

The Consistency of Probabilistic Databases with Independent Cells

Amir Gilad ✉

Duke University, Durham, NC, USA

Aviram Imber ✉

Technion – Israel Institute of Technology

Benny Kimelfeld ✉

Technion – Israel Institute of Technology

Abstract

A probabilistic database with attribute-level uncertainty consists of relations where cells of some attributes may hold probability distributions rather than deterministic content. Such databases arise, implicitly or explicitly, in the context of noisy operations such as missing data imputation, where we automatically fill in missing values, column prediction, where we predict unknown attributes, and database cleaning (and repairing), where we replace the original values due to detected errors or violation of integrity constraints. We study the computational complexity of problems that regard the selection of cell values in the presence of integrity constraints. More precisely, we focus on functional dependencies and study three problems: (1) deciding whether the constraints can be satisfied by any choice of values, (2) finding a most probable such choice, and (3) calculating the probability of satisfying the constraints. The data complexity of these problems is determined by the combination of the set of functional dependencies and the collection of uncertain attributes. We give full classifications into tractable and intractable complexities for several classes of constraints, including a single dependency, matching constraints, and unary functional dependencies.

2012 ACM Subject Classification Information systems → Data management systems

Keywords and phrases Probabilistic databases, attribute-level uncertainty, functional dependencies, most probable database

Funding The work of Amir Gilad was supported by NSF awards IIS-1703431, IIS-1552538, and IIS-2008107. The work of Aviram Imber and Benny Kimelfeld was supported by the Israel Science Foundation (ISF), Grant 768/19, and the German Research Foundation (DFG) Project 412400621 (DIP program).

1 Introduction

Various database tasks amount to reasoning about relations where attribute values are uncertain. To name a few, systems for *data cleaning* may detect errors and suggest alternative fixes with different confidence scores [14, 25, 26], approaches to *data repair* may suggest alternative values due to the violation of integrity constraints (e.g., key constraints and more general functional dependencies) [2, 32], and algorithms for *missing-data imputation* may suggest a probability distribution over possible completions of missing values [3, 22]. Such uncertainty is captured as a probabilistic database in the so called *attribute-level uncertainty* [27] (as opposed to the commonly studied *tuple-level uncertainty* [8]).

We refer to a relation of a probabilistic database in the attribute-level uncertainty as a Cell-Independent Relation (CIR). A CIR is a probabilistic database with a single relation, where the content of a cell is a distribution over possible values, and different cells are probabilistically independent. The CIR is the correspondent of a relation in the *Tuple-Independent Database* (TID) under the *tuple-level uncertainty*, where the existence of each tuple is uncertain (while its content is certain), and different tuples are probabilistically independent [27]. In contrast,

<i>tid</i>	room	?specialist	time
1	41	Bart(0.5) Lisa(0.5)	5 PM
2	163	Bart(0.7) Lisa(0.3)	5 PM
3	41	Bart(0.2) Maggie(0.8)	5 PM

(a) CIR U_1 with uncertain specialist.

$$F_1 := \{?specialist \text{ time} \rightarrow \text{room}\}$$

$$F_2 := \{?specialist \text{ time} \rightarrow \text{room}, \\ \text{room time} \rightarrow ?specialist\}$$

(b) Sets F_1 and F_2 of functional dependencies

<i>tid</i>	room	?specialist	time
1	41	Lisa	5 PM
2	163	Bart	5 PM
3	41	Maggie	5 PM

(c) Samples r (left) and r' (right) of U_1 .

<i>tid</i>	room	?specialist	time
1	41	Bart	5 PM
2	163	Lisa	5 PM
3	41	Bart	5 PM

■ **Figure 1** Running example: CIR, FDs, and samples.

the tuples of a CIR always exist, but their content is uncertain. For illustration, Figure 1a depicts a CIR with uncertain information about specialists attending rooms (e.g., since their attendance is determined by noisy sensors). Some attributes (here *room* and *business*) are certain and have deterministic values. The uncertain attributes (e.g., *?specialist*) are marked by a question mark and their cells have several options for values. We later explain how this distinction has crucial impact on the complexity of CIRs.

A natural scenario, studied by previous work for the TID model [11, 19], considers a probabilistic database in the presence of a given set of integrity constraints, and specifically, Functional Dependencies (FDs). Such a scenario gives rise to several interesting computational challenges, and we focus here on three basic ones. In the problem of *possible consistency*, the goal is to test for the existence of a possible world (with a nonzero probability) that satisfies the FDs. The problem of the *most probable database* (“MPD” [11]) is that of finding a possible world that satisfies the FDs and has the highest probability. In the problem of computing the *probability of consistency*, the goal is to calculate the above probability exactly (beyond just deciding whether it is nonzero), that is, the probability that (a random sample of) the given CIR satisfies the underlying FDs. We investigate the computational complexity of these three problems for the CIR model. Our results provide classifications of tractability for different classes of FDs. Importantly, we show that, for the studied classes, the complexity of these problems is determined by two factors: (1) the location of the uncertain attributes in the FDs (left or right side), and (2) the combination of the FDs in the given set of constraints.

The three problems relate to each other in the following manner. To solve MPD, we need to be able to solve the possible consistency problem. The analysis of the probability of consistency sheds light on the possible consistency problem (is it fundamentally harder to compute the probability than to just determine whether it is nonzero?), but its importance goes beyond that. As we explain in Section 4, computing this probability is useful to any type of constraints over CIRs, as the tractability of this probability implies that we can efficiently sample correctly from the conditional space of consistent samples.

Our study adopts the standard yardstick of *data complexity* [31], where we fix the relational schema and the set of functional dependencies. The schema mentions not only what attributes are in the header of the relation, but also *which attribute is certain and which attribute is uncertain*. The complexity of the problems can be different for different combinations of schema and constraints, and we aim for a detailed understanding of which combinations are tractable and which are not.

► **Example 1.** Consider again the CIR U_1 in Figure 1a along with the FD set F_1 of Figure 1b,

consisting of a single FD. The FD says that at a specific time, a specialist can be found in only one location. Figure 1c shows a consistent sample r of U_1 whose probability is $\Pr(r) = 0.5 \cdot 0.7 \cdot 0.8 = 0.28$. In particular, this probability is nonzero and, so, U_1 is possibly consistent. This sample has a maximal probability among the consistent samples; therefore, r is a most probable database for U_1 . Now, consider the FD set F_2 shown in Figure 1b, where the first FD is the one of F_1 and the second states that no two specialists should be in the same room at the same time. The sample r in Figure 1c is no longer consistent, but r' (in the same figure) is a consistent sample and also a most probable database. In fact, r' is the only consistent sample in this case, so the probability of consistency for F_2 turns out to be that of r' . ♦

In contrast to the state of affairs in the attribute-level uncertainty, for *tuple-level uncertainty* much more is known about MPD (i.e., finding the most likely instance of a probabilistic database conditioned on conformance to a set of FDs). In the case of tuple-independent databases, Gribkoff, Van den Broeck, and Suciu [11] established dichotomy in the complexity of MPD for sets of unary FDs. This dichotomy has been generalized by Livshits, Kimelfeld and Roy [19] to a full classification over all sets of FDs, where they also established that the problem is equivalent to finding a *cardinality repair* of an inconsistent database. Carmeli et al. [5] showed that two tractable cases, namely a *single FD* and a *matching constraint*, remain tractable even if the FDs are treated as *soft constraints* (where every violation incurs a cost). In this work, we aim to bring our understanding of attribute-level uncertainty closer to tuple-level uncertainty.

Results. We establish classification results on several classes of functional dependencies: singleton FDs, matching constraints (i.e., FD sets of the form $\{X \rightarrow Y, Y \rightarrow X\}$), and arbitrary sets of unary FDs. Each classification consists of three internal classifications—one for each of the three problems we study (possible consistency, most probable database, and the probability of consistency). In every case, finding a most probable database is tractable whenever possible consistency is tractable. There are cases where the probability of consistency is intractable in contrast to the tractability of the most probable database, but we did not find any case where the other direction holds (and we will be surprised if such case exists). We also establish some general conclusions beyond these classes. For example, in Theorem 13 (of Section 5) we claim that if we make no assumption that some attributes are certain, then possible consistency is hard for *every nontrivial set of FDs*.

► **Example 2.** Reconsider the CIR U_1 in Figure 1a along with the FD set F_1 of Figure 1b, consisting of a single FD. Our classification shows that, in general, finding a solution to the possible consistency problem for such an FD, with uncertain attributes on the left side, is NP-complete. Now, reconsider the FD set F_2 shown in Figure 1b, where the first FD is the one of F_1 . Thus, $F_1 \subset F_2$, however, interestingly, our results show that for sets with the structure of F_2 , finding an MPD (and, hence, also solving possible consistency) is in polynomial time. Intuitively, the additional FD in F_2 constrains the uncertain attribute on the left side of the first FD, making the problem tractable. Finally, computing the probability of consistency for sets with the structure of F_1 and F_2 is #P-hard (or more precisely FP^{#P}-complete). ♦

Related work. A most probable database is the same as the “Most Likely Intention” (MLI) in the framework of Probabilistic Unclean Databases (PUD) of De Sa et al. [26], in the special case where the *intention model* demands hard constraints and the *realization model*

applies random changes to cells independently in what they refer to as *parfactor/update* PUD. They showed that finding an MLI of a parfactor/update PUD generalizes the problem of finding an *update repair* of an inconsistent database with a minimum number of value changes. In turn, finding a minimal update repair has been studied in the literature and several complexity results are known for special cases of FDs, such as hardness (e.g., for the FD set $\{A \rightarrow B, B \rightarrow C\}$ due to Kolahi and Lakshmanan [17]) and tractability (e.g., for lhs-chains such as $\{A \rightarrow B, AD \rightarrow C\}$ due to Livshits et al. [19]). There are, though, substantial differences between finding a most probable consistent sample of a CIR and finding an optimal update repair of an inconsistent database, at least in the variations where complexity results are known. First, they allow to select *any* value (from an infinite domain) for a cell, in contrast to the distributions of the CIR that can limit the space of allowed values; indeed, this plays a major role in past repair algorithms (e.g., Proposition 5.6 of [19] and Algorithm FINDVREPAIR of [17]). Second, they allow to change the value of *any* attribute and do not distinguish between uncertain attributes (where changes are allowed) and certain ones, as we do here; this is critical since, again, without such assumptions the problem is intractable for every nontrivial set of FDs (Theorem 13).

The problem of possible consistency does not have a nontrivial correspondence in the tuple-independent database model since, there, if there is any consistent sample then the subset that consists of all deterministic tuples (i.e., ones with probability one) is such a sample. The probability of consistency might be reminiscent of the problem of *repair counting* that was studied for subset repairs [4, 20]. Besides the fact that subset repairs are about tuple-level uncertainty (and no probabilities are involved), here we do not have any notion of *maximality* (while a repair is required to be a maximal consistent subset).

