univariate polynomial: for a graph G , let $P_G(x)$ be

$$P_G(x) = T(B(G); x, 1).$$

Notice that this is indeed a polynomial and that its degree is at most $|E|$ (the degree is exactly $|E|$ iff G is itself a pseudoforest). If we could compute efficiently the coefficients of P_G , then we could in particular compute the value $P_G(1) = T(B(G); 1, 1)$, which is #P-hard by the previous proposition. We recall that to compute the coefficients of a polynomial of degree n , it is enough to know its value on $n + 1$ distinct points; in fact, given these values in $n + 1$ distinct points, it is possible to efficiently compute the coefficients of the polynomial by using standard interpolation techniques (for example, by using Lagrange polynomials).

We need one last definition.

Definition B.6. Let G be a graph. For $k \in \mathbb{N}$, let $s_k(G)$ be the graph obtained from G by replacing each edge of G by a path of length k ; this graph is called the k -stretch of G .

Then, using a result attributed to Brylawski (see Reference [32]), the authors of [26] obtain that, “up to a trivial factor,” we have

$$T(B(s_k(G)); 2, 1) \simeq T(B(G); 2^k, 1).$$

A careful inspection of [32] reveals¹³ that, in fact, we have

$$T(B(s_k(G)); 2, 1) = (2^k - 1)^{|E| - \text{rk}_{B(G)}(E)} \times T(B(G); 2^k, 1).$$

Notice that $\text{rk}_{B(G)}(E)$ is the size (number of edges) of a pseudoforest of G that is maximal by inclusion of edges, which we can compute in polynomial time.¹⁴

With this, the authors of [26] can conclude the proof that computing $T(B(G); 2, 1)$ is hard for (non-necessarily bipartite) graphs, i.e., that #PF is #P-hard. Indeed, given as input $G = (V, E)$, we can construct in polynomial time the graphs $s_k(G)$ for $|E| + 1$ distinct values of k , then use oracle calls to obtain the numbers $T(B(s_k(G)); 2, 1)$, which gives us the value of P_G on $|E| + 1$ distinct points. With that, we can recover the coefficients of P_G and compute $P_G(1) = T(B(G); 1, 1)$ as argued above, thus proving hardness for general graphs. To obtain hardness for bipartite graphs, it is enough to observe that when k is even then the k -stretch of G is bipartite (even if G is not bipartite). Hence, to obtain a proof of Proposition 4.7 for bipartite graphs, we can simply change that proof and specify that we make $|E| + 1$ calls to the oracle $T(B(s_k(G)); 2, 1)$ for $|E| + 1$ distinct even values of k .

¹³To be precise, we use Equations (7.1) and (7.2) of [32] with $x = 1$, $y = 0$, and Equation (2.2) with $x = 2$, $y = 1$.

¹⁴This is because, since $B(G)$ is a matroid, any two such pseudoforests have the same number of edges. We can then simply start from the empty subgraph and iteratively add edges until it is not possible to add an edge such that the resulting graph is a pseudoforest. This also relies on the fact that we can check in polynomial time whether a graph is a pseudoforest.

B.2 Proof of Theorem 4.10

In this section, we prove the tractability part of the following:

THEOREM 4.10 (DICHOTOMY). *Let q be an sjfBCQ. If $R(x, x)$ or $R(x, y)$ is a pattern of q , then $\#Comp^u(q)$ and $\#Comp_{Cd}^u(q)$ are $\#P$ -hard. Otherwise, these problems are in FP.*

Let q be an sjfBCQ not containing any of these two patterns. Then, as observed in Section 4.2, q is a conjunction of basic singletons query. Let $\sigma = \{R_1, \dots, R_l\}$ be the set of relation symbols of q , and D be an incomplete database over these relations, with dom the uniform domain of the nulls and d its size. For every $s \subseteq \sigma$, $s \neq \emptyset$, let:

- C_s be the set of constants that occur in all relations of s and in none of the others; c_s be its size;
- N_s be the set of nulls that occur in all relations of s and in none of the others; n_s be its size.

