



# Probabilistic Query Evaluation with Bag Semantics

Martin Grohe  

RWTH Aachen University, Germany

Peter Lindner  

École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Christoph Standke  

RWTH Aachen University, Germany

---

## Abstract

We initiate the study of probabilistic query evaluation under *bag semantics* where tuples are allowed to be present with duplicates. We focus on self-join free conjunctive queries, and probabilistic databases where occurrences of different facts are independent, which is the natural generalization of tuple-independent probabilistic databases to the bag semantics setting. For set semantics, the data complexity of this problem is well understood, even for the more general class of unions of conjunctive queries: it is either in polynomial time, or  $\#P$ -hard, depending on the query (Dalvi & Suciu, JACM 2012).

Due to potentially unbounded multiplicities, the bag probabilistic databases we discuss are no longer finite objects, which requires a treatment of representation mechanisms. Moreover, the answer to a Boolean query is a probability distribution over non-negative integers, rather than a probability distribution over  $\{\text{true}, \text{false}\}$ . Therefore, we discuss two flavors of probabilistic query evaluation: computing expectations of answer tuple multiplicities, and computing the probability that a tuple is contained in the answer at most  $k$  times for some parameter  $k$ . Subject to mild technical assumptions on the representation systems, it turns out that expectations are easy to compute, even for unions of conjunctive queries. For query answer probabilities, we obtain a dichotomy between solvability in polynomial time and  $\#P$ -hardness for self-join free conjunctive queries.

**2012 ACM Subject Classification** Theory of computation  $\rightarrow$  Incomplete, inconsistent, and uncertain databases

**Keywords and phrases** Probabilistic Query Evaluation, Probabilistic Databases, Bag Semantics

**Digital Object Identifier** 10.4230/LIPIcs.ICDT.2023.20

**Related Version** *Extended Version*: <https://arxiv.org/abs/2201.11524> [20]

**Funding** The work of Martin Grohe and Christoph Standke has been funded by the German Research Foundation (DFG) under grants GR 1492/16-1 and GRK 2236 (UnRAVeL).

## 1 Introduction

Probabilistic databases (PDBs) provide a framework for managing uncertain data. In database theory, they have been intensely studied since the late 1990s [30, 31]. Most efforts have been directed towards tuple-independent relational databases *under a set semantics*. Many relational database systems, however, use a bag semantics, where identical tuples may appear several times in the same relation. Despite receiving little attention so far, bag semantics are also a natural setting for probabilistic databases. For example, they naturally enter the picture when aggregation is performed, or when statistics are computed (e.g., by random sampling, say, without replacement). Either case might involve computing projections without duplicate elimination first. Even when starting from a tuple-independent probabilistic database with set semantics, this typically gives rise to (proper) bags. A bag semantics for PDBs has only been considered recently [19, 18], in the context of infinite



© Martin Grohe, Peter Lindner, and Christoph Standke;  
licensed under Creative Commons License CC-BY 4.0

26th International Conference on Database Theory (ICDT 2023).

Editors: Floris Geerts and Brecht Vandevoort; Article No. 20; pp. 20:1–20:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

PDBs. Even in the traditional setting of a PDB where only finitely many facts appear with non-zero probability, under a bag semantics we have to consider infinite probability spaces, simply because there is no a priori bound on the number of times a fact may appear in a bag. In general, while the complexity landscape of query answering is well understood for simple models of PDBs under set semantics, the picture for *bag semantics* is still unexplored. In this work, we take first steps to address this.

Formally, probabilistic databases are probability distributions over conventional database instances. In a database instance, the answer to a Boolean query under *set semantics* is either **true** (1) or **false** (0). In a probabilistic database, the answer to such a query becomes a  $\{0, 1\}$ -valued random variable. The problem of interest is *probabilistic query evaluation*, that is, computing the probability that a Boolean query returns **true**, when given a probabilistic database. The restriction to *Boolean* queries comes with no loss of generality: To compute the probability of any tuple in the result of a non-Boolean query, all we have to do is replace the free variables of the query according to the target tuple, and solve the problem for the resulting Boolean query.

Under a bag semantics, a Boolean query is still just a query without free variables, but the answer to Boolean query can be any non-negative integer, which can be interpreted as the multiplicity of the empty tuple in the query answer, or more intuitively as the number of different ways in which the query is satisfied. In probabilistic query evaluation, we then get  $\mathbb{N}$ -valued answer random variables. Still, the reduction from the non-Boolean to the Boolean case works as described above. *Therefore, without loss of generality, we only discuss Boolean queries in this paper.*

As most of the database theory literature, we study the *data complexity* of query evaluation [32], that is, the complexity of the problem, when the query  $Q$  is fixed, and the PDB is the input. The standard model for complexity theoretic investigations is that of *tuple-independent* PDBs, where the distinct facts constitute independent events. Probabilistic query evaluation is well-understood for the class of *unions of conjunctive queries* (UCQs) on PDBs that are tuple-independent (see the related works section below). However, all existing results discuss the problem under set semantics. Here, on the contrary, we discuss the probabilistic query evaluation under *bag semantics*.

For tuple-independent (set) PDBs, a variety of representation systems have been proposed (cf. [17, 30]), although for complexity theoretic discussions, it is usually assumed that the input is just given as a table of facts, together with their marginal probabilities [31]. In the bag version of tuple-independent PDBs [18], different facts are still independent. Yet, the individual facts (or, rather, their multiplicities) are  $\mathbb{N}$ -valued, instead of Boolean, random variables. As this, in general, rules out the naive representation through a list of facts, multiplicities, and probabilities, it is necessary to first define suitable representation systems before the complexity of computational problems can be discussed.

Once we have settled on a suitable class of representations, we investigate the problem of probabilistic query evaluation again, subject to representation system  $\mathcal{R}$ . Under bag semantics, there are now *two* natural computational problems regarding query evaluation:  $\text{EXPECTATION}_{\mathcal{R}}(Q)$ , which is computing the expected outcome, and  $\text{PQE}_{\mathcal{R}}(Q, k)$  which is computing the probability that the outcome is at most  $k$ . Notably, these two problems coincide for set semantics, because the expected value of a  $\{0, 1\}$ -valued random variables coincides with the probability that the outcome is 1. Under a bag semantics, however, the two versions exhibit quite different properties.

Recall that using a set semantics, unions of conjunctive queries can either be answered in polynomial time, or are  $\sharp\text{P}$ -hard [12]. Interestingly, computing expectations using a bag semantics is extraordinarily easy in comparison: With only mild assumptions on the repre-

sentation, the expectation of any UCQ can be computed in polynomial time. Furthermore, the variance of the random variable can also be computed in polynomial time, which via Chebyshev’s inequality gives us a way to estimate the probability that the query answer is close to its expectation. These results contrast the usual landscape of computational problems in uncertain data management, which are rarely solvable efficiently.

The computation of probabilities of concrete answer multiplicities, however, appears to be less accessible, and in fact, in its properties is more similar to the set semantics version of probabilistic query evaluation. Our main result states that for Boolean conjunctive queries without self-joins, we have a dichotomy between polynomial time and  $\#P$ -hardness of the query. This holds whenever efficient access to fact probabilities is guaranteed by the representation system and is independent of  $k$ . Although the proof builds upon ideas and notions introduced for the set semantics dichotomy [9, 10, 12], we are confronted with a number of completely new and intricate technical challenges due to the change of semantics. On the one hand, the bag semantics turns disjunctions and existential quantification into sums. This facilitates the computation of expected values, because it allows us to exploit linearity. On the other hand, the new semantics keep us from directly applying some of the central ideas from [12] when analyzing  $\text{PQE}_{\mathcal{R}}(Q, k)$ , thus necessitating novel techniques. The bag semantics dichotomy for answer count probabilities is, hence, far from being a simple corollary from the set semantics dichotomy. From the technical perspective, the most interesting result is the transfer of hardness from  $\text{PQE}_{\mathcal{R}}(Q, 0)$  to  $\text{PQE}_{\mathcal{R}}(Q, k)$ . In essence, we need to find a way to compute the probability that  $Q$  has 0 answers, with only having access to the probability that  $Q$  has at most  $k$  answers for any single fixed  $k$ . This reduction uses new non-trivial techniques: By manipulating the input table, we can construct multiple instances of the  $\text{PQE}_{\mathcal{R}}(Q, k)$  problem. We then transform the solutions to these problems, which are obtained through oracle calls, into function values of a polynomial (with a priori unknown coefficients) in such a way, that the solution to  $\text{PQE}_{\mathcal{R}}(Q, 0)$  on the original input is hidden in the leading coefficient of this polynomial. Using a technique from polynomial interpolation, we can find these leading coefficients, and hence, solve  $\text{PQE}_{\mathcal{R}}(Q, 0)$ .

## Related Work

The most prominent result regarding probabilistic query evaluation is the Dichotomy theorem by Dalvi and Suciu [12] that provides a separation between unions of conjunctive queries for which probabilistic evaluation is possible in polynomial time, and such where the problem becomes  $\#P$ -hard. They started their investigations with self-join free conjunctive queries [10] and later extended their results to general CQs [11] and then UCQs [12]. Beyond the queries they investigate, there are a few similar results for fragments with negations or inequalities [14, 26, 27], for homomorphism-closed queries [4] and others [28], and on restricted classes of PDBs [1]. Good overviews over related results are given in [31, 29]. In recent developments, the original dichotomies for self-join free CQs, and for general UCQs have been shown to hold even under severe restrictions to the fact probabilities that are allowed to appear [2, 23].

The bag semantics for CQs we use here is introduced in [7]. A detailed analysis of the interplay of bag and set semantics is presented in [8]. Considering multiplicities as semi-ring annotations [16, 21], embeds bag semantics into a broader mathematical framework.

In recent work (independent of ours), Feng et al. [13] analyze the fine-grained complexity of computing expectations of queries in probabilistic bag databases, albeit assuming finite multiplicity supports and hence still in the realm of finite probabilistic databases.