A CIR can be easily translated into a relation of a *block-independent-disjoint* (BID) probabilistic database [24]. In a BID, every relation is partitioned into independent blocks of mutually exclusive tuples, each associated with a probability. This model has also been studied under the terms *dirty database* [2] and *x-tuples* [6, 21, 23]. This translation implies that every upper bound for BIDs applies to CIRs, and the contrapositive: every hardness result that we establish (e.g., for the most probable database) extends immediately to BIDs; yet, it does not imply the other direction. Moreover, we are not aware of any positive results on inference over BIDs regarding integrity constraints. In addition, the translation from a CIR to a BID loses the information of which attributes are certain and which are uncertain, and as aforesaid, if we allow every attribute to be uncertain then the problem is hard for every nontrivial set of FDs (Theorem 13).

Organization. The remainder of the paper is organized as follows. We begin with preliminary definitions and notation (Section 2). We then define the CIR data model (Section 3) and the computational problems that we study (Section 4). Next, we describe our analysis for the case of singleton and matching dependencies (Section 5), and then the case of unary functional dependencies (Section 6). Lastly, we give concluding remarks (Section 7). Some of the proofs are omitted for lack of space, and they can be found in the Appendix.

2 Preliminaries

We begin with preliminary definitions and notation.

Relations. We assume countably infinite sets **Val** of values and **Att** of attributes. A *relation schema* is a finite set $R = \{A_1, \dots, A_k\}$ of attributes. An *R-tuple* is a function $t : R \rightarrow \mathbf{Val}$

that maps each attribute $A \in R$ to a value that we denote by $t[A]$. A *relation* r is associated with a relation schema, denoted $\mathbf{Att}(r)$, a finite set of tuple identifiers, denoted $\mathbf{tids}(r)$, and a mapping from $\mathbf{tids}(r)$ to $\mathbf{Att}(r)$ -tuples. (Note we allow for duplicate tuples, as we do not assume that the tuples of different identifiers are necessarily different.) We say that r is a relation *over* the relation schema $\mathbf{Att}(r)$. We denote by $r[i]$ the tuple that r maps to the identifier i . Hence, $r[i][A]$ is the value that tuple i has for the attribute A . As an example, Figure 1c (left) depicts a relation r with $\mathbf{Att}(r) = \{\text{room}, \text{?specialist}, \text{time}\}$ (for now, the question mark in ?specialist should be ignored.) Here, $\mathbf{tids}(r) = \{1, 2, 3\}$ and $r[1][\text{room}] = 41$.

Suppose that X is a set of attributes. We denote by $\pi_X r$ the projection of r onto X . More precisely, $\pi_X r$ is the relation r' such that $\mathbf{Att}(r') = X$, $\mathbf{tids}(r') = \mathbf{tids}(r)$, and $r'[i][A] = r[i][A]$ for every $A \in \mathbf{Att}(r) \cap X$. Observe that in our notation, $(\pi_X r)[i]$ is the projection of tuple i to X . As a shorthand notation, we write $r[i][X]$ instead of $(\pi_X r)[i]$. For example, in Figure 1c we have $r[2][\text{room ?specialist}] = (163, \text{Bart})$.

Functional dependencies. A *functional dependency*, or FD for short, is an expression of the form $X \rightarrow Y$ where X and Y are finite sets of attributes. We say that $X \rightarrow Y$ is *over* a relation schema R if R contains all mentioned attributes, that is, $X \cup Y \subseteq R$. A relation r satisfies the FD $X \rightarrow Y$ over $\mathbf{Att}(r)$ if every two tuples that agree on X also agree on Y . In our notation, we say that r satisfies $X \rightarrow Y$ if for every two tuple identifiers i and i' in $\mathbf{tids}(r)$ it holds that $r[i][Y] = r[i'][Y]$ whenever $r[i][X] = r[i'][X]$. A relation r satisfies a set F of FDs over $\mathbf{Att}(r)$, denoted $r \models F$, if r satisfies every FD in F .

We use the standard convention that in instances of X and Y we may remove curly braces and commas. To compactly denote a set of FDs, we may also intuitively combine multiple FD expressions and change the direction of the arrows. For example, the notation $A \leftrightarrow B \leftarrow CD$ is a shorthand notation of $\{A \rightarrow B, B \rightarrow A, CD \rightarrow B\}$.

An FD $X \rightarrow Y$ is *unary* if X consists of a single attribute, and it is *trivial* if $Y \subseteq X$ (i.e., it is satisfied by every relation). A *matching constraint* (as termed in past work [5]) is a constraint of the form $X \leftrightarrow Y$, that is, the set $\{X \rightarrow Y, Y \rightarrow X\}$.

The *closure* F^+ of a set F of FDs is the set of all FDs that are implied by F (or, equivalently, can be inferred by repeatedly applying the axioms of Armstrong). For example, F^+ includes all of the trivial FDs. The closure X_F^+ of a finite set X of attributes is the set of all attributes A such that $X \rightarrow A$ is in F^+ . Two finite attribute sets X and Y are *equivalent* (w.r.t. F) if $X_F^+ = Y_F^+$, or in other words, $X \rightarrow Y$ and $Y \rightarrow X$ are both in F^+ . By a slight abuse of notation, we say that two attributes A and B are *equivalent* if $\{A\}$ and $\{B\}$ are equivalent. Finally, if F is a set of FDs, then we denote by $\mathbf{Att}(F)$ the set of all attributes that occur in either the left or right sides of rules in F .

Probability distributions. We restrict our study in this paper to finite probability spaces (Ω, π) where Ω is a nonempty finite set of *samples* and $\pi : \Omega \rightarrow [0, 1]$ is a probability function satisfying $\sum_{o \in \Omega} \pi(o) = 1$. The *support* of $\delta = (\Omega, \pi)$, denoted $\text{supp}(\delta)$, is the set of samples $o \in \Omega$ such that $\pi(o) > 0$. We denote by $\text{Pr}_\delta(o)$ the probability $\pi(o)$. We may write just $\text{Pr}(o)$ when δ is clear from the context.

3 Cell-Independent Relations

A *Cell-Independent Relation*, or *CIR* for short, is similar to an ordinary relation, except that in certain attributes the values may be probabilistic; that is, instead of an ordinary value, each of them contains a probability distribution over values. One could claim that the model

<i>tid</i>	business	?spokesperson	?location
1	S. Propane	Mangione(0.6) Strickland(0.4)	Arlen(0.6) McMaynerberry(0.4)
2	Mega Lo Mart	Mangione(0.45) Thatherton(0.55)	Arlen(0.5) McMaynerberry(0.5)
3	Mega Lo Mart	Mangione(0.4) Buckley(0.6)	Arlen(0.55) McMaynerberry(0.45)
4	Get In Get Out	Peggy(1.0)	Arlen(0.35) McMaynerberry(0.3) Dallas(0.35)

■ **Figure 2** CIR U_2 with spokesperson and location as the uncertain attributes.

should allow every attribute to have uncertain values. However, knowing which attributes are certain has a major impact on the complexity of operations over CIRs. Formally, a CIR U is defined similarly to a relation, with the following differences:

- The schema of U , namely $\mathbf{Att}(U)$, has *marked attributes* where uncertain values are allowed. We denote a marked attribute using a leading question mark, as in $?A$, and the set of marked attributes by $?Att(U)$. (Note that $?Att(U)$ is a subset of $\mathbf{Att}(U)$.)
- For every $i \in \mathbf{tids}(U)$ and marked attribute $?A \in ?Att(U)$, the cell $U[i][?A]$ is a probability distribution over \mathbf{Val} .

By interpreting cells as probabilistically independent, a CIR U represents a probability distribution over ordinary relations. Specifically, a sample of U is a relation that is obtained from U by sampling a value for each uncertain cell. More formally, a sample of U is a relation r such that $\mathbf{Att}(r) = \mathbf{Att}(U)$, $\mathbf{tids}(r) = \mathbf{tids}(U)$, and for every $i \in \mathbf{tids}(r)$ and unmarked attribute A we have that $r[i][A] = U[i][A]$.

The probability $\Pr_U(r)$ of a sample r of U is the product of the probabilities of the values chosen for r :

$$\Pr_U(r) = \prod_{i \in \mathbf{tids}(U)} \prod_{?A \in ?Att(U)} \Pr_{U[i][?A]}(r[i][?A])$$

Note that $\Pr_{U[i][?A]}(r[i][?A])$ is the probability of the value $r[i][?A]$ (i.e., the value that tuple i of r has for the attribute $?A$) according to the distribution $U[i][?A]$ (i.e., the distribution that tuple i of U has for the attribute $?A$).

► **Example 3.** Figures 1a and 2 depict examples U_1 and U_2 , respectively, of CIRs. U_1 has been discussed in Example 1 and U_2 describes a CIR that stores businesses along with their spokespeople and headquarters locations. Some information in U_2 is noisy (e.g., since the rows are scraped from Web pages), and particularly the identity of the spokesperson and the business location. U_1 has a single uncertain attribute, namely $?specialist$, and U_2 has two uncertain attributes, namely $?spokesperson$ and $?location$. In particular, we have:

$$\mathbf{Att}(U_1) = \{\text{room}, ?specialist, \text{time}\} \quad ?Att(U_1) = \{?specialist\}$$

Distributions over values are written straightforwardly in the examples. For example, the distribution $U_1[2][?specialist]$ is the uniform distribution that consists of **Bart** and **Lisa**, each with probability 0.5.

The relations r and r' of Figure 1c are samples of U_1 . By the choices made in r , the probability $\Pr_U(r)$ is $0.5 \cdot 0.7 \cdot 0.8$. Note that the probability of r is smaller than the probability of the sample where the specialists are **Lisa**, **Bart** and **Maggie**, for instance, respectively. ♦

Simplified notation. In the analyses that we conduct in later sections, we may simplify the notation when defining a CIR U . When $\mathbf{Att}(U) = \{A_1, \dots, A_k\}$, we may introduce a new

tuple $t[i]$ with $t[i][A_\ell] = a_\ell$ simply as (a_1, \dots, a_k) , assuming that the attributes are naturally ordered alphabetically by their symbols. For example, if $\mathbf{Att}(U) = \{A, B, C\}$, then (a, b, c) corresponds to the tuple that maps A , B and C to a , b and c , respectively. We can also use a distribution δ instead of a value a_ℓ . In particular, we write $b_1 | \dots | b_t$ to denote a uniform distribution among the values b_1, \dots, b_t .

► **Example 4.** Continuing Example 3, in the simplified notation the tuple $U_1[2]$ can be written as $(163, \text{Bart} | \text{Lisa}, 5 \text{ pm})$ since the attributes are ordered lexicographically and, again, the distribution happens to be uniform. ♦

Consistency of CIRs

Let F be a set of FDs and let U be a CIR, both over the same schema. A *consistent sample* of U is a relation $r \in \text{supp}(U)$ such that $r \models F$. We say that U is *possibly consistent* if at least one consistent sample exists. By *the probability of consistency*, we refer to the probability $\Pr_{r \sim U}(r \models F)$ that a random sample of U satisfies F . As a shorthand notation, we denote this probability by $\Pr_U(F)$. Note that U is possibly consistent if and only if $\Pr_U(F) > 0$. A consistent sample r is a *most probable database* (using the terminology of Gribkoff, Van den Broeck and Suciu [11]) if $\Pr(r) \geq \Pr(r')$ for every other consistent sample r' .