We also define c as $\sum_{\emptyset \neq s \subseteq \sigma} c_s$. We can assume w.l.o.g. that $C_s \subseteq \text{dom}$ for all $\emptyset \neq s \subseteq \sigma$, otherwise, we can simply remove from D the corresponding facts. Let $L \stackrel{\text{def}}{=} 2^l - 1$, and let s_1, \dots, s_L be an arbitrary linear order of $\{s \subseteq \sigma \mid s \neq \emptyset\}$ (for instance, by non-decreasing size). We will follow the same steps as in the example of Section 4.3. The following lemma is the generalization of Claim 4.12, and explains how we can guide the computation so we do not count the same completion twice:

LEMMA B.7. *For a tuple $(I_{s_1}, \dots, I_{s_L})$ of subsets of dom satisfying (\star)*

$$I_s \subseteq (\text{dom} \setminus (C \cup \bigcup_{\substack{\emptyset \neq s' \subseteq \sigma \\ s' \neq s}} I_{s'})) \cup \bigcup_{\emptyset \neq s' \subseteq s} C_{s'}$$

for every $s \in (s_1, \dots, s_L)$ (in other words, all the sets I_s are mutually disjoint subsets of dom , and a set I_s can only contain a constant $b \in C$ if b is in one of the sets $C_{s'}$ for which s' is strictly included in s), let us define $P(I_{s_1}, \dots, I_{s_L})$ to be the complete database consisting of the following facts, for every $\emptyset \neq s \subseteq \sigma$:

- $R(a)$ for every $R \in s$ and $a \in I_s$ or $a \in C_s \setminus \bigcup_{s \supseteq s'} I_{s'}$

Then, for every two such tuples $(I_{s_1}, \dots, I_{s_L})$ and $(I'_{s_1}, \dots, I'_{s_L})$ satisfying (\star) and that are distinct, we have that $P(I_{s_1}, \dots, I_{s_L}) \neq P(I'_{s_1}, \dots, I'_{s_L})$.

PROOF. Let us write $P = P(I_{s_1}, \dots, I_{s_L})$ and $P' = P(I'_{s_1}, \dots, I'_{s_L})$. Assume that $P = P'$, and let us show that $(I_{s_1}, \dots, I_{s_L}) = (I'_{s_1}, \dots, I'_{s_L})$. Assume by way of contradiction that for some $\emptyset \neq s \subseteq \sigma$ we have $I_s \neq I'_s$. Then (w.l.o.g.) there exists $a \in I_s \setminus I'_s$. By the definition of P , we have that P contains all the facts $R(a)$ for $R \in s$. Let us show that P does not contain any fact $R(a)$ for $R \notin s$. Otherwise, assume that P contains $R(a)$ with $R \notin s$. Then there exists $s' \subseteq \sigma$ such that $R \in s'$ and such that $a \in I_{s'} \cup (C_{s'} \setminus \bigcup_{s' \supseteq s''} I_{s''})$. Since s does not contain R while s' does, we have $s' \not\subseteq s$. But then by (\star) , we have that I_s and $I_{s'} \cup C_{s'}$ are disjoint, which is a contradiction because a is supposed to be in both I_s and $I_{s'} \cup (C_{s'} \setminus \bigcup_{s' \supseteq s''} I_{s''})$. Therefore, it is indeed the case that P does not contain any fact $R(a)$ for $R \notin s$. Now, if P' contains a fact $R(a)$ for some $R \notin \sigma$, then we are done, since this would imply $P \neq P'$, a contradiction. Hence, we can assume that P' does not contain any fact $R(a)$ for $R \notin \sigma$. We will now prove that P' does not contain all the facts $R(a)$ for $R \in \sigma$, thus establishing a contradiction (because P does, so we would have $P \neq P'$) and concluding this proof. Assume by contradiction that P' contains all the facts $R(a)$ for $R \in s$. First, observe that we have $a \notin C_s$ because by (\star) , we have that I_s and C_s are disjoint, and we know that $a \in I_s$. Hence, the only way in which P' could contain all the facts $R(a)$ for $R \in s$ is if there exist s'_1, \dots, s'_k with $k \geq 1$ and $s'_j \subsetneq s$ for $1 \leq j \leq k$ such that $\bigcup_{1 \leq j \leq k} s'_j = s$ and such that for every $1 \leq j \leq k$, we have that (i) $a \in I_{s'_j} \cup (C_{s'_j} \setminus \bigcup_{s'_j \supseteq s''} I_{s''})$. Observe that there must exist $1 \leq j_1, j_2 \leq k$ such that s'_{j_1}

and s'_{j_2} are incomparable by inclusion (otherwise, since all s_j are strictly included in s , their union could not be equal to s). Also observe that by (\star) , we have that the sets $I_{s'_{j_1}} \cup C_{s'_{j_1}}$ and $I_{s'_{j_2}} \cup C_{s'_{j_2}}$ must be disjoint. But then (i) applied to j_1 and j_2 gives a contradiction (namely, these two sets are not disjoint, since they both contain a). This finishes the proof. \square