## 2 Preliminaries

We denote by  $\mathbb{N}$  and  $\mathbb{N}_+$  the sets of non-negative, and of positive integers, respectively. We denote open, closed and half open intervals of real numbers by  $(a, b)$ ,  $[a, b]$ ,  $[a, b)$  and  $(a, b]$ , respectively, where  $a \leq b$ . By  $\binom{n}{k}$  we denote the binomial coefficient and by  $\binom{n}{n_1, \dots, n_k}$  the multinomial coefficient.

Let  $\Omega$  be a non-empty finite or countably infinite set and let  $P: \Omega \rightarrow [0, 1]$  be a function satisfying  $\sum_{\omega \in \Omega} P(\omega) = 1$ . Then  $(\Omega, P)$  is a (*discrete*) *probability space*. Subsets  $A \subseteq \Omega$  are called *events*. We write  $\Pr_{\omega \sim \Omega}(\omega \in A)$  for the probability of a randomly drawn  $\omega \in \Omega$  (distributed according to  $P$ ) to be in  $A$ . More generally, we may write  $\Pr_{\omega \in \Omega}(\omega \text{ has property } \varphi)$  for the probability of a randomly drawn element to satisfy some property  $\varphi$ . *All probability spaces appearing in this paper are discrete.*

Functions  $X: \Omega \rightarrow \mathbb{R}$  on a probability space are called *random variables*. The expected value and variance of  $X$  are denoted by  $\mathbb{E}(X)$  and  $\text{Var}(X)$ , respectively. The values  $\mathbb{E}(X^k)$  for integers  $k \geq 2$  are called the higher-order *moments* of  $X$ .

### 2.1 Probabilistic Bag Databases

We fix a countable, non-empty set  $\text{dom}$  (the *domain*). A *database schema*  $\tau$  is a finite, non-empty set of relation symbols. Every relation symbol  $R$  has an arity  $\text{ar}(R) \in \mathbb{N}_+$ .

A *fact* over  $\tau$  and  $\text{dom}$  is an expression  $R(\mathbf{a})$  where  $\mathbf{a} \in \text{dom}^{\text{ar}(R)}$ . A (*bag*) *database instance*  $D$  is a bag (i.e. multiset) of facts. Formally, a bag (instance) is specified by a function  $\#_D$  that maps every fact  $f$  to its multiplicity  $\#_D(f)$  in  $D$ . The *active domain*  $\text{adom}(D)$  is the set of domain elements  $a$  from  $\text{dom}$  for which there exists a fact  $f$  containing  $a$  such that  $\#_D(f) > 0$ .

A *probabilistic (bag) database* (or, (*bag*) *PDB*)  $\mathcal{D}$  is a pair  $(\mathbb{D}, P)$  where  $\mathbb{D}$  is a set of bag instances and  $P: 2^{\mathbb{D}} \rightarrow [0, 1]$  is a probability distribution over  $\mathbb{D}$ . Note that, even when the total number of different facts is finite,  $\mathbb{D}$  may be infinite, as facts may have arbitrarily large multiplicities. We let  $\#_D(f)$  denote the random variable  $D \mapsto \#_D(f)$  for all facts  $f$ . If  $\mathcal{D} = (\mathbb{D}, P)$  is a PDB, then  $\text{adom}(\mathcal{D}) := \bigcup_{D \in \mathbb{D}} \text{adom}(D)$ . We call a PDB *fact-finite* if the set  $\{f: \#_D(f) > 0 \text{ for some } D \in \mathbb{D}\}$  is finite. In this case,  $\text{adom}(\mathcal{D})$  is finite, too.

A bag PDB  $\mathcal{D}$  is called *tuple-independent* if for all  $k \in \mathbb{N}$ , all pairwise distinct facts  $f_1, \dots, f_k$ , and all  $n_1, \dots, n_k \in \mathbb{N}$ , the events  $\#_D(f_i) = n_i$  are independent, i. e.,

$$\Pr_{D \sim \mathcal{D}}(\#_D(f_i) = n_i \text{ for all } i = 1, \dots, k) = \prod_{i=1}^k \Pr_{D \sim \mathcal{D}}(\#_D(f_i) = n_i).$$

*Unless it is stated otherwise, all probabilistic databases we treat in this paper are assumed to be fact-finite and tuple-independent.*

### 2.2 UCQs with Bag Semantics

Let  $\mathcal{V}$  be a countably infinite set of variables. An *atom* is an expression of the shape  $R(\mathbf{t})$  where  $R \in \tau$  and  $\mathbf{t} \in (\text{dom} \cup \mathcal{V})^{\text{ar}(R)}$ . A *conjunctive query (CQ)* is a formula  $Q$  of first-order logic (over  $\tau$  and  $\text{dom}$ ) of the shape

$$Q = \exists x_1 \dots \exists x_m : R_1(\mathbf{t}_1) \wedge \dots \wedge R_n(\mathbf{t}_n),$$

in which we always assume that the  $x_i$  are pairwise different, and that  $x_i$  appears in at least one of  $\mathbf{t}_1, \dots, \mathbf{t}_n$  for all  $i = 1, \dots, m$ . A CQ  $Q$  is *self-join free*, if every relation symbol occurs at most once within  $Q$ . In general, the *self-join width* of a CQ  $Q$  is the maximum

number of repetitions of the same relation symbol in  $Q$ . If  $Q$  is a CQ of the above shape, we let  $Q^*$  denote the quantifier-free part  $R_1(\mathbf{t}_1) \wedge \dots \wedge R_n(\mathbf{t}_n)$  of  $Q$ , and we call  $R_i(\mathbf{t}_i)$  an *atom* of  $Q$  for all  $i = 1, \dots, n$ . A *union of conjunctive queries (UCQ)* is a formula of the shape  $Q = Q_1 \vee \dots \vee Q_N$  where  $Q_1, \dots, Q_N$  are CQs. A query is called *Boolean*, if it contains no free variables (that is, there are no occurrences of variables that are not bound by a quantifier). *From now on, and throughout the remainder of the paper, we only discuss Boolean (UCQ)s.*

Recall that  $\#_D$  is the multiplicity function of the instance  $D$ . The bag semantics of (UCQ)s extends  $\#_D$  to queries. For Boolean CQs  $Q = \exists x_1 \dots \exists x_m : R_1(\mathbf{t}_1) \wedge \dots \wedge R_n(\mathbf{t}_n)$  we define

$$\#_D(Q) := \sum_{\mathbf{a} \in \text{adom}(D)^m} \prod_{i=1}^n \#_D(R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}])), \quad (1)$$

where  $\mathbf{x} = (x_1, \dots, x_m)$  and  $\mathbf{a} = (a_1, \dots, a_m)$ , and  $R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}])$  denotes the fact obtained from  $R_i(\mathbf{t}_i)$  by replacing, for all  $j = 1, \dots, m$ , every occurrence of  $x_j$  by  $a_j$ . If  $Q = Q_1 \vee \dots \vee Q_N$  is a Boolean UCQ, then each of the  $Q_i$  is a Boolean CQ. We define

$$\#_D(Q) := \#_D(Q_1) + \dots + \#_D(Q_N). \quad (2)$$

Whenever convenient, we write  $\#_D Q$  instead of  $\#_D(Q)$ . We emphasize once more, that the query being Boolean *does not* mean that its answer is 0 or 1 under bag semantics, but could be any non-negative integer.

► **Remark 2.1.** We point out that in (1), conjunctions should intuitively be understood as joins rather than intersections. Our definition (1) for the bag semantics of CQs matches the one that was given in [7]. This, and the extension (2) for UCQs, are essentially special cases of how semiring annotations of formulae are introduced in the provenance semiring framework [16, 21], the only difference being that we use the active domain semantics. For UCQs however, this is equivalent since the value of (1) stays the same when the quantifiers range over arbitrary supersets of  $\text{adom}(D)$ . ◻

Note that the result  $\#_D Q$  of a Boolean UCQ on a bag instance  $D$  is a non-negative integer. Thus, evaluated over a PDB  $\mathcal{D} = (\mathbb{D}, P)$ , this yields a  $\mathbb{N}$ -valued random variable  $\#_{\mathcal{D}} Q$  with

$$\Pr(\#_{\mathcal{D}} Q = k) = \Pr_{D \sim \mathcal{D}}(\#_D Q = k).$$

► **Example 2.2.** Consider the tuple-independent bag PDB over facts  $R(a)$  and  $S(a)$ , where  $R(a)$  has multiplicity 2 or 3, both with probability  $\frac{1}{2}$ , and  $S(a)$  has multiplicity 1, 2 or 3, with probability  $\frac{1}{3}$  each. Then, the probability of the event  $\#_{\mathcal{D}}(R(a) \wedge S(a)) = 6$  is given by  $\Pr(\#_{\mathcal{D}}(R(a)) = 2) \Pr(\#_{\mathcal{D}}(S(a)) = 3) + \Pr(\#_{\mathcal{D}}(R(a)) = 3) \Pr(\#_{\mathcal{D}}(S(a)) = 2) = \frac{1}{3}$ . ◻

There are now two straight-forward ways to formulate the problem of answering a Boolean UCQ over a probabilistic database. We could either ask for the expectation  $\mathbb{E}(\#_{\mathcal{D}} Q)$ , or compute the probability that  $\#_{\mathcal{D}} Q$  is at most / at least / equal to  $k$ . These two options coincide for set semantics, as  $\#_{\mathcal{D}} Q$  is  $\{0, 1\}$ -valued in this setting.<sup>1</sup> For bag PDBs, these are two separate problems to explore. Complexity-wise, we focus on *data complexity* [32]. That is, the query (and for the second option, additionally the number  $k$ ) is a parameter of the problem, so that the input is only the PDB. Before we can start working on these problems, we first need to discuss how bag PDBs are presented as an input to an algorithm. This is the purpose of the next section.

<sup>1</sup> In fact, in the literature both approaches have been used to introduce the problem of probabilistic query evaluation [30, 31].

