

Top-k Querying of Unknown Values under Order Constraints (Extended Version)

Antoine Amarilli¹, Yael Amsterdamer², Tova Milo³, and
Pierre Senellart^{4,5}

- 1 LTCI, Télécom ParisTech, Université Paris-Saclay; Paris, France
first.last@telecom-paristech.fr
- 2 Bar Ilan University; Ramat Gan, Israel
first.last@biu.ac.il
- 3 Tel Aviv University; Tel Aviv, Israel
last@cs.tau.ac.il
- 4 DI, École normale supérieure, PSL Research University; Paris, France
first.last@ens.fr
- 5 Inria Paris; Paris, France

Abstract

Many practical scenarios make it necessary to evaluate top- k queries over data items with partially unknown values. This paper considers a setting where the values are taken from a numerical domain, and where some *partial order constraints* are given over known and unknown values: under these constraints, we assume that all possible worlds are equally likely. Our work is the first to propose a principled scheme to derive the value distributions and expected values of unknown items in this setting, with the goal of computing estimated top- k results by interpolating the unknown values from the known ones. We study the complexity of this general task, and show tight complexity bounds, proving that the problem is intractable, but can be tractably approximated. We then consider the case of tree-shaped partial orders, where we show a constructive PTIME solution. We also compare our problem setting to other top- k definitions on uncertain data.

1 Introduction

Many data analysis tasks involve queries over ordered data, such as maximum and top- k queries, which must often be evaluated in presence of *unknown data values*. This problem occurs in many real-life scenarios: retrieving or computing exact data values is often expensive, but querying the partially unknown data may still be useful to obtain approximate results, or to decide which data values should be retrieved next. In such contexts, we can often make use of *order constraints* relating the data values, even when they are unknown: for instance, we know that object A is preferred to object B (though we do not know their exact rating).

This paper thus studies the following general problem. We consider a set of numerical values, some of which are unknown, and we assume a *partial order* on these values: we may know that $x \geq y$ should hold although the values x or y are unknown. Our goal is to *estimate* the unknown values, in a principled way, and to evaluate top- k queries, namely find the items with (estimated) highest values.


Without further information, one may assume that every valuation compatible with the order constraints is equally likely, i.e., build a probabilistic model where valuations are *uniformly distributed*. Indeed, uniform distributions in the absence of prior knowledge is a common assumption in probabilistic data management [1, 11, 33] for continuous distributions on data values within an interval; here we generalize to a uniform distribution over multiple unknown values. Though the distribution is uniform, the dependencies between values lead to non-trivial insights about unknown values and top- k results, as we will illustrate.



© Antoine Amarilli; Yael Amsterdamer; Tova Milo; Pierre Senellart;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

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



■ **Figure 1** Sample catalog taxonomy with compatibility scores

Illustrative example. We consider a specific application setting where our problem occurs. Consider a scenario where products are classified in a catalog taxonomy (Figure 1) using human input: the relevance of a product to any category is captured by a *compatibility score*. Assessing this compatibility is often left to human judgement rather than attempting to derive it from records or statistics [8, 39, 48]. Thus we assume scores are obtained via questions to domain experts or to a crowd of unqualified users¹. Our goal is to assign the product to the top- k most compatible categories among a set of *end categories* (in yellow with a solid border), as opposed to *virtual categories* (dashed border). The virtual categories generalize the end categories, and allow us to ask broader questions to experts, but we do not try to assign products to them, e.g., they do not have a dedicated page in our online store.

Imagine now that the product to classify is a *smartwatch*, and that we want to find the top-2 end categories for it. We asked an expert for its compatibility score with some categories (both end and virtual categories), which we indicate in Figure 1. Because expert input is costly, however, we wish to choose the top-2 end categories based on the incomplete information that we have. The naïve answer is to look only at categories with known scores, and to identify *Wearable Devices* and *Diving Watches* as the best end categories.

In this scenario, however, we can impose a natural partial order over the scores, both known and unknown: any product that belongs to a specific category (e.g., *Smartphones*) conceptually also belongs to each of the more general categories (e.g., *Cell Phones*). We can thus require that if a category x is a sub-category of y , then the product’s compatibility score for x should be at most its score for y . This can be explained in the instructions to the humans providing the scores, and enforced by the user interface. Beyond this constraint, we do not make any further assumption on the scores.