This next Lemma generalizes Claim 4.13 and tells us that by summing over all such tuples $(I_{s_1}, \dots, I_{s_L})$, we cannot miss a completion of D :

LEMMA B.8. *Let D' be a completion of D . Then there exists a tuple $(I_{s_1}, \dots, I_{s_L})$ of subsets of dom satisfying (\star) such that $D' = P(I_{s_1}, \dots, I_{s_L})$.*

PROOF. For $\emptyset \neq s \subseteq \sigma$, let us define D_s to be the set of constants that occur in all relation of s and in none of the others. Define the set I_s for $\emptyset \neq s \subseteq \sigma$ as follows: $I_s \stackrel{\text{def}}{=} D_s \setminus C_s$. It is then routine to check that $(I_{s_1}, \dots, I_{s_L})$ satisfies (\star) and is such that $D' = P(I_{s_1}, \dots, I_{s_L})$. \square

Lemma B.7 and B.8 allows us to express the result as

$$\sum_{I_{s_1} \subseteq \text{dom} \setminus C} \dots \sum_{I_{s_j} \subseteq (\text{dom} \setminus (C \cup \bigcup_{1 \leq k < j} I_{s_k})) \cup \bigcup_{\emptyset \neq s' \subseteq s} C_{s'}} \dots \sum_{I_{s_L} \subseteq \text{dom} \setminus (C_{s_L} \cup \bigcup_{1 \leq k < L} I_{s_k})} \text{check}(I_{s_1}, \dots, I_{s_L}) \quad (8)$$

where $\text{check}(I_{s_1}, \dots, I_{s_L}) \in \{0, 1\}$ is defined by

$$\text{check}(I_{s_1}, \dots, I_{s_L}) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } P(I_{s_1}, \dots, I_{s_L}) \text{ is a completion of } D \text{ that satisfies } q \\ 0 & \text{otherwise} \end{cases}$$

As such, we cannot evaluate this expression in P. The next step is to show that the value of $\text{check}(I_{s_1}, \dots, I_{s_L})$ only depends on $(|I_{s_1}|, \dots, |I_{s_L}|)$, which would allow us to rewrite the result as

$$\sum_{0 \leq i_{s_1}, \dots, i_{s_L} \leq d} \prod_{1 \leq j \leq L} \binom{d - c - \sum_{1 \leq k < j} i_{s_k} + \sum_{\emptyset \neq s' \subseteq s} C_{s'}}{i_{s_j}} \times \text{check}(i_{s_1}, \dots, i_{s_L}) \quad (9)$$

We give here the necessary and sufficient conditions for $P(I_{s_1}, \dots, I_{s_L})$ to be a completion of D that satisfies q .

LEMMA B.9. *We have $\text{check}(I_{s_1}, \dots, I_{s_L}) = 1$ if and only if the following conditions hold:*

- (1) *for every basic singleton query $C_i(x)$ of q , letting s be its sets of relation symbols, there exists $s \subseteq s' \subseteq \sigma$ such that we have $|I_{s'}| \geq 1$ or $c_{s'} \geq 1$.*
- (2) *for every $\emptyset \neq s \subseteq \sigma$, if $n_s \geq 1$ and $|\bigcup_{s' \supseteq s} C_{s'} \cup \bigcup_{s' \supseteq s} I_{s'}| = 0$ then $|I_s| \neq 0$.*
- (3) *consider the following system of equations, with integer variables between 0 and d :*

- *for every two sets A, A' of subsets of $\{\emptyset \neq s \subseteq \sigma\}$, we have a variable $z_{N_s}^{A, A'}$ for every $s \in A$ and a variable $z_{C_s}^{A, A'}$ for every $s \in A'$. For instance, if $\sigma = \{R, S, T, U\}$ and if $A = \{\{R, S\}, \{S, T\}\}$ and $A' = \{\{U\}\}$, then we have the variables*

$$z_{N_{\{R, S\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}} \text{ and } z_{N_{\{S, T\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}} \text{ and } z_{C_{\{U\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}}.$$

The intuition is that we will use $z_{N_{\{R, S\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}}$ of the nulls in $N_{\{R, S\}}$ and combine them with $z_{N_{\{S, T\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}}$ of the nulls in $N_{\{S, T\}}$ and with $z_{C_{\{U\}}}^{\{\{R, S\}, \{S, T\}\}, \{\{U\}\}}$ of the constants in $C_{\{U\}}$ to obtain constants in $I_{\{R, S, T, U\}}$. Let us write V this set of variables. (we note here that we are using slightly different notation than for the example in Section 4.3; this is for readability reasons only.)