Relation $R$	Parameter	Relation $S$	Parameter
$R(1, 1)$	(Bernoulli, $1/2$ )	$S(1)$	(Geometric, $1/3$ )
$R(1, 2)$	(Binomial, $10, 1/3$ )	$S(2)$	(Poisson, $3$ )
$R(2, 2)$	$(0 \mapsto 1/4; 1 \mapsto 1/4; 5 \mapsto 1/2)$		

■ **Figure 1** Example of a parameterized TI representation.

### 3 Representation Systems

For the set version of probabilistic query evaluation, the default representation system represents tuple-independent PDBs by specifying all facts together with their marginal probability. The distinction between a PDB and its representation is then usually blurred in the literature. This does not easily extend to bag PDBs, as the distributions of  $\#_D(f)$  for facts  $f$  may have infinite support.

► **Example 3.1.** Let  $\mathcal{D} = (\mathbb{D}, P)$  be a bag PDB over a single fact  $f$  with multiplicity distribution  $\#_D(f) \sim \text{Geometric}(\frac{1}{2})$ , i.e.,  $\Pr_{D \sim \mathcal{D}}(\#_D(f) = k) = 2^{-k}$ . Then the instances of  $\mathcal{D}$  with positive probability are  $\{\emptyset\}, \{f\}, \{f, f\}, \dots$ , so  $\mathcal{D}$  is an infinite PDB. ◻

To use such PDBs as inputs for algorithms, we introduce a suitable class of representation systems (RS) [17]. All computational problems are then stated with respect to an RS.

► **Definition 3.2** (cf. [17]). A *representation system* (RS) for bag PDBs is a pair  $(\mathbf{T}, \llbracket \cdot \rrbracket)$  where  $\mathbf{T}$  is a non-empty set (the elements of which we call *tables*), and  $\llbracket \cdot \rrbracket$  is a function that maps every  $T \in \mathbf{T}$  to a probabilistic database  $\llbracket T \rrbracket$ . ◻

Given an RS, we abuse notation and also use  $T$  to refer to the PDB  $\llbracket T \rrbracket$ . Note that Definition 3.2 is not tailored to tuple-independence yet and requires no independence assumptions. For representing tuple-independent bag PDBs, we introduce a particular subclass of RS's where facts are labeled with the parameters of parameterized distributions over multiplicities. For example, a fact  $f$  whose multiplicity is geometrically distributed with parameter  $\frac{1}{2}$  could be annotated with  $(\text{Geometric}, 1/2)$ , representing  $\frac{1}{2}$  using two integers.

► **Definition 3.3.** A *parameterized TI representation system* (in short: TIRS) is a tuple  $\mathcal{R} = (\Lambda, \mathbf{P}, \Sigma, \mathbf{T}, \langle \cdot \rangle, \llbracket \cdot \rrbracket)$  where  $\Lambda \neq \emptyset$  is a set (the *parameter set*);  $\mathbf{P}$  is a family  $(P_\lambda)_{\lambda \in \Lambda}$  of probability distributions  $P_\lambda$  over  $\mathbb{N}$ ;  $\Sigma \neq \emptyset$  is a finite set of symbols (the *encoding alphabet*);  $\langle \cdot \rangle: \Lambda \rightarrow \Sigma^*$  is an injective function (the *encoding function*); and  $(\mathbf{T}, \llbracket \cdot \rrbracket)$  is an RS where

- $\mathbf{T}$  is the family of all finite sets  $T$  of pairs  $(f, \langle \lambda_f \rangle)$  with pairwise different facts  $f$  of a given schema and  $\lambda_f \in \Lambda$  for all  $f$ ; and
- $\llbracket \cdot \rrbracket$  maps every  $T \in \mathbf{T}$  to the tuple-independent bag PDB  $\mathcal{D}$  with multiplicity probabilities  $\Pr(\#_D f = k) = P_{\lambda_f}(k)$  for all  $(f, \langle \lambda_f \rangle) \in T$ . ◻

Whenever a TIRS  $\mathcal{R}$  is given, we assume  $\mathcal{R} = (\Lambda_{\mathcal{R}}, \mathbf{P}_{\mathcal{R}}, \Sigma_{\mathcal{R}}, \mathbf{T}_{\mathcal{R}}, \langle \cdot \rangle_{\mathcal{R}}, \llbracket \cdot \rrbracket_{\mathcal{R}})$  by default.

► **Example 3.4.** Figure 1 shows a table  $T$  from a TIRS  $\mathcal{R}$ , illustrating how the parameters can be used to encode several multiplicity distributions. Four of the distributions are standard parameterized distributions, presented using their symbolic name together with their parameters. The multiplicity distribution for  $R(2, 2)$  is a generic distribution with finite support  $\{0, 1, 5\}$ . The annotation (Binomial,  $10, 1/3$ ) of  $R(1, 2)$  in the table specifies that  $\#_T R(1, 2) \sim \text{Binomial}(10, \frac{1}{3})$ . That is,

$$\Pr(\#_T R(1, 2) = k) = \begin{cases} \binom{10}{k} (\frac{1}{3})^k (\frac{2}{3})^{10-k} & \text{if } 0 \leq k \leq 10 \text{ and} \\ 0 & \text{if } k > 10. \end{cases}$$

The multiplicity probabilities of the other facts are given analogously in terms of the Bernoulli, geometric, and Poisson distributions, respectively. The supports of the multiplicity distributions are  $\{0, 1\}$  for the Bernoulli,  $\{0, \dots, n\}$  for the Binomial, and  $\mathbb{N}$  for both the geometric and Poisson distributions (and finite sets for explicitly encoded distributions). For the first three parameterized distributions, multiplicity probabilities always stay rational if the parameters are rational. This is not the case for the Poisson distribution.  $\lrcorner$

While Definition 3.2 seems abstract, this level of detail in the encoding of probability distributions allows us to rigorously discuss computational complexity without resorting to a very narrow framework that only supports some predefined distributions. Our model also comprises tuple-independent set PDBs: The traditional representation system can be recovered from Definition 3.3 using only the Bernoulli distribution. Moreover, we remark that we can always represent facts that are present with probability 0, by just omitting them from the tables (for example, fact  $R(2, 1)$  in Figure 1).

► **Remark 3.5.** In this work, we focus on TIRS's where the values needed for computation (moments in Section 4 and probabilities in Section 5) are rational. An extension to support irrational values is possible through models of real complexity [25, 5]. A principled treatment requires a substantial amount of introductory overhead that would go beyond the scope of this paper, and which we therefore leave for future work.  $\lrcorner$

## 4 Expectations and Variances

Before computing the probabilities of answer counts, we discuss the computation of the expectation and the variance of the answer count. Recall that in PDBs without multiplicities, the answer to a Boolean query (under set semantics) is either 0 (i. e., false) or 1 (i. e., true). That is, the answer count is a  $\{0, 1\}$ -valued random variable there, meaning that its expectation coincides with the probability of the answer count being 1. Because of this correspondence, the semantics of Boolean queries on (set) PDBs are sometimes also defined in terms of the expected value [31]. For bag PDBs, the situation is different, and this equivalence no longer holds. Thus, computing expectations, and computing answer count probabilities have to receive a separate treatment. Formally, we discuss the following problems in this section:

<b>Problem</b>	<b>EXPECTATION<math>_{\mathcal{R}}(Q)</math></b>	<b>Problem</b>	<b>VARIANCE<math>_{\mathcal{R}}(Q)</math></b>
PARAMETER:	A Boolean UCQ $Q$ .	PARAMETER:	A Boolean UCQ $Q$ .
INPUT:	A table $T \in \mathbf{T}_{\mathcal{R}}$ .	INPUT:	A table $T \in \mathbf{T}_{\mathcal{R}}$ .
OUTPUT:	The expectation $\mathbb{E}(\#_T Q)$ .	OUTPUT:	The variance $\text{Var}(\#_T Q)$ .

### 4.1 Expected Answer Count

We have pointed out above that computing expected answer counts for set PDBs and set semantics is equivalent to computing the probability that the query returns true. There are conjunctive queries, for example,  $Q = \exists x \exists y: R(x) \wedge S(x, y) \wedge T(y)$ , for which the latter problem is  $\sharp\text{P}$ -hard [15, 9]. Under a set semantics, disjunctions and existential quantifiers semantically correspond to taking maximums instead of adding multiplicities. Under a bag semantics, we are now able to exploit the linearity of expectation to easily compute expected values, which was not possible under a set semantics.

► **Lemma 4.1.** *Let  $\mathcal{D}$  be a tuple-independent PDB and let  $Q$  be a Boolean CQ,  $Q = \exists x_1 \dots \exists x_m: R_1(\mathbf{t}_1) \wedge \dots \wedge R_n(\mathbf{t}_n)$ . For every  $\mathbf{a} \in \text{adom}(\mathcal{D})^m$ , we let  $F(\mathbf{a})$  denote the set of facts appearing in  $Q^*[\mathbf{x}/\mathbf{a}]$ , and for every  $f \in F(\mathbf{a})$ , we let  $\nu(f, \mathbf{a})$  denote the number of times  $f$  appears in  $Q^*[\mathbf{x}/\mathbf{a}]$ . Then*

$$\mathbb{E}(\#_{\mathcal{D}}Q) = \sum_{\mathbf{a} \in \text{adom}(\mathcal{D})^m} \prod_{f \in F(\mathbf{a})} \mathbb{E}\left(\left(\#_{\mathcal{D}}f\right)^{\nu(f, \mathbf{a})}\right). \quad (3)$$

**Proof.** By definition, we have

$$\#_{\mathcal{D}}Q = \sum_{\mathbf{a} \in \text{adom}(D)} \#_D(Q^*[\mathbf{x}/\mathbf{a}]) = \sum_{\mathbf{a} \in \text{adom}(D)} \#_D(Q^*[\mathbf{x}/\mathbf{a}])$$

for every individual instance  $D$  of  $\mathcal{D}$ . The last equation above holds because, as  $Q^*$  is assumed to contain every quantified variable,  $\#_D(Q^*[\mathbf{x}/\mathbf{a}]) = 0$  whenever the tuple  $\mathbf{a}$  contains an element that is not in the active domain of  $D$ . By linearity of expectation, we have

$$\mathbb{E}(\#_{\mathcal{D}}Q) = \sum_{\mathbf{a} \in \text{adom}(\mathcal{D})} \mathbb{E}(\#_D(Q^*[\mathbf{x}/\mathbf{a}])).$$

Recall, that  $Q^*[\mathbf{x}/\mathbf{a}]$  is a conjunction of facts  $R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}])$ . Thus,  $\#_D(\bigwedge_{i=1}^n R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}])) = \prod_{i=1}^n \#_D(R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}]))$ . Because  $\mathcal{D}$  is tuple-independent, any two facts in  $F(\mathbf{a})$  are either equal, or independent. Therefore,

$$\mathbb{E}\left(\prod_{i=1}^n \#_D R_i(\mathbf{t}_i[\mathbf{x}/\mathbf{a}])\right) = \prod_{f \in F(\mathbf{a})} \mathbb{E}\left(\left(\#_{\mathcal{D}}f\right)^{\nu(f, \mathbf{a})}\right),$$

as the expectation of a product of independent random variables is the product of their expectations. Together, this yields the expression from (3). ◀

By linearity, the expectation of a UCQ is the sum of the expectations of its CQs.