► **Example 5.** Consider the CIR U_1 of Figure 1a. Let F_1 be that of Figure 1b, saying that at a specific time, a specialist can be found in only one location. Figure 1c (left) shows a consistent sample r of U_1 . Then $\Pr(r) = 0.5 \cdot 0.7 \cdot 0.8 = 0.28$. In particular, this probability is nonzero, hence U_1 is possibly consistent. The reader can verify that r has a maximal probability among the consistent samples (and, in fact, among all samples); therefore, r is a most probable database for U_1 . To calculate the probability of consistency, we will take the complement of the probability of *inconsistency*. An inconsistent sample can be obtained in two ways: (1) selecting **Lisa** in both the first and second tuples, or (2) selecting **Bart** in the second tuple and in at most one of the first and the third (which we can compute as the complement of the product of the probabilities of selecting the others). Therefore,

$$\Pr_{U_1}(F_1) = 1 - (0.3 \cdot 0.5 + 0.7 \cdot (1 - 0.5 \cdot 0.8)).$$

Now suppose that we use F_2 of Figure 1b saying that, in addition to F_1 , a room can host only one specialist at a specific time. In this case, r is no longer a consistent sample since Room 41 hosts different specialists at 5 PM, namely **Lisa** and **Maggie**. The reader can verify that the only consistent sample now is r' of Figure 1c. In particular, U_1 remains possible consistent, the sample r' is the most probable database, and the probability of consistency is the probability of r' , namely $0.5 \cdot 0.3 \cdot 0.2$. ♦

4 Consistency Problems

We study three computational problems in the paper, as in the following definition.

► **Definition 6.** Fix a schema R and a set F of FDs over R . In each of the following problems, we are given as input a CIR U over R :

1. Possible-consistency: determine whether $\Pr_U(F) > 0$.
2. Most probable database: find a consistent sample with a maximum probability.
3. Probability of consistency: calculate $\Pr_U(F)$.

Observe that these problems include the basics of probabilistic inference: *maximum likelihood* computation and *marginal probability* calculation. An MPD can be viewed as an optimal completion of missing values, or an optimal correction of values suspected of being erroneous, assuming the independence of cells (as a *prior* distribution) and conditioned on satisfying the constraints (as a *posterior* distribution). A necessary condition for the tractability of the most probable database is possible consistency, where we decide whether at least one consistent sample exists. The problem of computing the probability of consistency can be thought of as a basic problem that sheds light on possible consistency. For example, if possible consistency is decidable in polynomial time in some case, is it because we can, generally, compute the probability of consistency or because there is something fundamentally easier with feasibility? We will see cases that feature both phenomena.

A more technical reason to why we wish to be able to compute the probability of consistency is that it guarantees the ability to *sample* soundly from the conditional probability distribution (the posterior), that is, have an efficient randomized algorithm that produces a consistent sample r with the probability $\Pr_U(r \models F)$. The idea is quite simple and applies to every condition F over databases, regardless of being FDs (and was used in different settings, e.g., [7]). For completeness of presentation, we give the details in the Appendix.

As aforesaid, the second and third problems are at least as hard as the first one: finding a most probable database of U requires knowing whether U is possibly consistent, and calculating the exact probability is at least as hard as determining whether it is nonzero. There is no reason to believe a-priori that their complexities are comparable. Yet, our analysis will show that the third has the same or higher complexity in the situations that we study.

4.1 Complexity Assumptions

In our complexity analysis, we will restrict the discussion to uncertain cells that are finite distributions represented explicitly by giving a probability for each value in the support. Note that if all uncertain cells of U have a finite distribution, then U has a finite set of samples. Yet, its size can be exponential in the number of rows of U (and also in the number of columns of U , though we will treat this number as fixed as we explain next), even if each cell distribution is binary (i.e., has only two nonzero options). Every probability is assumed to be a rational number that is represented using the numerator and the denominator.

We will focus on the *data complexity* of problems, which means that we will make the assumption that the schema R of the CIR and the set F of FDs are both fixed. Hence, every combination (R, F) defines a separate computational problem, and different pairs (R, F) can potentially have different complexities.

4.2 Preliminary Observations

In the following sections, we study the complexity of the three consistency problems that we defined in Definition 6. Before we move on to the actual results, let us state some obvious general observations.

- Possible consistency is in NP, since we can verify a “yes” instance U in polynomial time by verifying that a relation r is a consistent sample.
- If possible consistency is NP-complete for some schema R and set F of FDs, then it is NP-hard to find a most probable database, and it is NP-hard to compute the probability of consistency.
- We will show that the probability of consistency can be #P-hard, or more precisely

■ **Table 1** Complexity of the consistency problems for a binary schema. “Possibility” refers to possible consistency, “MPD” refers to the most probable database problem, and “Probability” refers to the probability of consistency.

FDs	Possibility	MPD	Probability	Propositions
$A \rightarrow ?B$	PTime	PTime	PTime	7
$?A \rightarrow B$	NP-complete	NP-hard	FP ^{#P} -complete	10
$A \leftrightarrow ?B$	PTime	PTime	FP ^{#P} -complete	8 (PTime), 9 (FP ^{#P} -c.)
$?A \leftrightarrow ?B$	NP-complete	NP-hard	FP ^{#P} -complete	10

FP^{#P}-complete.¹ Membership in FP^{#P} of the probability of consistency is based on our assumption that probabilities are represented as rational numbers, and it can be shown using standard techniques (e.g., [1, 10]) that we do not repeat here.

We will take the above for granted and avoid repeating the statements throughout the paper.

5 Singleton and Matching Constraints

In this section, we investigate the complexity of the three problems we study in two special cases: a singleton constraint $\{X \rightarrow Y\}$ and a *matching constraint* $X \leftrightarrow Y$ (as it has been termed in past work [5]). We give full classifications of when such constraints are tractable and intractable for the three problems. We note that we leave open the classification of the entire class of FD sets, but we provide it for the general case of unary FDs in Section 6.

We begin with the case of a binary schema, where every set of FDs is equivalent to either a singleton or a matching constraint.

5.1 The Case of a Binary Schema

Throughout this section, we assume that the schema is $\{A, B\}$. The complexity of the different cases of FDs is shown in Table 1. To explain the entries of the table, let us begin with the tractable cases.

5.1.1 Algorithms

In this section, we show algorithms for $A \rightarrow ?B$ and for $A \leftrightarrow ?B$.

For $A \rightarrow ?B$, we need to determine a value b for each value a of the attribute A . The idea is that we do so independently for each a . Let V_A be the active domain of the attribute A of U , and V_B be the set of all values in the supports of the distributions of B . Formally:

$$V_A := \{U[i][A] \mid i \in \text{tids}(U)\} \quad V_B := \bigcup \{\text{supp}(U[i][B]) \mid i \in \text{tids}(U)\}$$

A consistent sample r selects a value $b_a \in V_B$ for each $a \in V_A$, and then $\Pr_U(r) = \prod_{a \in V_A} p(a, b_a)$ where $p(a, b)$ is given by:

$$p(a, b) := \prod_{i: U[i][A]=a} \Pr_{U[i][B]}(b)$$

¹ Recall that FP^{#P} is the class of functions that are computable in polynomial time with an oracle to a problem in #P (e.g., counting the number of satisfying assignments of a propositional formula). This class is considered intractable, and above the polynomial hierarchy [28].

Therefore, to find a most probable database, we consider each $a \in V_A$ independently, and find a $b \in V_B$ that maximizes $p(a, b)$. This b will be used for the tuples with the value a in A . In addition, we have the following formula that gives us immediately a polynomial-time algorithm (via a direct computation) for the probability of consistency:

$$\Pr_U(\{A \rightarrow B\}) = \prod_{a \in V_A} \sum_{b \in V_B} p(a, b)$$

Where $\sum_{b \in V_B} p(a, b)$ is the probability that the tuples with the value a for A agree on their B attribute. In summary, we have established the following.

► **Proposition 7.** *All three problems in Definition 6 are solvable in polynomial time for $A \rightarrow ?B$.*

Next, we discuss $A \leftrightarrow ?B$. Let U be a CIR. A consistent sample r of U entails the matching of each A value a to each B value b , so that no two a values occur with the same b , and no two b s occur with the same a . Therefore, we can solve this problem using an algorithm for minimum-cost perfect matching, as follows. Let V_A , V_B and $p(a, b)$ be as defined earlier in the section for $A \rightarrow ?B$. We construct a complete bipartite graph G as follows.

- The left-side vertex set is V_A and the right-side vertex set is V_B .
- The cost of every edge (a, b) is $(-\log p(a, b))$; we use this weight as our goal is to translate a maximum product into a minimum sum.²

Note that $|V_A|$ and $|V_B|$ are not necessarily of the same cardinality. If $|V_A| > |V_B|$, then U has no consistent sample at all. If $|V_A| < |V_B|$, then we add to the left side of the graph dummy vertices a' that are connected to all V_B vertices using the same cost, say 1. With this adjustment, we can now find a most probable database by finding a minimum-cost perfect matching in G . In summary, we have established the following.

► **Proposition 8.** *For $A \leftrightarrow ?B$, a most probable database can be found in polynomial time.*

It turns out that the third problem, the probability of consistency, is intractable. We show it in the next section.

5.1.2 Hardness

We now discuss the hardness results of Table 1. We begin with $A \leftrightarrow ?B$. Recall that possible consistency and the most probable database are solvable in polynomial time (Proposition 8). The probability of consistency, however, is hard.

► **Proposition 9.** *For $A \leftrightarrow ?B$, it is $\text{FP}^{\#\text{P}}$ -complete to compute the probability of consistency.*

The proof of Proposition 9 is by a reduction from counting the perfect matchings of a bipartite graph, which is known to be $\#\text{P}$ -complete [30]. The next proposition addresses the case of $?A \rightarrow B$ and the case of $?A \leftrightarrow ?B$.

► **Proposition 10.** *For each of $?A \rightarrow B$ and $?A \leftrightarrow ?B$:*

1. *Possible consistency is NP-complete.*
2. *The probability of consistency is $\text{FP}^{\#\text{P}}$ -complete.*

² We assume that the computational model for finding a minimum-cost perfect matching can handle the representation of logarithms, including $\log 0 = -\infty$. As an alternative, we could use directly an algorithm for maximizing the product of the edges in the perfect matching [29].