- Now, for every $\emptyset \neq s \subseteq \sigma$ we have the constraint

$$\sum_{z_{N_s}^{A,A'} \in V} z_{N_s}^{A,A'} \leq n_s$$

as well as the constraint

$$\sum_{z_{C_s}^{A,A'} \in V} z_{C_s}^{A,A'} \leq c_s$$

intuitively expressing that we do not use more nulls and constants than there are available.

- for every $\emptyset \neq s \subseteq \sigma$ we have a constraint

$$\sum_{\substack{A,A' \subseteq \{\emptyset \neq s \subseteq \sigma\} \\ A \cup A' = s}} \min_{z_*^{A,A'} \in V} z_*^{A,A'} \geq I_s$$

intuitively meaning that we have allocated the groups of nulls and constants in a way that allows us to fill the set I_s .

Then this system of equations must have a solution.

PROOF. The idea is the same as in Claim 4.14. The only difference is that we added condition (1), which ensures that the guessed completion indeed satisfies the query. \square

As in the example of Section 4.3, this implies that the value of $\text{check}(I_{s_1}, \dots, I_{s_l})$ only depends on $(|I_{s_1}|, \dots, |I_{s_l}|)$ and can be computed in FP (by testing all assignments of the z_* variables; because the schema is fixed so there are only a fixed number of such variables). But then we can compute the result in FP by evaluating the expression 9, which finishes the proof.

ACKNOWLEDGMENTS

We thank Antoine Amarilli for suggesting to use #BIS in the proof of Proposition 3.11, as well as the anonymous reviewers for their careful proofreading.

REFERENCES

- [1] Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Vol. 8. Addison-Wesley Reading. Retrieved from <http://webdam.inria.fr/Alice/>.
- [2] Serge Abiteboul, Paris Kanellakis, and Gösta Grahne. 1991. On the representation and querying of sets of possible worlds. *Theoret. Comput. Sci.* 78, 1 (1991), 159–187. Retrieved from <https://www.sciencedirect.com/science/article/pii/0304397551900072>.
- [3] Adam. 2013. Number of surjective functions. Retrieved from <https://math.stackexchange.com/a/615136/378365>.
- [4] Parag Agrawal, Anish Das Sarma, Jeffrey Ullman, and Jennifer Widom. 2010. Foundations of uncertain-data integration. *Proc. VLDB Endow.* 3, 1-2 (2010), 1080–1090. Retrieved from <http://ilpubs.stanford.edu:8090/846/2/main.pdf>.
- [5] Carme Álvarez and Birgit Jenner. 1993. A very hard log-space counting class. *Theoret. Comput. Sci.* 107, 1 (1993), 3–30. Retrieved from <https://www.sciencedirect.com/science/article/pii/0304397593902520>.
- [6] Periklis Andritsos, Ariel Fuxman, and Renee J. Miller. 2006. Clean answers over dirty databases: A probabilistic approach. In *Proceedings of the 22nd International Conference on Data Engineering (ICDE'06)*. IEEE, 30–30. Retrieved from <ftp://ftp.cs.toronto.edu/csrg-technical-reports/513/tr513.pdf>.
- [7] Lyublena Antova, Christoph Koch, and Dan Olteanu. 2009. $10^{(10^6)}$ worlds and beyond: Efficient representation and processing of incomplete information. *VLDB J.* 18, 5 (2009), 1021–1040. Retrieved from <https://arxiv.org/abs/cs/0606075>.
- [8] Marcelo Arenas, Pablo Barceló, and Mikaël Monet. 2020. Counting problems over incomplete databases. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 165–177. Retrieved from <https://arxiv.org/abs/1912.11064>.