► **Lemma 4.2.** *Let  $\mathcal{D}$  be a PDB and let  $Q = \bigvee_{i=1}^N Q_i$  be a Boolean UCQ. Then we have  $\mathbb{E}(\#_{\mathcal{D}}Q) = \sum_{i=1}^N \mathbb{E}(\#_{\mathcal{D}}Q_i)$ .*

Given that we can compute the necessary moments of fact multiplicities efficiently, Lemmas 4.1 and 4.2 yield a polynomial time procedure to compute the expectation of a UCQ. The order of moments we need is governed by the self-join width of the individual CQs.

► **Definition 4.3.** A TIRS  $\mathcal{R}$  has *polynomially computable moments* up to order  $k$ , if for all  $\lambda \in \Lambda_{\mathcal{R}}$ , we have  $\sum_{n=0}^{\infty} n^k \cdot P_{\lambda}(n) < \infty$  and the function  $\langle \lambda \rangle \mapsto \sum_{n=0}^{\infty} n^{\ell} \cdot P_{\lambda}(n)$  can be computed in polynomial time in  $|\langle \lambda \rangle|$  for all  $\ell \leq k$ . ◻

Before giving the main statement, let us revisit Example 3.4 for illustration.

► **Example 4.4.** Let  $\mathcal{R}$  be the TIRS from Example 3.4. The moments of  $X \sim \text{Bernoulli}(p)$  are  $\mathbb{E}(X^k) = p$  for all  $k \geq 1$ . Direct calculation shows that for  $X \sim \text{Binomial}(n, p)$ , the moment  $\mathbb{E}(X^k)$  is given by a polynomial in  $n$  and  $p$ . In general, for most of the common distributions, one of the following cases applies. Either, as above, a closed form expression for  $\mathbb{E}(X^k)$  is known, or, the moments of  $X$  are characterized in terms of the moment generating function (mgf)  $\mathbb{E}(e^{tX})$  of  $X$ , where  $t$  is a real-valued variable. In the latter case,  $\mathbb{E}(X^k)$  is obtained by taking the  $k$ th derivative of the mgf and evaluating it at  $t = 0$  [6, p. 62]. An inspection of the mgfs of the geometric, and the Poisson distributions [6, p. 621f] reveals that their  $k$ th moments are polynomials in their respective parameters as well. Together,  $\mathcal{R}$  has polynomially computable moments up to order  $k$  for all  $k \in \mathbb{N}_+$ . ◻



► **Proposition 4.5.** *Let  $Q = \bigvee_{i=1}^N Q_i$  be a Boolean UCQ, and let  $\mathcal{R}$  be a TIRS with polynomially computable moments up to order  $k$ , where  $k$  is the maximum self-join width among the  $Q_i$ . Then  $\text{EXPECTATION}_{\mathcal{R}}(\#_T Q)$  is computable in polynomial time.*

**Proof.** We plug (3) into the formula from Lemma 4.2. This yields at most  $\leq N \cdot |\text{adom}(T)|^{m \cdot m}$  terms (where  $m$  is the maximal number of atoms among the CQs  $Q_1, \dots, Q_N$ ). These terms only contain moments of fact multiplicities of order at most  $k$ . ◀

We emphasize that the number  $k$  from Proposition 4.5, that dictates which moments we need to be able to compute efficiently, comes from the fixed query  $Q$  and is therefore constant. More precisely, it is given through the number of self-joins in the query. In particular, if all CQs in  $Q$  are self-join free, it suffices to have efficient access to the expectations of the multiplicities.

## 4.2 Variance of the Answer Count

With the ideas from the previous section, we can also compute the variance of query answers in polynomial time. Naturally, to be able to calculate the variance efficiently, we need moments of up to the double order in comparison to the computation of the expected value.

► **Proposition 4.6.** *Let  $Q = \bigvee_{i=1}^N Q_i$  be a Boolean UCQ, and let  $\mathcal{R}$  be a TIRS with polynomially computable moments up to order  $2k$ , where  $k$  is the maximum self-join width among the  $Q_i$ . Then  $\text{VARIANCE}_{\mathcal{R}}(\#_T Q)$  is computable in polynomial time.*

As before, the main idea is to rewrite the variance in terms of the moments of fact multiplicities. This can be achieved by exploiting tuple-independence and linearity of expectation. The full proof is contained in the extended version of this paper [20].

Despite the fact that the variance of query answers may be of independent interest, it can be also used to obtain bounds for the probability that the true value of  $\#_T Q$  is close to its expectation, using the Chebyshev inequality [24, Theorem 5.11]. This can be used to derive bounds on  $\Pr(\#_T Q \leq k)$ , when the exact value is hard to compute.

► **Remark 4.7.** Proposition 4.6 extends naturally to higher-order moments: If  $\mathcal{R}$  is a TIRS with polynomially computable moments up to order  $\ell \cdot k$  and  $Q = \bigvee_{i=1}^N Q_i$  a Boolean UCQ where the maximum self-join width among the  $Q_i$  is  $k$ , then all centralized and all raw moments of order up to  $\ell$  of  $\#_T Q$  are computable in polynomial time. ◻

## 5 Answer Count Probabilities

In this section, we treat the alternative version of probabilistic query evaluation in bag PDBs using answer count probabilities rather than expected values. Formally, we discuss the following problem.

### Problem $\text{PQE}_{\mathcal{R}}(Q, k)$

PARAMETER: A Boolean (U)CQ  $Q$ , and  $k \in \mathbb{N}$ .

INPUT: A table  $T \in \mathbf{T}_{\mathcal{R}}$ .

OUTPUT: The probability  $\Pr(\#_T Q \leq k)$ .

This problem amounts to evaluating the cumulative distribution function of the random variable  $\#_T Q$  at  $k$ . The properties of this problem bear a close resemblance to the set version of probabilistic query evaluation, and we hence name this problem “PQE”.

► **Remark 5.1.** Instead of asking for  $\Pr(\#_T Q \leq k)$ , we could similarly define the problem of evaluating the probability that  $\#_T Q$  is *at least*, or *exactly* equal to  $k$ . This has no impact on complexity discussions, though, as these variants are polynomial time equivalent. ◻

Throughout this section, calculations involve the probabilities for the multiplicities of individual facts. However, we want to discuss the complexity of  $\text{PQE}_{\mathcal{R}}(Q, k)$  independently of the complexity, in  $k$ , of evaluating the multiplicity distributions. This motivates the following definition, together with taking  $k$  to be a parameter, instead of it being part of the input.

► **Definition 5.2.** A TIRS  $\mathcal{R}$  is called a *p-TIRS*, if for all  $k \in \mathbb{N}$  there exists a polynomial  $p_k$  such that for all  $\lambda \in \Lambda_{\mathcal{R}}$ , the function  $\langle \lambda \rangle \mapsto P_{\lambda}(k)$  can be evaluated in time  $\mathcal{O}(p_k(|\langle \lambda \rangle|))$ . ◻

Definition 5.2 captures reasonable assumptions for “efficient” TIRS’s with respect to the evaluation of probabilities: If the requirement from the definition is not given, then  $\text{PQE}_{\mathcal{R}}(\exists x: R(x), k)$  can not be solved in polynomial time, even on the class of tables that only contain a single annotated fact  $R(a)$ . This effect only arises due to the presence of unwieldy probability distributions in  $\mathcal{R}$ .

As it turns out, solving  $\text{PQE}_{\mathcal{R}}(Q, k)$  proves to be far more intricate compared to the problems of the previous section. For our investigation, we concentrate on self-join free conjunctive queries. While some simple results follow easily from the set semantics version of the problem, the complexity theoretic discussions quickly become quite involved and require the application of a set of interesting non-trivial techniques.