For possible consistency, we show reductions from variations of SAT, where we translate satisfying assignments of a formula to consistent samples of a CIR. For the probability of consistency we use a similar idea to Proposition 9. The proofs of Proposition 9 and Proposition 10 can be found in the Appendix.

We have now completed all results of Table 1. We will use these results for the extension to singleton, matching, and unary constraints.

5.2 Singleton and Matching Constraints

We generalize the results for the binary case to the more general case where the FD set is either a singleton or a matching constraint.

► **Theorem 11.** *Let X and Y be sets of attributes such that $X \not\subseteq Y$ and $Y \not\subseteq X$, and at least one attribute in $X \cup Y$ is uncertain.*

1. *In the case of $X \rightarrow Y$, if X consists of only certain attributes, then all three problems are solvable in polynomial time. Otherwise, possible consistency is NP-complete and the probability of consistency is $\text{FP}^{\#P}$ -complete.*
2. *In the case of $X \leftrightarrow Y$, if either X or Y consists of only certain attributes, then a most probable database can be found in polynomial time; otherwise, possible consistency is NP-hard. In any case, the probability of consistency is $\text{FP}^{\#P}$ -complete.*

Proof sketch. For the first part, the hardness side is due to a straightforward reduction from $?A \rightarrow B$, where hardness is stated in Proposition 10, and for the tractability side, we show a reduction to the case of $A \rightarrow ?B$, which is tractable due to Proposition 7. For the second part, the tractability side is via a reduction to the case of $A \leftrightarrow ?B$, which is tractable due to Proposition 8. The hardness of possible consistency relies on the cases of $?A \rightarrow B$ and $?A \leftrightarrow ?B$ from Proposition 10, and the hardness of probability of consistency relies on the case $A \leftrightarrow ?B$ from Proposition 9. The full proof can be found in the Appendix. ◀

► **Example 12.** Consider again the CIR U_1 of Figure 1a, and the following two constraints: $F_1 := \{?\text{specialist time} \rightarrow \text{room}\}$ and $F_2 := F_1 \cup \{\text{room time} \rightarrow ?\text{specialist}\}$. For F_1 , all three problems are hard, since the left hand side of the FD contains the uncertain attribute $?\text{specialist}$. For F_2 , a most probable database can be found in polynomial time, since F_2 is equivalent to $\text{room time} \leftrightarrow ?\text{specialist time}$, where one side (the left side) consists of only certain attributes. However, the probability of consistency remains $\text{FP}^{\#P}$ -hard. ♦

Note that in Theorem 11, the assumption that $X \not\subseteq Y$ and $Y \not\subseteq X$ does not lose generality, for the following reason. If $X \subseteq Y$, then the FD $X \rightarrow Y$ is equivalent to $X \rightarrow Y \setminus X$, the FD $Y \rightarrow X$ is trivial, and the matching constraint $X \leftrightarrow Y$ is equivalent to the singleton $\{X \rightarrow Y\}$ (which is covered in Part 1).

From Theorem 11 we can conclude that when all attributes are uncertain, possible consistency is hard, unless the FDs are all trivial (and then all three problems are clearly solvable in polynomial time); this is under the reasonable (and necessary) assumption that F has no *consensus FDs*, that is, the left hand side of every FD is nonempty [19]. We later discuss this assumption. This emphasizes the importance of having a data model that distinguishes between certain and uncertain attributes.

► **Theorem 13.** *Let F be a nontrivial set of FDs over a relation schema R , none being a consensus FD. Then possible consistency is NP-complete.*

The proof (in the Appendix) selects between a reduction from MPD with the singleton $\{?A \rightarrow ?B\}$ and a reduction from MPD with the matching constraint $?A \leftrightarrow ?B$, depending on the structure of the F .

We note that the assumption that F has no consensus FDs is necessary. For example, for $F = \{\emptyset \rightarrow ?A\}$, which is nontrivial, we can find a most probable database by considering every possible value a for $?A$, computing the probability of selecting a in all distributions, and finally using the value with the maximal probability.

From Theorem 13 we immediately conclude the hardness of the three problems on every nontrivial set of FDs in the block-independent-disjoint (BID) model of probabilistic databases [24], due to the translation mentioned in the Introduction.

6 General Sets of Unary Functional Dependencies

In Section 5.1, we studied the complexity of the three problems in the case of a binary schema, and we gave a full classification of the different possible sets of FDs. In this section, we extend these results to a general classification (dichotomy) for every set of *unary* FDs, that is, FDs with a single attribute on the left side. Our result uses a decomposition technique that we devise next.

6.1 Reduction by Decomposition

In this section, we devise a decomposition technique that allows us to reduce our computational problems from one set of FDs into multiple smaller subsets of the set. This technique is stated in the next theorem. After the theorem, we show several consequences that illustrate the use of the technique. Later, we will use these consequences to establish a full classification of complexity for the sets of unary FDs.

► **Theorem 14.** *Let F be a set of FDs over a relation schema R . Suppose that $F = F_1 \cup F_2$ and that all attributes in $\mathbf{Att}(F_1) \cap \mathbf{Att}(F_2)$ are certain (unmarked). Each of the three problems (in Definition 6) can be solved in polynomial time if its version with F_j and $\mathbf{Att}(F_j)$ is solvable in polynomial time for both $j = 1$ and $j = 2$.*

Proof sketch. Let $U_j = \pi_{\mathbf{Att}(F_j)} U$ for $j = 1, 2$. We show the following:

1. U is possibly consistent w.r.t. F if and only if U_1 and U_2 are possibly consistent w.r.t. F_1 and F_2 , respectively.
2. MPDs of U_1 and U_2 can be easily combined to produce an MPD of U .
3. $\Pr_U(F) = \Pr_{U_1}(F_1) \cdot \Pr_{U_2}(F_2)$.

The full details are provided in the Appendix. ◀

An immediate conclusion from Theorem 14 is that we can eliminate the FDs that involve only certain attributes if we know how to deal with the remaining FDs.

► **Corollary 15.** *Let F be a set of FDs over a relation schema R . Let $X \rightarrow Y$ be an FD in F , and suppose that all attributes in X and Y are certain. Then each of the three problems (in Definition 6) is polynomial-time reducible to its version with R and $F \setminus \{X \rightarrow Y\}$.*

► **Remark 16.** Eliminating the FDs over the certain attributes is not always beneficial, since these FDs might be needed for applying a polynomial-time algorithm. As an example, consider the following set of FDs: $\{?A \rightarrow B, B \rightarrow C, C \rightarrow ?A\}$. As we will show later, for this set of FDs we can find a most probable database in polynomial time. However, we will also show that possible consistency is NP-hard for the subset $\{?A \rightarrow B, C \rightarrow ?A\}$. Hence, $B \rightarrow C$ is needed for the polynomial-time algorithm. ◀

The following consequence of Theorem 14 identifies a general tractable case: the problems are solvable in polynomial time if uncertain attributes do not appear in the left side of the FDs (but they can appear in the right side or outside of the FDs).

► **Theorem 17.** *Let F be a set of FDs. If the left side of every FD includes only certain attributes, then each of the three problems (in Definition 6) is solvable in polynomial time.*

Proof sketch. Assume, without loss of generality, that each FD in F contains a single attribute on the right side. For every $A \in \mathbf{Att}(F)$, let F_A be the subset of F that contains all FDs with A being the right side (i.e., all FDs of the form $X \rightarrow A$). Then $F = \cup_{A \in \mathbf{Att}(F)} F_A$. Note that sets F_A and F_B , where $A \neq B$, share only certain attributes. This is true since our assumption implies that an uncertain attribute $?A$ can appear only in $F_{?A}$. Hence, we can apply Theorem 14 repeatedly and conclude that we need a polynomial-time solution for each $F_{?A}$. In the Appendix, we show that we can obtain that using a similar concept to the algorithm for $A \rightarrow ?B$ from Section 5.1.1. ◀

6.2 Classification

We now state the precise classification of the complexity of the problems in the case of unary FDs. The statement uses the following terminology. Let F be a set of unary FDs. Recall that two attributes A and B are *equivalent* if they have the same closure, that is, $\{A\}_F^+ = \{B\}_F^+$. An attribute A is called a *sink* if $\{A\}_F^+ = \{A\}$, that is, A does not appear in the left hand side of any nontrivial FD. In this section, we will prove the following classification (trichotomy) result, which is also illustrated in Figure 3.

► **Theorem 18.** *Let F be a set of unary FDs over a relation schema R . Then following hold.*

1. *If every uncertain attribute is either a sink or equivalent to a certain attribute, then a most probable database can be found in polynomial time; otherwise, possible consistency is NP-complete.*
2. *If every uncertain attribute is a sink, then the probability of consistency can be calculated in polynomial time; otherwise, it is $\text{FP}^{\#\text{P}}$ -complete.*

The following examples illustrate the instantiation of the theorem to specific scenarios.

► **Example 19.** We give several examples for the case of a ternary schema $\{A, B, C\}$. Consider the following sets of FDs:

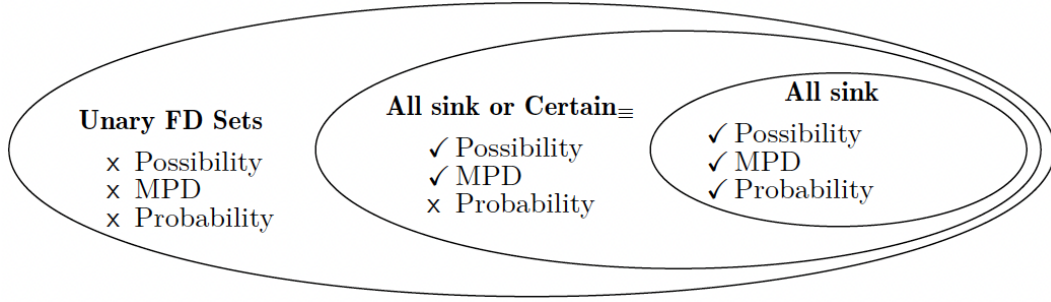
$$F_1 := \{A \rightarrow B \rightarrow ?C\} \quad F_2 := \{A \rightarrow ?B \rightarrow C\} \quad F_3 := \{A \leftrightarrow ?B \rightarrow ?C\}$$

Theorem 18 tells us that the following. All three problems are solvable in polynomial time in the case of F_1 , since the uncertain attribute $?C$ is a sink. In the case of F_2 , we can see that $?B$ is neither a sink nor equivalent to any certain attribute; hence, all three problems are intractable for F_2 . In the case of F_3 , it holds that $?C$ is a sink and $?B$ is not a sink but is equivalent to the certain attribute A . Hence, the consistency of F_3 is $\text{FP}^{\#\text{P}}$ -complete, but we can find a most probable database in polynomial time. ♦

► **Example 20.** Consider again the CIR U_2 of Figure 2. Consider the following constraints.