- [9] Marcelo Arenas, Leopoldo Bertossi, and Jan Chomicki. 1999. Consistent query answers in inconsistent databases. In *Proceedings of the 18th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 68–79. Retrieved from <http://marenas.sitios.ing.uc.cl/publications/pods99.pdf>.
- [10] Marcelo Arenas, Luis Alberto Croquevielle, Rajesh Jayaram, and Cristian Riveros. 2019. Efficient logspace classes for enumeration, counting, and uniform generation. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 59–73. Retrieved from <https://arxiv.org/abs/1906.09226>.
- [11] Omar Benjelloun, Anish Das Sarma, Alon Halevy, Martin Theobald, and Jennifer Widom. 2008. Databases with uncertainty and lineage. *VLDB J.* 17, 2 (2008), 243–264. Retrieved from <http://ilpubs.stanford.edu:8090/811/1/2007-26.pdf>.
- [12] Leopoldo Bertossi. 2011. Database repairing and consistent query answering. *Synth. Lect. Data Manag.* 3, 5 (2011), 1–121. Retrieved from <https://www.cs.ubc.ca/~laks/cpsc504/dc-leo.pdf>.
- [13] Leopoldo Bertossi. 2019. Database repairs and consistent query answering: Origins and further developments. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 48–58. Retrieved from <https://dl.acm.org/citation.cfm?id=3322190>.
- [14] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. 2012. Holographic reduction, interpolation and hardness. *Computat. Complex.* 21, 4 (2012), 573–604. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.145.8882&rep=rep1&type=pdf>.
- [15] Marco Calautti, Marco Console, and Andreas Pieris. 2019. Counting database repairs under primary keys revisited. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 104–118. Retrieved from <https://core.ac.uk/download/pdf/224804156.pdf>.
- [16] T. Codd. 1975. Understanding relations (installment# 7). *FDT Bull. ACM SIGMOD* 7 (1975), 23–28.
- [17] Marco Console, Paolo Guagliardo, and Leonid Libkin. 2016. Approximations and refinements of certain answers via many-valued logics. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*. 349–358. Retrieved from <https://homepages.inf.ed.ac.uk/libkin/papers/kr16.pdf>.
- [18] Nilesh Dalvi, Christopher Re, and Dan Suciu. 2011. Queries and materialized views on probabilistic databases. *J. Comput. Syst. Sci.* 77, 3 (2011), 473–490. Retrieved from <https://www-cs.stanford.edu/~chrismre/papers/jcss-probdb.pdf>.
- [19] Nilesh Dalvi and Dan Suciu. 2013. The dichotomy of probabilistic inference for unions of conjunctive queries. *J. ACM* 59, 6 (2013), 1–87. Retrieved from <https://homes.cs.washington.edu/~suciu/jacm-dichotomy.pdf>.
- [20] Martin Dyer, Alan Frieze, and Mark Jerrum. 2002. On counting independent sets in sparse graphs. *SIAM J. Comput.* 31, 5 (2002), 1527–1541. Retrieved from <http://yaroslavvb.com/papers/dyer-on.pdf>.
- [21] Jack Edmonds. 1965. Paths, trees, and flowers. *Can. J. Math.* 17 (1965), 449–467. Retrieved from https://math.nist.gov/~JBernal/p_t_f.pdf.
- [22] Genghua Fan, Yan Li, Ning Song, and Daqing Yang. 2015. Decomposing a graph into pseudoforests with one having bounded degree. *J. Combinat. Theor., Series B* 115 (2015), 72–95. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0095895615000581>.
- [23] Wenfei Fan, Floris Geerts, Xibei Jia, and Anastasios Kementsietsidis. 2008. Conditional functional dependencies for capturing data inconsistencies. *ACM Trans. Datab. Syst.* 33, 2 (2008), 1–48. Retrieved from <https://www.inf.ed.ac.uk/publications/online/0949.pdf>.
- [24] Stephen A. Fenner, Lance J. Fortnow, and Stuart A. Kurtz. 1994. Gap-definable counting classes. *J. Comput. Syst. Sci.* 48, 1 (1994), 116–148. Retrieved from <https://www.sciencedirect.com/science/article/pii/S002200005800248>.
- [25] Amélie Gheerbrant and Cristina Sirangelo. 2019. Best answers over incomplete data: Complexity and first-order rewritings. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*. Retrieved from <https://www.ijcai.org/proceedings/2019/236>.
- [26] Omer Giménez and Marc Noy. 2006. On the complexity of computing the Tutte polynomial of bicircular matroids. *Combinat., Probab. Comput.* 15, 3 (2006), 385–395. Retrieved from <https://www.cambridge.org/core/journals/combinatorics-probability-and-computing/article/on-the-complexity-of-computing-the-tutte-polynomial-of-bicircular-matroids/8B45505EDDDF91337E62D45B143EC5E5>.
- [27] Logan Grout and Benjamin Moore. 2019. On decomposing graphs into forests and pseudoforests. *arXiv preprint arXiv:1904.12435* (2019).
- [28] Sanjay Gupta. 1991. The power of witness reduction. In *Proceedings of the 6th Structure in Complexity Theory Conference*. IEEE, 43–59. Retrieved from <https://ieeexplore.ieee.org/abstract/document/160242>.
- [29] Sanjay Gupta. 1995. Closure properties and witness reduction. *J. Comput. Syst. Sci.* 50, 3 (1995), 412–432. Retrieved from <https://www.sciencedirect.com/science/article/pii/S002200008571032X>.
- [30] Tomasz Imielinski and Witold Lipski Jr. 1984. Incomplete information in relational databases. *J. ACM* 31, 4 (1984), 761–791. Retrieved from <https://cs.uwaterloo.ca/~david/cs848s14/il84.pdf>.
- [31] Neil Immerman. 2012. *Descriptive Complexity*. Springer Science & Business Media. Retrieved from https://people.cs.umass.edu/~immerman/book/ch0_1_2.pdf.