Our main result in this section is a dichotomy for Boolean CQs without self-joins. From now on, we employ nomenclature (like *hierarchical*) that was introduced in [11, 12]. If  $Q$  is a Boolean self-join free CQ, then for every variable  $x$ , we let  $\text{at}(x)$  denote the set of relation symbols  $R$  such that  $Q$  contains an  $R$ -atom that contains  $x$ . We call  $Q$  *hierarchical* if for all distinct  $x$  and  $y$ , whenever  $\text{at}(x) \cap \text{at}(y) \neq \emptyset$ , then  $\text{at}(x) \subseteq \text{at}(y)$  or  $\text{at}(y) \subseteq \text{at}(x)$ . This definition essentially provides the separation between easy and hard Boolean CQs without self-joins. In the bag semantics setting, however, there exists an edge case where the problem gets easy just due to the limited expressive power of the representation system. This edge case is governed only by the probabilities for multiplicity zero that appear in the representations. We denote this set by  $\text{zeroPr}(\mathcal{R})$ , i.e.,

$$\text{zeroPr}(\mathcal{R}) = \{p \in [0, 1] : P_{\lambda}(0) = p \text{ for some } \lambda \in \Lambda(\mathcal{R})\}.$$

If a p-TIRS satisfies  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$ , then it can only represent bag PDBs whose deduplication is deterministic. In this case, as we will show in the next subsection, the problem becomes easy even for arbitrary UCQs.<sup>2</sup>

► **Theorem 5.3.** *Let  $Q$  be a Boolean CQ without self-joins and let  $\mathcal{R}$  be a p-TIRS.*

1. *If  $Q$  is hierarchical or  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$ , then  $\text{PQE}_{\mathcal{R}}(Q, k)$  is solvable in polynomial time for all  $k \in \mathbb{N}$ .*
2. *Otherwise,  $\text{PQE}(Q, k)$  is  $\#P$ -hard for all  $k \in \mathbb{N}$ .*

► **Remark 5.4.** It is natural to ask what happens, if  $k$  is treated as part of the input. With a reduction similar to the proof of Proposition 5 in [28], it is easy to identify situations in which the corresponding problem is  $\#P$ -hard. For example, this is already the case for the

<sup>2</sup> Using the same definition in the set semantics version of the problem would come down to restricting the input tuple-independent PDB to only use 0 and 1 as marginal probabilities, so the problem would collapse to traditional (non-probabilistic) query evaluation. Under a bag semantics, there still exist interesting examples in this class, as the probability distribution over non-zero multiplicities is not restricted in any way.

simple query  $\exists x: R(x)$ , provided that  $k$  is encoded in binary and that the p-TIRS supports fair coin flips whose outcomes are either a positive integer, or zero. A full proof is shown in the extended version [20].  $\lrcorner$

The remainder of this section is dedicated to establishing Theorem 5.3.

## 5.1 Tractable Cases

Let us first discuss the case of p-TIRS's  $\mathcal{R}$  with  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$ . Here,  $\text{PQE}_{\mathcal{R}}(Q, 0)$  is trivial, because the problem essentially reduces to deterministic query evaluation. The following lemma generalizes this to all values of  $k$ .

► **Lemma 5.5.** *If  $\mathcal{R}$  is a p-TIRS with  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$ , then  $\text{PQE}_{\mathcal{R}}(Q, k)$  is solvable in polynomial time for all Boolean UCQs  $Q$ , and all  $k \in \mathbb{N}$ .*

**Proof.** Let  $\mathcal{R}$  be any p-TIRS with  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$  and let  $Q$  be any Boolean UCQ. If  $P_{\lambda}(0) = 1$  for all  $\lambda \in \Lambda_{\mathcal{R}}$ , then  $\mathcal{R}$  can only represent the PDB where the empty instance has probability 1. In this case,  $\#_T Q = 0$  almost surely, so  $\Pr(\#_T Q \leq k) = 1$  for all  $k \in \mathbb{N}$ .

In the general case, suppose  $Q = \bigvee_{i=1}^N Q_i$  such that  $Q_1, \dots, Q_N$  are CQs. Let  $A$  be the set of functions  $\alpha$  that map the variables of  $Q$  into the active domain of the input  $T$ . We call  $\alpha$  *good*, if there exists  $i \in \{1, \dots, N\}$  such that all the facts emerging from the atoms of  $Q_i$  by replacing every variable  $x$  with  $\alpha(x)$  have positive multiplicity in  $T$  (almost surely). If there are at least  $k + 1$  good  $\alpha$  in  $A$ , then  $\#_T Q > k$  with probability 1 and, hence, we return 0 in this case. Otherwise, when there are at most  $k$  good  $\alpha$ , we restrict  $T$  to the set of all facts that can be obtained from atoms of  $Q$  by replacing all variables  $x$  with  $\alpha(x)$  (and retaining the parameters  $\lambda$ ). The resulting table  $T'$  contains at most  $k$  times the number of atoms in  $Q$  many facts, which is independent of the number of facts in  $T$ . Hence, we can compute  $\Pr(\#_T Q \leq k) = \Pr(\#_{T'} Q \leq k)$  in time polynomial in  $T$  by using brute-force.  $\blacktriangleleft$

From now on, we focus on the structure of queries again. The polynomial time procedure for Boolean CQs without self-joins is reminiscent of the original algorithm for set semantics as described in [11]. Therefore, we need to introduce some more vocabulary from their work. A variable  $x$  is called *maximal*, if  $\text{at}(y) \subseteq \text{at}(x)$  for all  $y$  with  $\text{at}(x) \cap \text{at}(y) \neq \emptyset$ . With every CQ  $Q$  we associate an undirected graph  $G_Q$  whose vertices are the variables appearing in  $Q$ , and where two variables  $x$  and  $y$  are adjacent if they appear in a common atom. Let  $V_1, \dots, V_m$  be the vertex sets of the connected components of  $G_Q$ . We can then write the quantifier-free part  $Q^*$  of  $Q$  as  $Q^* = Q_0^* \wedge \bigwedge_{i=1}^m Q_i^*$  where  $Q_0^*$  is the conjunction of the constant atoms of  $Q$  and  $Q_1^*, \dots, Q_m^*$  are the conjunctions of atoms corresponding to the connected components  $V_1, \dots, V_m$ . We call  $Q_1^*, \dots, Q_m^*$  the *connected components* (short: *components*) of  $Q$ .

► **Remark 5.6.** If  $Q$  is hierarchical, then every component of  $Q$  contains a maximal variable.<sup>3</sup> Moreover, if  $x$  is maximal in a component  $Q_i^*$ , then  $x$  appears in all atoms of  $Q_i^*$ .  $\lrcorner$

► **Remark 5.7.** For every CQ  $Q$  with components  $Q_1^*, \dots, Q_m^*$ , and constant atoms  $Q_0^*$ , the answer on every instance  $D$  is given by the product of the answers of the queries  $Q_0, \dots, Q_m$ , where  $Q_i = \exists \mathbf{x}_i: Q_i^*$  (and  $Q_0 = Q_0^*$ ), and  $\mathbf{x}_i$  are exactly the variables appearing in the component  $Q_i^*$ . That is,  $\#_D Q = \#_D Q_0^* \cdot \prod_{i=1}^m \#_D (\exists \mathbf{x}_i: Q_i^*)$ . This is shown in the extended version of this paper [20]. If convenient, we therefore use  $Q_0 \wedge Q_1 \wedge \dots \wedge Q_m$  as an alternative representation of  $Q$ .  $\lrcorner$

<sup>3</sup> This is true, since the sets  $\text{at}(x)$  for the variables of any component have a pairwise non-empty intersection, meaning that they are pairwise comparable with respect to  $\subseteq$ .

## 20:12 Probabilistic Query Evaluation with Bag Semantics

The main result of this subsection is the following.

► **Theorem 5.8.** *Let  $\mathcal{R}$  be a  $p$ -TIRS, and let  $Q$  be a hierarchical Boolean CQ without self-joins. Then  $\text{PQE}_{\mathcal{R}}(Q, k)$  is solvable in polynomial time for each  $k \in \mathbb{N}$ .*

**Proof Sketch.** The theorem is established by giving a polynomial time algorithm that computes, and adds up the probabilities  $\Pr(\#_T Q = k')$  for all  $k' \leq k$ . The important observation is that (as under set semantics) the components  $Q_i$  of the query (and the conjunction  $Q_0$  of the constant atom) yield independent events, which follows since  $Q$  is self-join free. In order to compute the probability of  $\#_T Q = k'$ , we can thus sum over all decompositions of  $k'$  into a product  $k' = k_0 \cdot k_1 \cdots k_m$ , and reduce the problem to the computations of  $\Pr(\#_T Q_i = k_i)$ . Although the cases  $k = 0$ , and the conjunction  $Q_0$  have to be treated slightly different for technical reasons, we can proceed recursively: Every component contains a maximal variable, and setting this variable to any constant, the component potentially breaks up into a smaller hierarchical, self-join free CQ. Investigating the expressions shows that the total number of operations on the probabilities of fact probabilities is polynomial in the size of  $T$ . ◀

► **Remark 5.9.** The full proof of Theorem 5.8 can be found in the extended version [20]. As pointed out, the proof borrows main ideas from the algorithm for the probabilistic evaluation of hierarchical Boolean self-join free CQs on tuple-independent PDBs with set semantics, as presented in [12, p. 30:15] (originating in [9, 10]). The novel component is the treatment of multiplicities using bag semantics. In comparison to the algorithm of Dalvi and Suciu, existential quantifiers behave quite differently here, and we additionally need to argue about the possible ways to distribute a given multiplicity over subformulae or facts. ▽

### 5.2 Intractable Cases

We now show that in the remaining case (non-hierarchical queries and  $p$ -TIRS's with  $\text{zeroPr}(\mathcal{R}) \cap (0, 1) \neq \emptyset$ ), the problems  $\text{PQE}_{\mathcal{R}}(Q, k)$  are all hard to solve.

Let  $Q$  be a fixed query and let  $\text{PQE}^{\text{set}}(Q)$  denote the traditional set version of the probabilistic query evaluation problem. That is,  $\text{PQE}^{\text{set}}(Q)$  is the problem to compute the probability that  $Q$  evaluates to true under *set* semantics, on input a tuple-independent *set* PDB. We recall that the bag version  $\text{PQE}_{\mathcal{R}}(Q, k)$  of the problem (introduced at the beginning of the section) takes the additional parameter  $k$ , and depends on the representation system  $\mathcal{R}$ . Let us first discuss  $\text{PQE}_{\mathcal{R}}(Q, k)$  for  $k = 0$ . In this case, subject to very mild requirements on  $\mathcal{R}$ , we can lift  $\#P$ -hardness from the set version [12], even for the full class of UCQs.