1. $\text{business} \rightarrow ?\text{spokesperson}?\text{location}$
2. $?\text{spokesperson} \rightarrow ?\text{location}$
3. $\text{business} \leftrightarrow ?\text{spokesperson} \rightarrow ?\text{location}$

For the first constraint, all three problems are tractable since both $?spokesperson$ and $?location$ are sinks. For the second constraint, all three problems are intractable since $?spokesperson$ is neither a sink nor equivalent to any certain attribute. For the third constraint, a most



■ **Figure 3** Classification of the complexity of consistency problems for sets of unary FDs. (See Table 1 for the naming of the problems.) “All sink” refers to the case where every uncertain attribute is sink, and “All sink or Certain₌” refers to the case where every uncertain attribute is either a sink or equivalent to a certain attribute.

probable database can be found in polynomial time since ?spokesperson is equivalent to the certain business and ?location is a sink, but the probability of consistency is $\text{FP}^{\#P}$ -complete since ?spokesperson is not a sink. ◆

In the remainder of this section, we prove each of the two parts of Theorem 18 separately.

6.2.1 Part 1 of Theorem 18 (Possible Consistency and MPD)

We first prove the tractability side of Part 1 of the theorem.

► **Lemma 21.** *Let F be a set of unary FDs over a schema R . If every uncertain attribute is either a sink or equivalent to a certain attribute, then a most probable database can be found in polynomial time.*

Proof sketch. The idea to define a set $F_{?A}$ of FDs for every uncertain attribute $?A \in ?\mathbf{Att}(U)$, and a set F' of FDs where all left-side attributes are certain, such that F is equivalent to $F' \cup \bigcup_{?A \in ?\mathbf{Att}(U)} F_{?A}$. Then, we repeatedly apply Theorem 14 to reduce the original problem to instances that are solvable in polynomial time by Proposition 8 and Theorem 17. ◀

For the hardness side of Part 1 of Theorem 18, we will need the following lemma, which generalizes the case of $?A \leftrightarrow ?B$ from Proposition 10.

► **Lemma 22.** *Let $R = \{?A_1, \dots, ?A_k\}$ consist of $k > 1$ uncertain attributes, and suppose that F is a set of FDs stating that all attributes in R are equivalent. Then possible consistency is NP-complete.*

The next lemma states the hardness side of Part 1 of Theorem 18.

► **Lemma 23.** *Let F be a set of unary FDs over a schema R . If there is an uncertain attribute that is neither a sink nor equivalent to a certain attribute, then possible consistency is NP-complete.*

Proof sketch. Let $?A$ be an attribute that is neither a sink nor equivalent to a certain attribute. Let X be the closure of $?A$ and X' be $X \setminus \{?A\}$. Observe the following. First, X' must be nonempty since $?A$ is not a sink. Second, if any attribute in X' implies $?A$ then it is equivalent to $?A$, and then it is necessarily uncertain. We consider two cases:

1. No attribute in X' implies $?A$.

2. Some attribute in X' implies $?A$.

For the first case, we show a reduction from $?A \rightarrow B$, where possible consistency is NP-complete due to Proposition 10. For the second case, let $?B_1, \dots, ?B_\ell$ be the set of all attributes in X' that imply $?A$. As said above, each $?B_j$ must be uncertain. Then all of $?B_1, \dots, ?B_\ell, ?A$ are equivalent. We show a reduction from the problem of Lemma 22 where $k = \ell + 1$. The constructions of the reductions for the two cases and the proofs of correctness can be found in the Appendix. ◀

6.2.2 Part 2 Theorem 18 (Probability of Consistency)

We now move on to Part 2. The tractability side follows immediately from Theorem 17, since if all uncertain attributes are sinks, then all left-side attributes are certain (up to trivial FDs $?A \rightarrow ?A$ that can be ignored). Hence, it remains to prove the hardness side of Part 2 of Theorem 18. We start with the following lemma, where we use a reduction from the case of $A \leftrightarrow ?B$, where probability of consistency is $\text{FP}^{\#P}$ -complete by Proposition 9, to establish hardness for a more general case.

► **Lemma 24.** *Let F be a set of unary FDs over a schema R . If at least one uncertain attribute is equivalent to a certain attribute, then the probability of consistency is $\text{FP}^{\#P}$ -complete.*

We can now complete the proof of the hardness side of Part 2.

► **Lemma 25.** *Let F be a set of unary FDs over a schema R . If there is at least one uncertain attribute that is not a sink, then the probability of consistency is $\text{FP}^{\#P}$ -complete.*

Proof sketch. Let $?A$ be an uncertain attribute that is not a sink. Let $Y = (?A)_F^+ \setminus \{?A\}$. Note that Y is nonempty, since $?A$ is not a sink. If any attribute B in Y functionally determines $?A$, then we can use this attribute as a certain attribute (even if it is uncertain) and use Lemma 24, since $?A$ is equivalent to B . Otherwise, suppose that no attribute in Y determines $?A$. For this case, we show (in the Appendix) a reduction from $?A \rightarrow B$, where the probability of consistency is $\text{FP}^{\#P}$ -hard according to Proposition 10. ◀

6.2.3 Recap

We can now complete the proof of Theorem 18. For Part 1, the tractability side is given by Lemma 21, and the hardness is given by Lemma 23. As for Part 2, the tractability side follows immediately from Theorem 17, and the hardness side is stated in Lemma 25.

7 Conclusions

We defined the concept of a CIR and studied the complexity of three problems that relate to consistency under FDs: possible consistency, finding a most probable database, and the probability of consistency. A seemingly minor feature of the definition of a CIR is the distinction between certain and uncertain attributes; yet, this distinction turns out to be crucial for detecting tractable cases. We gave classification results for several classes of FD sets, including a single FD, a matching constraint, and arbitrary sets of unary FDs. We also showed that if all attributes are allowed to be uncertain, then the first two problems are intractable for every nontrivial set of FDs.

This work leaves many problems for future investigation. Within the model, we have not yet completed the classification for the whole class of FD sets, where the problem remains open. Recall that a full classification is known for the most probable database for

tuple-independent databases [19]. Moreover, as we hit hardness already for simple cases (e.g., $?A \rightarrow B$), it is important to identify realistic properties of the CIR that reduce the complexity of the problems and allow for efficient algorithms.

Going beyond the framework of this paper, we plan to study additional types of constraints that are relevant to data cleaning [9], such as conditional FDs, denial constraints, and foreign-key constraints (where significant progress has been recently made in the problem of consistent query answering [13]). Another useful direction is to consider *soft* or *approximate* versions of the constraints, where it suffices to be consistent to some quantitative extent [5, 15, 18]. Finally, we have made the assumption of probabilistic independence among the cells as this is the most basic setting to initiate this research. To capture realistic correlations in the database noise, it is important to extend this work to data models that allow for (learnable) probabilistic dependencies, such as Markov Logic [16].

References

- 1 Serge Abiteboul, T.-H. Hubert Chan, Evgeny Kharlamov, Werner Nutt, and Pierre Senellart. Aggregate queries for discrete and continuous probabilistic XML. In *ICDT*, ACM International Conference Proceeding Series, pages 50–61. ACM, 2010.
- 2 Periklis Andritsos, Ariel Fuxman, and Renée J. Miller. Clean answers over dirty databases: A probabilistic approach. In *ICDE*, page 30, 2006.
- 3 Felix Bießmann, Tammo Rukat, Philipp Schmidt, Prathik Naidu, Sebastian Schelter, Andrey Taptunov, Dustin Lange, and David Salinas. Datawig: Missing value imputation for tables. *J. Mach. Learn. Res.*, 20:175:1–175:6, 2019.
- 4 Marco Calautti, Ester Livshits, Andreas Pieris, and Markus Schneider. Counting database repairs entailing a query: The case of functional dependencies. In *PODS*, pages 403–412. ACM, 2022.
- 5 Nofar Carmeli, Martin Grohe, Benny Kimelfeld, Ester Livshits, and Muhammad Tibi. Database repairing with soft functional dependencies. In *ICDT*, volume 186 of *LIPICs*, pages 16:1–16:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- 6 Reynold Cheng, Jinchuan Chen, and Xike Xie. Cleaning uncertain data with quality guarantees. *Proc. VLDB Endow.*, 1(1):722–735, 2008.
- 7 Sara Cohen, Benny Kimelfeld, and Yehoshua Sagiv. Incorporating constraints in probabilistic XML. *ACM Trans. Database Syst.*, 34(3):18:1–18:45, 2009.
- 8 Nilesh N. Dalvi and Dan Suciu. Efficient query evaluation on probabilistic databases. In *VLDB*, pages 864–875. Morgan Kaufmann, 2004.
- 9 Wenfei Fan and Floris Geerts. *Foundations of Data Quality Management*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, 2012.
- 10 Erich Grädel, Yuri Gurevich, and Colin Hirsch. The complexity of query reliability. In *PODS*, pages 227–234. ACM Press, 1998.
- 11 Eric Gribkoff, Guy Van den Broeck, and Dan Suciu. The most probable database problem. In *BUDA*, 2014. URL: <http://www.sigmod2014.org/buda>.
- 12 Venkatesan Guruswami. Inapproximability results for set splitting and satisfiability problems with no mixed clauses. In *APPROX*, volume 1913 of *LNCS*, pages 155–166. Springer, 2000.
- 13 Miika Hannula and Jef Wijsen. A dichotomy in consistent query answering for primary keys and unary foreign keys. In *PODS*, pages 437–449. ACM, 2022.
- 14 Alireza Heidari, Joshua McGrath, Ihab F. Ilyas, and Theodoros Rekatsinas. HoloDetect: Few-shot learning for error detection. In *SIGMOD Conference*, pages 829–846. ACM, 2019.
- 15 Abhay Kumar Jha, Vibhor Rastogi, and Dan Suciu. Query evaluation with soft-key constraints. In *PODS*, pages 119–128. ACM, 2008.
- 16 Abhay Kumar Jha and Dan Suciu. Probabilistic databases with markovviews. *Proc. VLDB Endow.*, 5(11):1160–1171, 2012.
- 17 Solmaz Kolahi and Laks V. S. Lakshmanan. On approximating optimum repairs for functional dependency violations. In *ICDT*, volume 361 of *ACM*, pages 53–62. ACM, 2009.
- 18 Ester Livshits, Alireza Heidari, Ihab F. Ilyas, and Benny Kimelfeld. Approximate denial constraints. *Proc. VLDB Endow.*, 13(10):1682–1695, 2020.
- 19 Ester Livshits, Benny Kimelfeld, and Sudeepa Roy. Computing optimal repairs for functional dependencies. *ACM Trans. Database Syst.*, 45(1):4:1–4:46, 2020. doi:10.1145/3360904.
- 20 Ester Livshits, Benny Kimelfeld, and Jef Wijsen. Counting subset repairs with functional dependencies. *J. Comput. Syst. Sci.*, 117:154–164, 2021.
- 21 Dany Maslowski and Jef Wijsen. A dichotomy in the complexity of counting database repairs. *J. Comput. Syst. Sci.*, 79(6):958–983, 2013.
- 22 Chris Mayfield, Jennifer Neville, and Sunil Prabhakar. ERACER: a database approach for statistical inference and data cleaning. In *SIGMOD Conference*, pages 75–86. ACM, 2010.
- 23 Luyi Mo, Reynold Cheng, Xiang Li, David W. Cheung, and Xuan S. Yang. Cleaning uncertain data for top-k queries. In *ICDE*, pages 134–145, 2013.