- [32] François Jaeger, Dirk L. Vertigan, and Dominic J. A. Welsh. 1990. On the computational complexity of the Jones and Tutte polynomials. In *Math. Proc. Cambridge Phil. Soc.*, Vol. 108. Cambridge University Press, 35–53. Retrieved from <https://www.cambridge.org/core/journals/mathematical-proceedings-of-the-cambridge-philosophical-society/article/on-the-computational-complexity-of-the-jones-and-tutte-polynomials/0EA052341269A2816C36B15380B8AA02>.
- [33] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. 1986. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.* 43 (1986), 169–188. Retrieved from <http://www2.stat.duke.edu/~scs/Courses/Stat376/Papers/ConvergeRates/RandomizedAlgs/JerrumValiantVaziraniTCS1986.pdf>.
- [34] Ker-I. Ko. 1982. Some observations on the probabilistic algorithms and NP-hard problems. *Inf. Process. Lett.* 14, 1 (1982), 39–43. Retrieved from <https://www.sciencedirect.com/science/article/pii/0020019082901399>.
- [35] Johannes Köbler, Uwe Schöning, and Jacobo Torán. 1989. On counting and approximation. *Acta Inform.* 26, 4 (1989), 363–379. Retrieved from https://www.researchgate.net/publication/226508658_On_counting_and_approximation.
- [36] Łukasz Kowalik. 2006. Approximation scheme for lowest outdegree orientation and graph density measures. In *Proceedings of the International Symposium on Algorithms and Computation*. Springer, 557–566. Retrieved from <https://www.mimuw.edu.pl/~kowalik/papers/orient.pdf>.
- [37] Leonid Libkin. 2014. Incomplete data: What went wrong, and how to fix it. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 1–13. Retrieved from <https://homepages.inf.ed.ac.uk/libkin/papers/pods14.pdf>.
- [38] Leonid Libkin. 2018. Certain answers meet zero-one laws. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 195–207. Retrieved from <https://homepages.inf.ed.ac.uk/libkin/papers/pods18.pdf>.
- [39] Meena Mahajan, Thomas Thierauf, and N. V. Vinodchandran. 1994. A note on SpanP functions. *Inf. Process. Lett.* 51, 1 (1994), 7–10. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.24.9933&rep=rep1&type=pdf>.
- [40] Dany Maslowski and Jef Wijsen. 2013. A dichotomy in the complexity of counting database repairs. *J. Comput. Syst. Sci.* 79, 6 (2013), 958–983. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0022000013000214>.
- [41] Dany Maslowski and Jef Wijsen. 2014. Counting database repairs that satisfy conjunctive queries with self-joins. In *Proceedings of the International Conference on Database Theory*. 155–164. Retrieved from http://www.openproceedings.org/ICDT/2014/paper_17.pdf.
- [42] Mitsunori Ogiwara and Lane A. Hemachandra. 1993. A complexity theory for feasible closure properties. *J. Comput. Syst. Sci.* 46, 3 (1993), 295–325. Retrieved from <https://www.sciencedirect.com/science/article/pii/002200009390006I>.
- [43] James Oxley. 2003. What is a matroid. *Cubo Matem. Educac.* 5, 3 (2003), 179–218. Retrieved from <https://www.math.lsu.edu/~oxley/survey4.pdf>.