► **Proposition 5.10.** *Let  $S \subseteq [0, 1]$  be finite and let  $\mathcal{R}$  be a  $p$ -TIRS such that  $1 - p \in \text{zeroPr}(\mathcal{R})$  for all  $p \in S$ . Let  $Q$  be a Boolean UCQ. If  $\text{PQE}^{\text{set}}(Q)$  is  $\#P$ -hard on tuple-independent (set) PDBs with marginal probabilities in  $S$ , then  $\text{PQE}_{\mathcal{R}}(Q, 0)$  is  $\#P$ -hard.*

**Proof.** Let  $\mathcal{D}$  be an input to  $\text{PQE}^{\text{set}}(Q)$  with fact set  $F$  where all marginal probabilities are in  $S$ , given by the list of all facts  $f$  with their marginal probability  $p_f$ . For all  $p \in S$ , pick  $\lambda_p \in \Lambda$  such that  $P_{\lambda}(0) = 1 - p$ . Let  $T = \bigcup_{f \in F} \{(f, \langle \lambda_{p_f} \rangle)\}$ , and let  $\delta$  be the function that maps every instance  $D$  of  $T$  to its deduplication  $D'$  (which is an instance of  $\mathcal{D}$ ). Then, by the choice of the parameters, we have  $\Pr_{D \sim [T]}(\delta(D) = D') = \Pr_{\mathcal{D}}(\{D'\})$  for all  $D'$ . Moreover,  $\#_D Q > 0$  if and only if  $\delta(D) \models Q$ . Thus,

$$\Pr_{D \sim [T]}(\#_D Q > 0) = \Pr_{D \sim [T]}(\delta(D) \models Q) = \Pr_{D' \sim \mathcal{D}}(D' \models Q).$$

Therefore,  $\text{PQE}^{\text{set}}(Q)$  over tuple-independent PDBs with marginal probabilities from  $S$  can be solved by solving  $\text{PQE}_{\mathcal{R}}(Q, 0)$ . ◀

Remarkably, [23, Theorem 2.2] shows that Boolean UCQs for which  $\text{PQE}^{\text{set}}(Q)$  is  $\sharp\text{P}$ -hard are already hard when the marginal probabilities are restricted to  $S = \{c, 1\}$ , for any rational  $c \in (0, 1)$ . Hence,  $\text{PQE}_{\mathcal{R}}(Q, 0)$  is also  $\sharp\text{P}$ -hard on these queries, as soon as  $\{0, 1 - c\} \subseteq \text{zeroPr}(\mathcal{R})$ .

Our goal is now to show that if  $\mathcal{R}$  is a p-TIRS, then for any Boolean CQ  $Q$  without self-joins,  $\sharp\text{P}$ -hardness of  $\text{PQE}_{\mathcal{R}}(Q, 0)$  transfers to  $\text{PQE}_{\mathcal{R}}(Q, k)$  for all  $k > 0$ .

► **Theorem 5.11.** *Let  $\mathcal{R}$  be a p-TIRS and let  $Q$  be a Boolean self-join free CQ. Then, if  $\text{PQE}_{\mathcal{R}}(Q, 0)$  is  $\sharp\text{P}$ -hard,  $\text{PQE}_{\mathcal{R}}(Q, k)$  is  $\sharp\text{P}$ -hard for each  $k \in \mathbb{N}$ .*

Proving Theorem 5.11 is quite involved, and is split over various lemmas in the remainder of this subsection. Let  $\mathcal{R}$  be any fixed p-TIRS and let  $Q$  be a Boolean CQ without self-joins. We demonstrate the theorem by presenting an algorithm that solves  $\text{PQE}_{\mathcal{R}}(Q, 0)$  in polynomial time, when given an oracle for  $\text{PQE}_{\mathcal{R}}(Q, k)$  for *any* positive  $k$ .

Clearly, we cannot simply infer  $\Pr(\sharp_T Q = 0)$  from  $\Pr(\sharp_T Q \leq k)$ . Naively, we would want to shift the answer count of  $Q$  by  $k$ , so that the problem could be answered immediately. However, this is not possible in general. Our way out is to use the oracle several times, on manipulated inputs. Since the algorithms we describe are confined to the p-TIRS  $\mathcal{R}$ , we are severely restricted in the flexibility of manipulating the probabilities of fact multiplicities: Unless further assumptions are made, we can only work with the annotations that are already present in the input  $T$  to the problem. We may, however, also drop entries from  $T$  or introduce copies of facts using new domain elements.

For a given table  $T$  and a fixed single-component query  $Q$ , we exploit this idea in Algorithm 1 in order to construct a new table  $T^{(m)}$ , called the *inflation* of  $T$  of order  $m$ . It has the property that  $\sharp_{T^{(m)}} Q$  is the sum of answer counts of  $Q$  on  $m$  independent copies of  $T$ . A small example for the result of running Algorithm 1 for  $m = 2$  is shown in Figure 3. We will later use oracle answers on several inflations in order to interpolate  $\Pr(\sharp_T Q_i = 0)$  per component  $Q_i$  individually, and then combine the results together.

■ **Algorithm 1**  $\text{inflate}_Q(T, m)$ .

---

**Parameter:** Boolean self-join free CQ  $Q$  with a single component and no constant atoms

**Input:**  $T \in \mathbf{T}_{\mathcal{R}}$ ,  $m \in \mathbb{N}$

**Output:** Inflation of order  $m$  of  $T$ :  $T^{(m)} = \bigcup_{i=1}^m T_{m,i} \in \mathbf{T}_{\mathcal{R}}$  such that

(O1) for all  $i \neq j$  we have  $T_{m,i} \cap T_{m,j} = \emptyset$ ,

(O2) for all  $i = 1, \dots, m$  we have  $\sharp_{T_{m,i}} Q \sim \sharp_T Q$  i.i.d., and

(O3)  $\sharp_{T^{(m)}} Q = \sum_{i=1}^m \sharp_{T_{m,i}} Q$ .

---

```

1: Initialize  $T_{m,1}, \dots, T_{m,m}$  to be empty.
2: For each domain element  $a$ , introduce new pairwise distinct elements  $a^{(1)}, \dots, a^{(m)}$ .
3: for all relation symbols  $R$  appearing in  $Q$  do
4:   Let  $R(t_1, \dots, t_r)$  be the unique atom in  $Q$  with relation symbol  $R$ .
5:   for all pairs of the form  $(R(a_1, \dots, a_r), \langle \lambda \rangle) \in T$  do
6:     for all  $i = 1, \dots, m$  do
7:       Add  $(R(a_{i,1}, \dots, a_{i,k}), \langle \lambda \rangle)$  to  $T_{m,i}$  where  $a_{i,j} = \begin{cases} a_j^{(i)}, & \text{if } t_j \text{ is a variable;} \\ a_j, & \text{if } t_j \text{ is a constant.} \end{cases}$ 
8:     end for
9:   end for
10: end for
11: return  $T^{(m)} := \bigcup_{i=1}^m T_{m,i}$ 

```

---

Table $T$		Table $T^{(2)}$	
Relation $R$	Parameter	Relation $R$	Parameter
$R(a, a, a)$	(Binomial, 10, 1/3)	$R(a^{(1)}, a^{(1)}, a)$	(Binomial, 10, 1/3)
$R(a, b, c)$	(Geometric, 1/2)	$R(a^{(1)}, b^{(1)}, c)$	(Geometric, 1/2)
		$R(a^{(2)}, a^{(2)}, a)$	(Binomial, 10, 1/3)
		$R(a^{(2)}, b^{(2)}, c)$	(Geometric, 1/2)

■ **Figure 3** Example of a table  $T$  and its inflation  $T^{(2)}$  for the query  $\exists x, y: R(x, y, a)$ .

► **Lemma 5.12.** *For every fixed Boolean self-join free CQ  $Q$  with a single component and no constant atoms, Algorithm 1 runs in time  $\mathcal{O}(|T| \cdot m)$ , and satisfies the output conditions (O1), (O2) and (O3).*

The proof of Lemma 5.12 is contained in the extended version [20]. The assumption that the input to Algorithm 1 is self-join free with just a single connected component and no constant atoms is essential to establish (O1) and (O3), because it eliminates any potential co-dependencies among the individual tables  $T_{m,1}, \dots, T_{m,m}$  we create. The following example shows that this assumption is inevitable, as the conditions of Lemma 5.12 can not be established in general.

► **Example 5.13.** Let  $\mathcal{R}$  be a TIRS with  $\Lambda_{\mathcal{R}} = \{\lambda\}$  such that  $P_{\lambda}(2) = P_{\lambda}(3) = \frac{1}{2}$  (and  $P_{\lambda}(k) = 0$  for all  $k \notin \{2, 3\}$ ). Consider  $Q = \exists x \exists y: R(x) \wedge S(y)$ , and  $T = ((R(1), \langle \lambda \rangle), (S(1), \langle \lambda \rangle)) \in \mathbf{T}_{\mathcal{R}}$ . Note that  $Q$  has two components and, hence, does not satisfy the assumptions of Lemma 5.12. Then  $\#_T Q$  takes the values 4, 6 and 9, with probabilities  $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$ . Thus, if  $X, Y \sim \#_T Q$  i.i.d., then  $X + Y$  is 13 with probability  $\frac{1}{16} + \frac{1}{16} = \frac{1}{8}$ . However, for every  $T' \in \mathbf{T}_{\mathcal{R}}$ , the random variable  $\#_{T'} Q$  almost surely takes composite numbers, as it is equal to the sum of all multiplicities of  $R$ -facts, times the sum of all multiplicities of  $S$ -facts, both of these numbers being either 0 or at least 2. Thus, there exists no  $T' \in \mathbf{T}_{\mathcal{R}}$  such that  $\#_{T'} Q = X + Y$ . ⌋

For this reason, our main algorithm will call Algorithm 1 independently, for each connected component  $Q_i$  of  $Q$ . Then, Algorithm 1 does not inflate the whole table  $T$ , but only the part  $T_i$  corresponding to  $Q_i$ . If  $Q = Q' \wedge Q_i$  and we denote  $\#_{T_i} Q_i$  by  $X$  and  $\#_{T \setminus T_i} Q'$  by  $Y$ , then replacing  $T_i$  in  $T$  with its inflation of order  $n$  yields a new table with answer count  $(\#_{T \setminus T_i} Q') \cdot (\#_{T_i^{(n)}} Q_i) = Y \cdot \sum_{i=1}^n X_i$ , where  $X_1, \dots, X_n \sim X$  i.i.d. Before further describing the reduction, we first explore some algebraic properties of the above situation in general.