- 24 Christopher Ré and Dan Suciu. Materialized views in probabilistic databases for information exchange and query optimization. In *Proc. VLDB Endow.*, pages 51–62, 2007.
- 25 Theodoros Rekatsinas, Xu Chu, Ihab F. Ilyas, and Christopher Ré. HoloClean: Holistic data repairs with probabilistic inference. *PVLDB*, 10(11):1190–1201, 2017.
- 26 Christopher De Sa, Ihab F. Ilyas, Benny Kimelfeld, Christopher Ré, and Theodoros Rekatsinas. A formal framework for probabilistic unclean databases. In *ICDT*, volume 127 of *LIPICs*, pages 6:1–6:18, 2019.
- 27 Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch. *Probabilistic Databases. Synthesis Lectures on Data Management*. Morgan & Claypool Publishers, 2011.
- 28 Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- 29 Frank S.C. Tseng, Wei-Pang Yang, and Arbee L.P. Chen. Finding a complete matching with the maximum product on weighted bipartite graphs. *Computers & Mathematics with Applications*, 25(5):65–71, 1993.
- 30 Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979.
- 31 Moshe Y. Vardi. The complexity of relational query languages (extended abstract). In *STOC*, pages 137–146. ACM, 1982.
- 32 Jef Wijsen. Database repairing using updates. *ACM Trans. Database Syst.*, 30(3):722–768, 2005.

A Sampling Consistent Relations

As stated in Section 4, we now explain how computing the probability of consistency can help to sample correctly (i.e., with the correct probability) from the conditional probability where the condition is the satisfaction of a set F of constraints. Suppose that X_1, \dots, X_N are the random elements that represent the distributions of the uncertain cells of a CIR U . To produce a random sample r with probability $\Pr(r \mid F)$, we sample from the X_i one by one, from X_1 to X_N . When we sample from X_j , the probability of each value a is adjusted to be $\Pr(X_i = a \mid F \wedge X_1, \dots, X_{j-1} = a_1, \dots, a_{j-1})$ where a_1, \dots, a_{j-1} are the values chosen already for X_1, \dots, X_{j-1} .³ We can compute these adjusted probabilities if we know how to compute the probability of consistency. Specifically, by application of the Bayes rule, the adjusted probability can be represented as:

$$\frac{\Pr(X_1, \dots, X_j = a_1, \dots, a_{j-1}, a) \cdot \Pr(F \mid X_1, \dots, X_j = a_1, \dots, a_{j-1}, a)}{\Pr(X_1, \dots, X_{j-1} = a_1, \dots, a_{j-1}) \cdot \Pr(F \mid X_1, \dots, X_{j-1} = a_1, \dots, a_{j-1})}$$

Then, the probabilities $\Pr(F \mid X_1, \dots, X_\ell = b_1, \dots, b_\ell)$ are simply $\Pr_{U'}(F)$ where U' is obtained from U by replacing each X_j with the deterministic b_j .

B Proofs for Section 5 (Singleton and Matching Constraints)

► **Proposition 9.** *For $A \leftrightarrow ?B$, it is $\text{FP}^{\#P}$ -complete to compute the probability of consistency.*

Proof. We show a reduction from the problem of counting the perfect matchings of a bipartite graph (which is the same as calculating the permanent of a 0/1-matrix). This problem is known to be $\#P$ -complete [30]. We are given a bipartite graph $G = (V_L, V_R, E)$ such that $|V_L| = |V_R|$ and the goal is to compute the number of perfect matchings that G has. We construct a CIR U as follows. For each vertex $v \in V_L$ we collect the set $N_v \subseteq V_R$ of neighbors of v . Let $N_v = \{u_1, \dots, u_\ell\}$. We add to U the tuple $(v, u_1 \mid \dots \mid u_\ell)$.

Observe that every consistent sample induces a perfect matching (due to $A \leftrightarrow ?B$), and vice versa. Hence, the number of consistent samples of U is the same as the number of perfect matchings of G . Since we used only uniform probabilities, every sample of U has the same probability, namely $1/(\prod_{v \in V_L} |N_v|)$. Therefore, the number of perfect matchings is $\Pr_U(A \leftrightarrow ?B) \cdot \prod_{v \in V_L} |N_v|$. ◀

► **Lemma 26.** *For $?A \rightarrow B$:*

1. *Possible consistency is NP-complete.*
2. *It is $\text{FP}^{\#P}$ -complete to compute the probability of consistency.*

Proof. We prove each part separately.

Part 1. We show a reduction from *non-mixed satisfiability* (NM-SAT), where each clause contains either only positive literals (“positive clause”) or only negative literals (“negative clause”). This problem is known to be NP-complete [12].

We are given a formula $c_1 \wedge \dots \wedge c_m$ over x_1, \dots, x_n . We construct an uncertain table as follows. For each positive $c_i = y_1 \vee \dots \vee y_\ell$ we have in the table the tuple

$$(y_1 \mid \dots \mid y_\ell, \mathbf{true}),$$

³ This is true because the probability of a sample r of U can be written as $\Pr(X_1, \dots, X_N = a_1, \dots, a_N \mid F)$, which is equal to $\prod_{j=1}^N \Pr(X_j = a_j \mid F \wedge X_1, \dots, X_{j-1} = a_1, \dots, a_{j-1})$.

that is, a tuple with a distinct identifier i such that $U[i][A]$ is a uniform distribution over $\{y_1, \dots, y_\ell\}$ and $U[i][B]$ is the value **true**. Similarly, for each negative clause $c_i = \neg y_1 \vee \dots \vee \neg y_\ell$ we have in the table the tuple

$$(y_1 | \dots | y_\ell, \mathbf{false}).$$

Hence, for each positive clause we need to select one satisfying variable, for each negative clause we need to select one satisfying variable, and we cannot select the same variable to satisfy both a positive and a negative clause. The correctness of the reduction is fairly obvious.

Part 2. To prove Part 2, we use a reduction from counting the perfect matchings, similarly to the proof of Proposition 9, except that now we reverse the order of the attributes: Instead of adding the tuple $(v, u_1 | \dots | u_\ell)$, we add the tuple $(u_1 | \dots | u_\ell, v)$. The reader can easily verify that each consistent sample again encodes a unique perfect matching, and vice versa. ◀

► **Lemma 27.** For $?A \leftrightarrow ?B$:

1. Possible consistency is NP-hard.
2. It is $\text{FP}^{\#P}$ -complete to compute the probability of consistency.

Proof. We prove each part separately.

Part 1. We need to show the NP-hardness of possible consistency. We show a reduction from standard SAT, where we are given a formula $\varphi = c_1 \wedge \dots \wedge c_m$ over x_1, \dots, x_n , and we construct a CIR U over $\{A, B\}$ as follows. For each clause $c = d_1 \vee \dots \vee d_\ell$ we add to U the tuple

$$(c, \langle c, d_1 \rangle | \dots | \langle c, d_\ell \rangle).$$

Note that the values of U are clauses c and pairs $\langle c, d \rangle$ where d is a literal. In addition to these tuples, we collect every two pairs $\langle c, d \rangle$ and $\langle c', d' \rangle$ such that d and d' are in conflict, that is, if $d = x$ then $d' = \neg x$ and if $d = \neg x$ then $d' = x$. For each such pair, we add to U the tuple

$$(\langle c, d \rangle | \langle c', d' \rangle, \langle c, d \rangle | \langle c', d' \rangle).$$

This completes the reduction. Next, we prove the correctness of the reduction, that is, φ is satisfiable if and only if U is possibly consistent.

For the “only if” direction, suppose that τ is a satisfying truth assignment for φ . We construct a consistent sample r as follows. For every tuples of the form $(c, \langle c, d_1 \rangle | \dots | \langle c, d_\ell \rangle)$, we choose for B a value $\langle c, d_i \rangle$ such that $\tau(d_i) = \mathbf{true}$. In the case of tuples of the form $(\langle c, d \rangle | \langle c', d' \rangle, \langle c, d \rangle | \langle c', d' \rangle)$, we choose the pair $\langle c', d' \rangle$ such that $\tau(d') = \mathbf{false}$ for both attributes A and B . We need to show that r satisfies $?A \leftrightarrow ?B$. It is easy to see why the left attribute determines the right attribute, and so, $?A \rightarrow ?B$ holds. Regarding $?B \rightarrow ?A$, we need to see verify that we do not have any conflicting tuples $(c, \langle c, d \rangle)$ and $(\langle c', d' \rangle, \langle c', d' \rangle)$ where $c = c'$ and $d = d'$. This is due to the fact that $\tau(d) = \mathbf{true}$ and $\tau(d') = \mathbf{false}$.

For the “if” direction, suppose that r is a consistent sample. We define a satisfying truth assignment τ as follows. Suppose that r contains $(c, \langle c, d \rangle)$. Then r necessarily contains $(\langle c', d' \rangle, \langle c', d' \rangle)$ for every c' that contains the negation d' of d . Therefore, r does not contain any $(c', \langle c', d' \rangle)$ where d' contradicts d . So, we choose τ such that $\tau(d) = \mathbf{true}$. If needed, we complete τ to the remaining variables arbitrarily. From the construction of τ it holds that every clause c is satisfied. This completes the proof of Part 1.

Part 2. Note that this part is immediate from Proposition 9, since every instance of $A \leftrightarrow ?B$ can be viewed as an instance of $?A \leftrightarrow ?B$ where all A values are known. ◀

► **Proposition 10.** For each of $?A \rightarrow B$ and $?A \leftrightarrow ?B$:

1. Possible consistency is NP-complete.
2. The probability of consistency is $\text{FP}^{\#\text{P}}$ -complete.

Proof. We proved each constraint separately in Lemma 26 and Lemma 27, respectively. ◀

► **Theorem 11.** Let X and Y be sets of attributes such that $X \not\subseteq Y$ and $Y \not\subseteq X$, and at least one attribute in $X \cup Y$ is uncertain.

1. In the case of $X \rightarrow Y$, if X consists of only certain attributes, then all three problems are solvable in polynomial time. Otherwise, possible consistency is NP-complete and the probability of consistency is $\text{FP}^{\#\text{P}}$ -complete.
2. In the case of $X \leftrightarrow Y$, if either X or Y consists of only certain attributes, then a most probable database can be found in polynomial time; otherwise, possible consistency is NP-hard. In any case, the probability of consistency is $\text{FP}^{\#\text{P}}$ -complete.

Proof. We prove each part separately.