- [44] J. Scott Provan and Michael O. Ball. 1983. The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM J. Comput.* 12, 4 (1983), 777–788. Retrieved from <https://epubs.siam.org/doi/abs/10.1137/0212053>.
- [45] Raymond Reiter. 1978. *On Closed World Data Bases*. Springer US, 55–76. Retrieved from <https://pdfs.semanticscholar.org/d82a/786f460a5d5c2c6d97aa60f0ead0e70dc67e.pdf>.
- [46] Anish Das Sarma, Omar Benjelloun, Alon Halevy, Shubha Nabar, and Jennifer Widom. 2009. Representing uncertain data: Models, properties, and algorithms. *VLDB J.* 18, 5 (2009), 989–1019. Retrieved from <http://ilpubs.stanford.edu:8090/924/1/uncertainData.pdf>.
- [47] Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. 2011. *Probabilistic Databases*. Morgan & Claypool. Retrieved from <https://www.morganclaypool.com/doi/abs/10.2200/S00362ED1V01Y201105DTM016>.
- [48] Seinosuke Toda and Osamu Watanabe. 1992. Polynomial-time 1-Turing reductions from # PH to # P. *Theoret. Comput. Sci.* 100, 1 (1992), 205–221. Retrieved from <https://www.sciencedirect.com/science/article/pii/030439759290369Q>.
- [49] Leslie G. Valiant. 1976. Relative complexity of checking and evaluating. *Inf. Process. Lett.* 5, 1 (1976), 20–23. Retrieved from <https://www.sciencedirect.com/science/article/pii/0020019076900971>.
- [50] Leslie G. Valiant. 1979. The complexity of computing the permanent. *Theoret. Comput. Sci.* 8, 2 (1979), 189–201. Retrieved from <https://core.ac.uk/download/pdf/82500417.pdf>.
- [51] Leslie G. Valiant. 1979. The complexity of enumeration and reliability problems. *SIAM J. Comput.* 8, 3 (1979), 410–421. Retrieved from <https://www.math.cmu.edu/~af1p/Teaching/MCC17/Papers/enumerate.pdf>.
- [52] Ron van der Meyden. 1998. Logical approaches to incomplete information: A survey. In *Logics for Databases and Information Systems*. Springer, 307–356. Retrieved from https://link.springer.com/chapter/10.1007/978-1-4615-5643-5_10.
- [53] Moshe Y. Vardi. 1982. The complexity of relational query languages. In *Proceedings of the 14th ACM Symposium on Theory of Computing*. 137–146. Retrieved from <http://www.dis.uniroma1.it/~degiacon/didattica/semingsoft/SIS05-06/materiale/1-query-congiuntive/riferimenti/vardi-1982.pdf>.

- [54] Dominic Welsh. 1999. The Tutte polynomial. *Rand. Struct. Algor.* 15, 3-4 (1999), 210–228. Retrieved from [https://onlinelibrary.wiley.com/doi/pdf/10.1002/\(SICI\)1098-2418\(199910/12\)15:3/43C210::AID-RSA23E3.0.CO;2-R](https://onlinelibrary.wiley.com/doi/pdf/10.1002/(SICI)1098-2418(199910/12)15:3/43C210::AID-RSA23E3.0.CO;2-R).
- [55] Mingji Xia, Peng Zhang, and Wenbo Zhao. 2007. Computational complexity of counting problems on 3-regular planar graphs. *Theoret. Comput. Sci.* 384, 1 (2007), 111–125. Retrieved from <https://core.ac.uk/download/pdf/82063901.pdf>.
- [56] Thomas Zaslavsky. 1982. Bicircular geometry and the lattice of forests of a graph. *Quart. J. Math.* 33, 4 (1982), 493–511. Retrieved from <https://academic.oup.com/qjmath/article-abstract/33/4/493/1498307?redirectedFrom=PDF>.

Received November 2020; accepted April 2021