► **Lemma 5.14.** *Let  $X$  and  $Y$  be independent random variables with values in  $\mathbb{N}$  and let  $k \in \mathbb{N}$ . Suppose  $X_1, X_2, \dots$  are i.i.d. random variables with  $X_1 \sim X$ . Let  $p_0 := \Pr(X = 0)$  and  $q_0 := \Pr(Y = 0)$ . Then, there exist  $z_1, \dots, z_k \geq 0$  such that for all  $n \in \mathbb{N}$  we have*

$$\Pr\left(Y \cdot \sum_{i=1}^n X_i \leq k\right) = q_0 + (1 - q_0) \cdot p_0^n + \sum_{j=1}^k \binom{n}{j} \cdot p_0^{n-j} \cdot z_j.$$

This is demonstrated in the extended version [20]. We now describe how  $p_0 = \Pr(X = 0)$  can be recovered from the values of  $\Pr\left(Y \cdot \sum_{i=1}^n X_i \leq k\right)$  and  $q_0 = \Pr(Y = 0)$  whenever  $q_0 < 1$  and  $p_0 > 0$ . With the values  $z_1, \dots, z_n$  from Lemma 5.14, and  $z_0 := 1 - q_0$ , we define a function

$$g(n) := \Pr\left(Y \cdot \sum_{i=1}^n X_i \leq k\right) - q_0 = \sum_{j=0}^k \binom{n}{j} \cdot p_0^{n-j} \cdot z_j. \quad (4)$$

Now, for  $m \in \mathbb{N}$  and  $x = 0, 1, \dots, m$ , we define

$$h_m(x) := g(m+x) \cdot g(m-x) = \sum_{j_1, j_2=0}^k \binom{m+x}{j_1} \cdot \binom{m-x}{j_2} \cdot p_0^{2m-j_1-j_2} \cdot z_{j_1} \cdot z_{j_2}. \quad (5)$$

Then, for every fixed  $m$ ,  $h_m$  is a polynomial in  $x$  with domain  $\{0, \dots, m\}$ . As it will turn out, the leading coefficient  $\text{lc}(h_m)$  of  $h_m$  can be used to recover the value of  $p_0$  as follows: Let  $j_{\max}$  be the maximum  $j$  such that  $z_j \neq 0$ . Since for fixed  $m$ , both  $\binom{m+x}{j}$  and  $\binom{m-x}{j}$  are polynomials of degree  $j$  in  $x$ , the degree of  $h_m$  is  $2j_{\max}$  and its leading coefficient is

$$\text{lc}(h_m) = (-1)^{j_{\max}} \cdot (j_{\max}!)^{-2} \cdot p_0^{2m-2j_{\max}} \cdot z_{j_{\max}}^2,$$

which yields

$$p_0 = \sqrt{\frac{\text{lc}(h_{m+1})}{\text{lc}(h_m)}}. \quad (6)$$

Thus, it suffices to determine  $\text{lc}(h_m)$  and  $\text{lc}(h_{m+1})$ . However, we neither know  $j_{\max}$ , nor  $z_{j_{\max}}$ , and we only have access to the values of  $h_m$  and  $h_{m+1}$ . To find the leading coefficients anyway, we employ the method of finite differences, a standard tool from polynomial interpolation [22, chapter 4]. For this, we use the difference operator  $\Delta$  that is defined as  $\Delta f(x) := f(x+1) - f(x)$  for all functions  $f$ . When  $f$  is a (non-zero) polynomial of degree  $n$ , the difference operator reduces its degree by one and its leading coefficient is multiplied by  $n$ . Therefore, after taking differences  $n$  times, starting from subsequent values of a polynomial  $f$ , we are left with the constant function  $\Delta^n f = n! \cdot \text{lc}(f) \neq 0$ . In particular, taking differences more than  $n$  times yields the zero function. Hence, we can determine  $\text{lc}(f)$  by finding the largest  $\ell$  for which  $\Delta^\ell f(0) \neq 0$ .

■ **Algorithm 2**  $\text{solveComponent}_Q(T, i)$ .

**Parameter:** Boolean self-join free CQ  $Q = Q_0 \wedge \bigwedge_{i=1}^r Q_i$  with connected components  $Q_1 \dots, Q_r$ .

**Oracle Access:** Oracle for  $\text{PQE}_{\mathcal{R}}(Q, k)$  that, on input  $\tilde{T}$ , returns  $\Pr(\#\_{\tilde{T}} Q \leq k)$

**Input:**  $T \in \mathbf{T}_{\mathcal{R}}$ ,  $i \in \{1, \dots, r\}$

**Output:**  $\Pr(\#\_{T_i} Q_i = 0)$

---

```

1: if  $P_\lambda(0) = 1$  for all  $\lambda \in \Lambda$  then return 0 end if
2: Fix  $\lambda$  with  $P_\lambda(0) < 1$  and suppose  $Q = Q' \wedge Q_i$  (cf. Remark 5.7).
3: if  $Q'$  is empty then
4:   Set  $q_0 := 0$  and  $g(0) := 1$ .
5: else
6:   Let  $T' \in \mathbf{T}_{\mathcal{R}}$  be the canonical database for  $Q'$ , with  $\lambda_f := \lambda$  for all facts  $f$  in  $T'$ .
7:   Calculate  $q_0 := \Pr(\#\_{T'} Q' = 0)$  and set  $g(0) := 1 - q_0$ .
8: end if
9: for  $n = 1, 2, \dots, 4k + 1$  do
10:   Set  $T_i^{(n)} := \text{inflate}_{Q_i}(T_i, n)$ .
11:   Set  $g(n) := \Pr(\#\_{T' \cup T_i^{(n)}} Q \leq k) - q_0$ , using the oracle.
12: end for
13: if  $g(k+1) = 0$  then return 0 end if
14: for  $x = 0, 1, \dots, 2k$  and  $m = 2k, 2k + 1$  do  $h_m(x) := g(m+x) \cdot g(m-x)$  end for
15: Initialize  $\ell := 2k$ .
16: while  $\Delta^\ell h_{2k}(0) = 0$  do  $\ell := \ell - 1$  end while
17: return  $\sqrt{\Delta^\ell h_{2k+1}(0) / \Delta^\ell h_{2k}(0)}$ 

```

---

## 20:16 Probabilistic Query Evaluation with Bag Semantics

The full procedure that uses the above steps to calculate  $p_0$  yields Algorithm 2. Recall that it focuses on a single connected component. To ensure easy access to the value of  $q_0$ , we utilize a table that encodes the canonical database of the remainder of the query.<sup>4</sup> Note that  $k$  is always treated as a fixed constant, and our goal is to reduce  $\text{PQE}_{\mathcal{R}}(Q, 0)$  to  $\text{PQE}_{\mathcal{R}}(Q, k)$ .

► **Lemma 5.15.** *Algorithm 2 runs in polynomial time and yields the correct result.*

**Proof.** With the notation introduced in the algorithm, we let  $Y = \#_{T'}Q'$  (or  $Y = 1$  if  $Q'$  is empty) and  $X = \#_T Q_i = \#_{T_i} Q_i$ . Then,  $q_0 = \Pr(Y = 0)$  as in Lemma 5.14 and the aim of the algorithm is to return  $p_0$ .

First, line 1 covers the edge case that  $\mathcal{R}$  can only represent the empty database instance. In all other cases, we fix  $\lambda$  with  $P_\lambda(0) > 0$ . As  $q_0 = 1 - (1 - P_\lambda(0))^t$  where  $t$  is the number of atoms of  $Q'$ , we have  $q_0 < 1$ . From Lemma 5.12, we see that  $\#_{T' \cup T_i^{(n)}} Q = Y \cdot \sum_{i=1}^n X_i$ , so we are in the situation of Lemma 5.14. Hence,  $g$  and  $h_m$  are as in (4) and (5). Now, as  $g(k+1) = p_0 \cdot \sum_{j=0}^k \binom{n}{j} \cdot p_0^{k-j} \cdot z_j$  with  $z_0 = 1 - q_0 > 0$ , we find that  $p_0$  is zero if and only if  $g(k+1)$  is zero. This is checked in line 13. Finally, the paragraphs following Lemma 5.14 apply, and we determine the degree of  $h_{2k}$  using the method of finite differences by setting  $\ell$  to the maximum possible degree and decreasing it step-by-step as long as  $\Delta^\ell h_{2k}(0) = 0$  in lines 15 and 16. Then, we have  $\ell = 2j_{\max}$  and return

$$\sqrt{\frac{\Delta^\ell h_{2k+1}(0)}{\Delta^\ell h_{2k}(0)}} = \sqrt{\frac{\ell! \text{lc}(h_{m+1})}{\ell! \text{lc}(h_m)}} \stackrel{(6)}{=} p_0.$$

Concerning the runtime, since  $\mathcal{R}$  is a p-TIRS, all answers of the oracle calls are of polynomial size in the input. Since  $k$  is fixed, the algorithm performs a constant number of computation steps and each term in the calculations is either independent of the input or of polynomial size, yielding a polynomial runtime. ◀

**Proof of Theorem 5.11.** Let  $k > 0$  and suppose that we have an oracle for  $\text{PQE}_{\mathcal{R}}(Q, k)$ . Let  $Q = Q_0 \wedge \bigwedge_{i=1}^m Q_i$  be the partition of  $Q$  into components, with  $Q_0$  being the conjunction of the constant atoms. Then the  $\#_T Q_i$  are independent and  $\#_T Q = \#_T Q_0 \cdot \prod_{i=1}^m \#_T Q_i$ . Therefore,

$$\Pr(\#_T Q = 0) = 1 - \Pr(\#_T Q_0 \neq 0) \cdot \prod_{i=1}^m (1 - \Pr(\#_T Q_i = 0)).$$

As  $\Pr(\#_T Q_0 \neq 0)$  is easy to compute and Algorithm 2 computes  $\Pr(\#_T Q_i = 0)$  for  $i = 1, \dots, k$  with oracle calls for  $\text{PQE}_{\mathcal{R}}(Q, k)$ , this yields a polynomial time Turing-reduction from  $\text{PQE}_{\mathcal{R}}(Q, 0)$  to  $\text{PQE}_{\mathcal{R}}(Q, k)$ . ◀

With the results from the previous subsections, this completes the proof of Theorem 5.3.