Part 1. The hardness side is due to straightforward reduction from $?A \rightarrow B$ where hardness is stated in Lemma 26. For the tractability side, we show a reduction to the case of $A \rightarrow ?B$, which is tractable due to Proposition 7. Without loss of generality, suppose that all attributes of Y are uncertain (and we assume that all attributes of X are certain). Given an input U for $X \rightarrow Y$, we construct an instance U_0 for $A \rightarrow ?B$ with $\text{tids}(U_0) = \text{tids}(U)$ by converting each tuple of U into a tuple of U_0 . The construction is simple: for each identifier $i \in \text{tids}(U)$, the tuple $U_0[i]$ is obtained as follows:

- $U_0[i](A) = \pi_X U[i]$ (i.e., the concatenation of the values in the attributes of X);
- $U_0[i](?B)$ is the distribution over all possible tuples that can be generated from $\pi_Y U[i]$, each with its probability (i.e., the product of the values of the different attributes of Y).

The correctness of the reduction is straightforward.

Part 2. The tractability side is via a reduction to the case of $A \leftrightarrow ?B$, which is tractable due to Proposition 8. The reduction is the same as the one to $A \rightarrow ?B$ that we showed in Part 1.

For the hardness side, let us begin with possible consistency. We consider two cases. If X and Y share an uncertain attribute $?A$, then there are easy reductions from the case of $?A \leftrightarrow ?AB$ (by fixing all other attributes to a single constant), and $?A \leftrightarrow ?AB$ is equivalent to $?A \rightarrow B$ where hardness is due to Lemma 26. If X and Y do not share an uncertain attribute, then X and Y contain distinct uncertain attributes $?A$ and $?B$, and then we have easy reductions (again fixing all other attributes different from $?A$ and $?B$ to a single constant) from the case of $?A \leftrightarrow ?B$, where hardness is due to Lemma 27.

The above arguments also imply $\text{FP}^{\#\text{P}}$ -completeness of the probability of consistency, except for the case where all attributes of X are certain. In the latter case, we apply our easy reduction from $A \leftrightarrow ?B$ and use Proposition 9. ◀

► **Theorem 13.** Let F be a nontrivial set of FDs over a relation schema R , none being a consensus FD. Then possible consistency is NP-complete.

Proof. Let $X \rightarrow Z$ be a nontrivial FD in F such that X is minimal (with respect to containment) among the left hand sides of FDs in F . Let $Y = X_F^+ \setminus X$. Note that X is nonempty due to the assumption of the theorem. Also note that Y is nonempty since the closure of X contains every attribute in Z . We consider two cases:

1. $F \models Y \rightarrow ?A$ for some $?A \in X$.
2. $Y_F^+ = Y$.

Note that these are the only possible two cases, since $X \cup Y$ is closed under F .

We first select an attribute $?A \in X$ as follows. In the first case, we use the $?A$ of the case; in the second case, we choose an arbitrary $?A \in X$ (which exists due to the assumption that X is nonempty). In the first case we apply a reduction from $?A \leftrightarrow ?B$, and in the second we apply a reduction from $?A \rightarrow ?B$. Note that the two problems are hard for both constraints, according to Theorem 11. The reduction is detailed next.

Let U be an input CIR over $\{?A, ?B\}$. We construct a CIR U' over R with $\text{tids}(U') = \text{tids}(U)$ by transforming every tuple $U[i]$ into a tuple $U'[i]$, as follows. Let $c \in \mathbf{Val}$ be an arbitrary constant value. Let δ be a uniform distribution over all pairs $\langle a, b \rangle$ where $a \in \mathbf{Val}$ and $b \in \mathbf{Val}$ are values in the (ranges of) distributions in $U[?A]$ and $U[?B]$, respectively. For every attribute $?C$ of R ,

$$U'[?C] =: \begin{cases} U[i][?A] & \text{if } ?C = ?A; \\ U[i][?B] & \text{if } ?C \in Y; \\ c & \text{if } ?C \in X \setminus \{?A\}; \\ \delta & \text{otherwise.} \end{cases}$$

The correctness of the reduction is due to the following. First, if r' is a consistent sample of U' , then $\pi_{\{?A, ?B\}} r'$ is a consistent sample of U due to the construction of U' .

Second, every consistent sample r of U can be extended into a consistent sample r' of U' by choosing the value $\langle r[i][?A], r[i][?B] \rangle$ for each distribution δ in a cell of the tuple i . To see that $r' \models F$, consider a nontrivial FD $X' \rightarrow Y'$ and let i and j be two tuple identifiers in $\text{tids}(r')$. Suppose that $r'[i]$ and $r'[j]$ agree on the tuples of X' . Note that X' cannot be a strict subset of X due to the minimality of X . If X' includes at least one attribute outside of Y , then $r'[i][X]$ and $r'[j][X]$ include an occurrence of $r[i][?A]$ and, hence, they must be equal (due to the construction of r' and the consistency of r). If $X' \subseteq Y$ and we are in the first case, then $r'[i][X]$ and $r'[j][X]$ include an occurrence of $r[i][?B]$ and, hence, they must be equal. If $X' \subseteq Y$ and we are in the second case, then $Y' \subseteq Y$ and, hence, $r'[i][Y']$ and $r'[j][Y']$ are again the same tuples (due to the construction of r'). ◀

C Proofs for Section 6 (General Sets of Unary Functional Dependencies)

► **Theorem 14.** *Let F be a set of FDs over a relation schema R . Suppose that $F = F_1 \cup F_2$ and that all attributes in $\mathbf{Att}(F_1) \cap \mathbf{Att}(F_2)$ are certain (unmarked). Each of the three problems (in Definition 6) can be solved in polynomial time if its version with F_j and $\mathbf{Att}(F_j)$ is solvable in polynomial time for both $j = 1$ and $j = 2$.*

Proof. We start with the problem of possible consistency. Let $U_j = \pi_{\mathbf{Att}(F_j)} U$ for $j = 1, 2$, and let $U' = \pi_{\mathbf{Att}(U) \setminus \mathbf{Att}(F)} U$. We show that U is possibly consistent w.r.t. F if and only if U_j is possibly consistent w.r.t. F_j for $j = 1, 2$. If r is a consistent sample of U then $\pi_{\mathbf{Att}(F_1)} r$ is a consistent sample of U_1 w.r.t. F_1 and $\pi_{\mathbf{Att}(F_2)} r$ is a consistent sample of U_2 w.r.t. F_2 .

Next, assume that r_1 is a consistent sample of U_1 w.r.t. F_1 and r_2 is a consistent sample of U_2 w.r.t. F_2 . Let r' be a sample in the support of U' . We show that $r_1 + r_2 + r'$ is a consistent sample of U w.r.t. F , where $r_1 + r_2 + r'$ is the relation r such that $\text{tids}(r) = \text{tids}(r_1)$, that $\mathbf{Att}(r) = \mathbf{Att}(r_1) \cup \mathbf{Att}(r_2) \cup \mathbf{Att}(r')$, and that $r[i] = r_1[i] \cup r_2[i] \cup r_3[i]$ (i.e., the natural join of the tuples $r_1[i], r_2[i]$ and $r'[i]$) for all $i \in \text{tids}(r)$. Observe that

$$\text{tids}(r_1) = \text{tids}(r_2) = \text{tids}(r') = \text{tids}(U_1) = \text{tids}(U_2) = \text{tids}(U') = \text{tids}(U),$$

and that for every tuple identifier i it is the case that the three relations (r_1, r_2 and r') agree on the common attributes of tuple i , since these are certain attributes. Therefore, r_1, r_2 and r' can be naturally combined to produce a relation. This relation is consistent since every FD is covered by one of the F_j s. Moreover, the addition of r' (that has a disjoint set of attributes) does not change the consistency.

For the problem of most probable database, if r_1 and r_2 are most probable databases of U_1 and U_2 , respectively, then by similar arguments, $r_1 + r_2 + r'$ is a most probable database of U for every sample r' in the support of U' .

Finally, for the probability of consistency, we show that $\Pr_U(F) = \Pr_{U_1}(F_1) \cdot \Pr_{U_2}(F_2)$. Let r denote the random element that corresponds to a sample of U , let r_1 be the random element $\pi_{\mathbf{Att}(F_1)}r$, and r_2 be the random element $\pi_{\mathbf{Att}(F_2)}r$. Then $r \models F$ if and only if $r_1 \models F_1$ and $r_2 \models F_2$. Moreover, r_1 and r_2 are probabilistically independent, since they involve disjoint sets of distributions, and in particular, their consistencies are probabilistically independent. This completes the proof. \blacktriangleleft

► Theorem 17. *Let F be a set of FDs. If the left side of every FD includes only certain attributes, then each of the three problems (in Definition 6) is solvable in polynomial time.*

Proof. Following the proof sketch, we show that we can solve each of the three problems for each $F_{?A}$ desperately. Consider the undirected graph G that has as vertices the tuple identifiers of U , and an edge between two identifiers i and i' whenever the two agree on all attributes of the left hand side of some dependency in $F_{?A}$. Hence, an edge means that the two tuples should have the same value in every consistent sample. Therefore, we can compute the most probable database, and calculate the probability of consistency, in a similar manner to the algorithms for $A \rightarrow ?B$ from Section 5.1.1, except that now we consider entire connected components rather than simple groups by the attribute A . \blacktriangleleft

► Lemma 21. *Let F be a set of unary FDs over a schema R . If every uncertain attribute is either a sink or equivalent to a certain attribute, then a most probable database can be found in polynomial time.*

Proof. Note that a sink uncertain attribute $?A$ cannot appear in the left side of any FD in F , unless it is the trivial FD $?A \rightarrow ?A$. We assume that F does not contain such trivial FDs, and so, a sink attribute can appear only on the right side of a rule. For every uncertain attribute $?A \in ?\mathbf{Att}(U)$ we define a set $F_{?A}$ of FDs as follows.

- If $?A$ is a sink or not at all in $\mathbf{Att}(F)$, then $F_{?A} = \emptyset$.
- If $?A$ is non-sink, then we find an equivalent certain attribute B and set $F_{?A} = \{?A \leftrightarrow B\}$.

Then F is equivalent to the set

$$F' \cup \bigcup_{?A \in ?\mathbf{Att}(U)} F_{?A} \tag{1}$$

where F' is obtained from F by replacing every non-sink uncertain attribute $?A$ by its partner B in $F_{?A}$. Now, observe the following.

- For each $F_{?A}$ we can find a most probable database in polynomial time, according to Proposition 8.
- In F' we have that all left-side attributes are certain, and so, we can find a most probable database according to Theorem 17.

Finally, observe that every uncertain attribute appears in at most one subset in Equation (1). Hence, we can repeatedly apply Theorem 14 and complete the proof. ◀

► **Lemma 22.** *Let $R = \{?A_1, \dots, ?A_k\}$ consist of $k > 1$ uncertain attributes, and suppose that F is a set of FDs stating that all attributes in R are equivalent. Then possible consistency is NP-complete.*

Proof. The case of $k = 2$, namely $?A_1 \leftrightarrow ?A_2$, has been shown in Lemma 27. For $k > 2$, we will show a reduction from $?A_1 \leftrightarrow ?A_2$. Given an input U_0 for $\{?A_1, ?A_2\}$, we construct an instance U for $\{?A_1, \dots, ?A_k\}$ by adding to U_0 columns to the right. We need $k - 2$ such columns. We set in each cell of these $k - 2$ columns the same distribution: the uniform distribution over m distinct values, where $m = |\text{tids}(U_0)|$ is the number of tuples in U_0 .

The correctness of the reduction is due to the following observations. First, if r is a consistent sample for U , then $\pi_{\{?A_1, ?A_2\}}r$ is a consistent sample for U_0 . Second, every consistent sample r_0 of U_0 can be extended into a feasible consistent sample of U by adding to each tuple a value that is uniquely determined by the values in the row. Third, the probability of every sample r of U is determined only by the values in $?A_1$ and $?A_2$:

$$\Pr_U(r) = \Pr_{U_0}(\pi_{\{?A_1, ?A_2\}}r) \cdot m^{-m(k-2)}$$

In particular, finding a consistent sample r for U is the same problem as finding a consistent sample r_0 for U_0 . ◀

► **Lemma 23.** *Let F be a set of unary FDs over a schema R . If there is an uncertain attribute that is neither a sink nor equivalent to a certain attribute, then possible consistency is NP-complete.*

Proof. We consider the two cases that are specified in the proof sketch of Lemma 23 (in the main body of the paper).

Case 1: No attribute in X' implies $?A$. We show a reduction from $?A \rightarrow B$, where possible consistency is NP-complete due to Proposition 10. Let U_0 be an input for $\{?A, B\}$ and $?A \rightarrow B$. We construct an input U for R and F as follows. The two CIRs have the same set of identifiers, that is, $\text{tids}(U) = \text{tids}(U_0)$. For each $i \in \text{tids}(U)$ we define the tuple $U[i]$ by:

- $U[i][?A] = U_0[i][?A]$;
- $U[i][B] = U_0[i][B]$ for every $B \in X'$ (even if B is uncertain);
- $U[i][C] = i$ for every $C \in R \setminus X$.

Note that we are using the tuple identifier i as a value to assure that every two tuples have different values in the corresponding C attributes. We will prove that U is possibly consistent if and only if U_0 is possibly consistent. From a consistent sample r for U we can get a consistent sample r_0 for U_0 by projecting on $?A$ and some attribute $C \in X$, and renaming C as B . Note that r_0 is a consistent sample of U_0 since r satisfies $?A \rightarrow C$.

For the other direction, suppose that r_0 is a consistent sample for U_0 . To obtain a consistent sample for U , we select each $?A$ value $U[i][?A]$ to be one chosen by r_0 , namely $r_0[i][?A]$. Denote the result by r . We need to show that r satisfies F . Let $D \rightarrow E$ an FD in F . If $D = ?A$ then $E \in X$ and then the FD is satisfied since r_0 satisfies $?A \rightarrow B$. If $D \in X'$ then $E \in X$, and $E \neq ?A$ since D and $?A$ are not equivalent. Hence, $E \in X'$ as well. It

thus follows that the E column is equal to the D column and, therefore, $D \rightarrow E$ is satisfied. Finally, if $D \notin X$ then no two tuples agree on D , and then $D \rightarrow E$ is satisfied in a vacuous manner.

Case 2: Some attribute in X' implies $?A$. Let $?B_1, \dots, ?B_\ell$ be the set of all attributes in X' that imply $?A$. As said above, each $?B_j$ must be uncertain. Then all of $?B_1, \dots, ?B_\ell, ?A$ are equivalent. We will show a reduction from the problem of Lemma 22 where $k = \ell + 1$. Let U_0 be an input for that problem; hence, U_0 is a CIR over $\{?A_1, \dots, ?A_k\}$. We construct a CIR U over R , as follows. Again, $\text{tids}(U) = \text{tids}(U_0)$ and for each $i \in \text{tids}(U)$ we define the tuple $U[i]$ according to the following rules. We use an arbitrary constant value that we denote by c .

- $U[i][?A] = U_0[i][?A_k]$;
- $U[i][?B_j] = U_0[i][?A_j]$ for $j = 1, \dots, \ell$;
- $U[i][C] = c$ for every $C \in X' \setminus \{?B_1, \dots, ?B_\ell\}$;
- $U[i][C] = i$ for every $C \in R \setminus X$.

Again, we show that U is possibly consistent if and only if U_0 is possibly consistent. From a consistent sample r for U we can get a consistent sample r_0 for U_0 by projecting on $\{?B_1, \dots, ?B_\ell, ?A\}$, renaming each $?B_j$ as $?A_j$ and $?A$ as $?A_k$. Note that r_0 is a consistent sample for U_0 (and in particular r_0 satisfies the FDs of Lemma 22) since r satisfies $?B_1 \leftrightarrow \dots \leftrightarrow ?B_\ell \leftrightarrow ?A$.

For the other direction, let r_0 be a consistent sample of U_0 . For each tuple identifier $i \in \text{tids}(U)$ we need to state the choice for the value $U[i][?A]$ and $U[i][?B_j]$ for $j = 1, \dots, \ell$. We set the former to $r_0[i][?A_k]$ and the latter to $r_0[i][?A_j]$. Let the result be r . We need to show that r satisfies F . Let $D \rightarrow E$ be an FD in F . We have the following cases:

- $D \notin X$. Then no two tuples agree on D , and then $D \rightarrow E$ is satisfied in a vacuous manner.
- $E \notin X$. Then $D \notin X$ (previous case) since X is a closure.
- $E \in X \setminus \{?A, ?B_1, \dots, ?B_\ell\}$. Then $D \rightarrow E$ holds since all values of the attribute E are equal (to c) in r .
- $D \in X \setminus \{?A, ?B_1, \dots, ?B_\ell\}$. Then $E \in X$ and E cannot be in $\{?A, ?B_1, \dots, ?B_\ell\}$, or otherwise D is also equivalent to $?A$ (while we know that only $?B_1, \dots, ?B_\ell$ are equivalent to $?A$).
- $\{D, E\} \subseteq \{?A, ?B_1, \dots, ?B_\ell\}$. Then $D \rightarrow E$ is satisfied since r_0 satisfies the constraint $?A_1 \leftrightarrow \dots \leftrightarrow ?A_k$.

We conclude that r satisfies F , and so, r is a consistent sample. This concludes the proof. ◀

► **Lemma 24.** *Let F be a set of unary FDs over a schema R . If at least one uncertain attribute is equivalent to a certain attribute, then the probability of consistency is $\text{FP}^{\#P}$ -complete.*

Proof. Let A be a certain attribute, and let $?B$ be an uncertain attribute that is equivalent to A . We will use Proposition 9 and show a reduction from $R_0 = \{A, ?B\}$ and $F_0 = \{A \leftrightarrow ?B\}$. Let U_0 be an input for R_0 and F_0 . We will produce an input U for R and F by adding columns with certain values, regardless of whether the attribute is certain or not, and we will show that the probability of consistency is the same in U and U_0 . The construction is straightforward: we simply copy the A column. In notation, for each identifier $i \in \text{tids}(U)$ and attribute $C \in R \setminus \{A, B\}$ we set $U[i][C] = U_0[i][A]$.

Observe that every sample r_0 of U_0 has a unique extension r_0^+ of U , and every sample of U is the unique extension r_0^+ of some sample r_0 of U_0 . To complete the proof, we show that both of the following hold:

1. $\Pr_{U_0}(r_0) = \Pr_U(r_0^+)$;
2. $r_0 \models F_0$ if and only if $r_0^+ \models F$.

Part 1 is straightforward from the construction, since all values of U are certain except for those in $?B$. So we now prove Part 2. The “if” direction is due to the fact that F implies both $A \rightarrow ?B$ and $?B \rightarrow A$. For the “only if” direction, suppose that $r_0 \models F_0$ and let $r = r_0^+$. We show that $r \models F$. Let $C \rightarrow D$ be an FD in F , let i and j be two identifiers, and suppose that $r[i][C] = r[j][C]$. We need to show that $r[i][D] = r[j][D]$. Due to our construction, $r[i][C] = r[j][C]$ implies that either $r[i][A] = r[j][A]$ or $r[i][?B] = r[j][?B]$ (in case C is $?B$). Since r satisfies $A \rightarrow ?B$, it holds that both $r[i][A] = r[j][A]$ and $r[i][?B] = r[j][?B]$ hold. Finally, observe that D is either $?B$ or a copy of A (from our construction). Therefore, $r[i][D] = r[j][D]$, as claimed. ◀

► **Lemma 25.** *Let F be a set of unary FDs over a schema R . If there is at least one uncertain attribute that is not a sink, then the probability of consistency is $\text{FP}^{\#P}$ -complete.*

Proof. Following the proof sketch, we assume that no attribute in Y determines $?A$. For this case, we show a reduction from $?A \rightarrow B$, where the probability of consistency is $\text{FP}^{\#P}$ -hard according to Proposition 10.

Let $R_0 = \{?A, B\}$ and $F_0 = \{?A \rightarrow B\}$. Let U_0 be an input for R_0 and U_0 . We will again produce an input U for R and F by adding columns with certain values, regardless of whether the attribute is certain or not, and we will show that the probability of consistency is the same in U and U_0 . The construction is as follows: we copy the B column to every attribute in Y , and use a unique fresh value for every other attribute. In notation, for each identifier $i \in \text{tids}(U)$ and attribute $C \in Y$ we set $U[i][C] = U_0[i][B]$, and for every attribute $C \in R \setminus (Y \cup \{?A\})$ we set $U[i][C] = i$.

We show the correctness of the reduction, using a similar argument as in the proof of Lemma 24. Note that our construction is such that every sample r_0 of U_0 has a unique extension r_0^+ of U , and every sample of U is the unique extension r_0^+ of some sample r_0 of U_0 . Then, the following hold:

1. $\Pr_{U_0}(r_0) = \Pr_U(r_0^+)$.
2. $r_0 \models F_0$ if and only if $r_0^+ \models F$;

Again, Part 1 is straightforward from the construction, so we prove Part 2. The “if” direction holds since $F \models ?A \rightarrow B$. For the “only if” direction, suppose that $r_0 \models F_0$ and let $r = r_0^+$. We show that $r \models F$. Let $C \rightarrow D$ be an FD in F , let i and j be two different identifiers, and suppose that $r[i][C] = r[j][C]$. We need to show that $r[i][D] = r[j][D]$. Note that the construction implies that C must be in the closure of $?A$, since $i \neq j$. If C is $?A$, then D belongs to Y , and then $r[i][D] = r[j][D]$ since r_0 is consistent. If C belongs to Y , then D also belongs to Y (since D cannot be $?A$), and then $r[i][D] = r[j][D]$ by construction. This completes the proof. ◀