**Proof of Theorem 5.3.** For p-TIRS's with  $\text{zeroPr}(\mathcal{R}) \subseteq \{0, 1\}$ , the statement is given by Lemma 5.5. By Theorem 5.8,  $\text{PQE}_{\mathcal{R}}(Q, k)$  is solvable in polynomial time for hierarchical Boolean CQs without self-joins. For the case of  $Q$  being non-hierarchical (and  $\text{zeroPr}(\mathcal{R}) \cap (0, 1) \neq \emptyset$ ), let  $p \in (0, 1)$  such that  $p \in \text{zeroPr}(\mathcal{R})$ . By [2, Theorem 3.4], the set version  $\text{PQE}^{\text{set}}(Q)$  is already hard on the class of tuple-independent set PDBs where all probabilities are equal to  $1 - p$ . It follows from Proposition 5.10 that  $\text{PQE}_{\mathcal{R}}(Q, 0)$  is  $\#P$ -hard. By Theorem 5.11, so is  $\text{PQE}_{\mathcal{R}}(Q, k)$  for all  $k \in \mathbb{N}_+$ . ◀

<sup>4</sup> The canonical database belonging to a self-join free CQ is the instance containing the atoms appearing in the query, with all variables being treated as constants.



## 6 Conclusion

The results of our paper extend the understanding of probabilistic query evaluation into a new direction by discussing bag semantics. We investigated two principal computational problems: computing expectations, and computing the probability of answer counts. Interestingly, even though these problems are equivalent for set semantics, they behave quite differently under bag semantics. Our findings show that generally, computing expectations is the easier problem. For computing answer count probabilities, in the case of self-join free CQs, we obtain a polynomial time vs.  $\sharp\text{P}$ -hard dichotomy, depending on whether the query is hierarchical. This transfers the corresponding results of [9, 10] from set to bag semantics.

While our results for the expectation problem concern UCQs, the complexity of computing answer count probabilities remains open beyond self-join free CQs. It is also unclear, how the problem behaves on bag versions of other well-representable classes of set PDBs. A more detailed analysis of the complexity of  $\text{PQE}_{\mathcal{R}}(Q, k)$  in terms of  $k$  remains open as well.

To formally argue about the complexity of some natural distributions such as the Poisson distribution, irrational probabilities or parameters have to be supported. This yields non-trivial complexity theoretic questions that we leave for future work.

---

### References

- 1 Antoine Amarilli, Pierre Bourhis, and Pierre Senellart. Tractable Lineages on Treelike Instances: Limits and Extensions. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2016)*, pages 355–370. Association for Computing Machinery, 2016. doi:10.1145/2902251.2902301.
- 2 Antoine Amarilli and Benny Kimelfeld. Uniform Reliability of Self-Join-Free Conjunctive Queries, 2021. Extended version of [3]. arXiv:1908.07093v6.
- 3 Antoine Amarilli and Benny Kimelfeld. Uniform Reliability of Self-Join-Free Conjunctive Queries. In *24th International Conference on Database Theory (ICDT 2021)*, volume 186 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.ICDT.2021.17.
- 4 Antoine Amarilli and İsmail İlkan Ceylan. A Dichotomy for Homomorphism-Closed Queries on Probabilistic Graphs. In *23rd International Conference on Database Theory (ICDT 2020)*, volume 155 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:20, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICDT.2020.5.
- 5 Mark Braverman and Stephen Cook. Computing over the Reals: Foundations for Scientific Computing. *Notices of the AMS*, 53(3):318–329, 2006.
- 6 George Casella and Roger L. Berger. *Statistical Inference*. Thomson Learning, Pacific Grove, CA, USA, 2nd edition, 2002.
- 7 Surajit Chaudhuri and Moshe Y. Vardi. Optimization of Real Conjunctive Queries. In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 1993)*, pages 59–70. ACM Press, 1993. doi:10.1145/153850.153856.
- 8 Sara Cohen. Equivalence of Queries That Are Sensitive to Multiplicities. *The VLDB Journal*, 18(3):765–785, 2009. doi:10.1007/s00778-008-0122-1.
- 9 Nilesh Dalvi and Dan Suciu. Efficient Query Evaluation on Probabilistic Databases. In *Proceedings of the Thirtieth International Conference on Very Large Data Bases (VLDB 2004)*, volume 30, pages 864–875. VLDB Endowment, 2004.
- 10 Nilesh Dalvi and Dan Suciu. Efficient Query Evaluation on Probabilistic Databases. *The VLDB Journal*, 16(4):523–544, 2007. doi:10.1007/s00778-006-0004-3.

- 11 Nilesch Dalvi and Dan Suciu. The Dichotomy of Conjunctive Queries on Probabilistic Structures. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2007)*, pages 293–302, New York, NY, USA, 2007. Association for Computing Machinery. doi:10.1145/1265530.1265571.
- 12 Nilesch Dalvi and Dan Suciu. The Dichotomy of Probabilistic Inference for Unions of Conjunctive Queries. *Journal of the ACM*, 59(6):1–87, 2012. doi:10.1145/2395116.2395119.
- 13 Su Feng, Boris Glavic, Aaron Huber, Oliver Kennedy, and Atri Rudra. Computing Expected Multiplicities for Bag-TIDBs with Bounded Multiplicities, 2022. arXiv:2204.02758v3.
- 14 Robert Fink and Dan Olteanu. Dichotomies for Queries with Negation in Probabilistic Databases. *ACM Transactions on Database Systems*, 41(1):4:1–4:47, 2016. doi:10.1145/2877203.
- 15 Erich Grädel, Yuri Gurevich, and Colin Hirsch. The Complexity of Query Reliability. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 1998)*, pages 227–234. ACM Press, 1998. doi:10.1145/275487.295124.
- 16 Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance Semirings. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2007)*, pages 31–40. Association for Computing Machinery, 2007. doi:10.1145/1265530.1265535.
- 17 Todd J. Green and Val Tannen. Models for Incomplete and Probabilistic Information. In *Current Trends in Database Technology – EDBT 2006*, pages 278–296, Berlin, Germany, 2006. Springer-Verlag Berlin Heidelberg. doi:10.1007/11896548\_24.
- 18 Martin Grohe and Peter Lindner. Independence in Infinite Probabilistic Databases. *J. ACM*, 69(5):37:1–37:42, 2022. doi:10.1145/3549525.
- 19 Martin Grohe and Peter Lindner. Infinite Probabilistic Databases. *Logical Methods in Computer Science*, Volume 18, Issue 1, 2022. doi:10.46298/lmcs-18(1:34)2022.
- 20 Martin Grohe, Peter Lindner, and Christoph Standke. Probabilistic Query Evaluation with Bag Semantics, 2022. arXiv:2201.11524.
- 21 Erich Grädel and Val Tannen. Semiring Provenance for First-Order Model Checking, 2017. arXiv:1712.01980v1.
- 22 Francis Begnaud Hildebrand. *Introduction to Numerical Analysis*. Courier Corporation, 1987.
- 23 Batya Kenig and Dan Suciu. A Dichotomy for the Generalized Model Counting Problem for Unions of Conjunctive Queries. In *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS 2021)*, pages 312–324, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3452021.3458313.
- 24 Achim Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer-Verlag London, London, UK, 2nd edition, 2014. Translation from the German language edition. doi:10.1007/978-1-4471-5361-0.
- 25 Ker-I Ko. *Complexity Theory of Real Functions*. Progress in Theoretical Computer Science. Birkhäuser Boston, 1991. doi:10.1007/978-1-4684-6802-1.
- 26 Dan Olteanu and Jiewen Huang. Using OBDDs for Efficient Query Evaluation on Probabilistic Databases. In *SUM 2008: Scalable Uncertainty Management*, Lecture Notes in Computer Science, pages 326–340, Berlin, Heidelberg, 2008. Springer. doi:10.1007/978-3-540-87993-0\_26.
- 27 Dan Olteanu and Jiewen Huang. Secondary-Storage Confidence Computation for Conjunctive Queries with Inequalities. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, pages 389–402. ACM, 2009. doi:10.1145/1559845.1559887.
- 28 Christopher Ré and Dan Suciu. The Trichotomy Of HAVING Queries On A Probabilistic Database. *The VLDB Journal*, 18(5):1091–1116, July 2009. doi:10.1007/s00778-009-0151-4.
- 29 Dan Suciu. Probabilistic Databases for All. In *Proceedings of the 39th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 19–31. ACM, 2020. doi:10.1145/3375395.3389129.

- 30 Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. *Probabilistic Databases*, volume 3.2 of *Synthesis Lectures on Data Management*. Morgan & Claypool, San Rafael, CA, USA, 2011. Lecture #16. doi:10.2200/S00362ED1V01Y201105DTM016.
- 31 Guy Van den Broeck and Dan Suciu. Query Processing on Probabilistic Data: A Survey. *Foundations and Trends® in Databases*, 7(3–4):197–341, 2017. doi:10.1561/1900000052.
- 32 Moshe Y. Vardi. The Complexity of Relational Query Languages. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing (STOC 1982)*, pages 137–146, New York, NY, USA, 1982. ACM Press. doi:10.1145/800070.802186.