Now, note that due to order constraints, the scores of *Watches* and *Diving Gear*, while unknown, cannot be lower than that of *Diving Watches*; so either of the two could replace *Diving Watches* in the top-2 answer. To choose between these categories, we observe that the score of *Diving Gear* must be exactly 0.5 (which bounds it from above and below). In contrast, as the score of *Wearable Devices* is 0.9, *Clothing* has a score of at least 0.9, so the score of *Watches* can be anything between 0.5 and an unknown value which is ≥ 0.9 . A better top-2 answer is thus *Wearable Devices* and *Watches*, the latter being likely to have a higher score than *Diving Watches* (or *Diving Gear*).

Other application domains. Beyond the crowdsourcing example that we illustrate, our setting is relevant to many other application domains involving different types of unknown data. For instance, in the domain of Web services [45], one may wish to find the k -most

¹ In the latter case, we aggregate the answers of multiple workers to obtain the compatibility score.

relevant apartments for a user, given user-provided criteria on the apartment (price range, location, size, etc.) but no precise way to aggregate them into a relevance function. In this case, order constraints may be imposed on the unknown relevance of apartments, whenever one apartment *dominates* another (i.e., is better for each criterion); exact-value constraints may capture user-provided ratings to viewed apartments; and a top- k query could be used, e.g., to select the most relevant apartments among those available for sale.

As another example, in the context of top- k queries over sensor data [24, 34], one may wish to find the k -fastest drivers in a certain region given partial data from speedometers and street cameras; comparing the progress and locations of vehicles may yield partial order constraints on their speed. Other domains include, e.g., data mining [3], managing preference data [47], or finding optimal values of black-box functions expressing cost, running time, etc.

Contributions. As previously mentioned, we assume a uniform probability distribution over valuations of unknown items, which we capture formally in our model via possible-world semantics. We then use the *expected values* of items as an estimate of their unknown values, for the sake of top- k computation (see Section 2). Our work presents three main contributions using this general model, as follows.

First, in Section 3 we present a *general and principled scheme to interpolate unknown values from known ones under partial order constraints*, and thereby obtain the top- k such values. We implement this in an algorithm that is polynomial in the number of possible item orderings, and consequently show that in the worst case it is in $\text{FP}^{\#\text{P}}$ in the size of the input.² The problem of finding expected values has a geometric characterization as centroid computation in high-dimensional polytopes (as we explain further); however, our $\text{FP}^{\#\text{P}}$ membership result goes beyond existing computational geometry results since the constraints that we consider, namely, partial order and exact-value constraints, correspond to special classes of polytopes not studied in the context of geometry. Indeed, centroid computation is generally not in $\text{FP}^{\#\text{P}}$ [31, 41]. Our work also departs from previous work on top- k queries over incomplete or probabilistic data [14, 27, 45]: we do not make the simplifying assumption that item distributions are independent and given, but rather study the effect of constraints on individual item distributions.

Our second main contribution, in Section 4, is to establish *hardness results*, and specifically, a matching lower $\text{FP}^{\#\text{P}}$ bound for top- k computation. While the $\#\text{P}$ -hardness of computing expected values follows from the geometric characterization of the problem [41], we show that top- k is hard even without computing expected values. Hence the $\text{FP}^{\#\text{P}}$ bound is tight for both interpolation and top- k ; this shows that the assumption regarding variable independence in previous work, that enables PTIME solutions [14, 27, 45], indeed simplifies the problem. To complete the picture we discuss possible approximation schemes for the interpolation and top- k problems, again by the connection to centroid computation.

Our third main contribution, in Section 5, is the study of *tractable cases*, following the hardness of the general problem and the high complexity of approximation. We devise a PTIME algorithm to compute expected values and top- k results when the order constraints are *tree-shaped* or decomposable to trees. This class of constraints is commonly encountered in the context of unknown values (e.g., taxonomies of products are often trees rather than DAGs); yet, to our knowledge, the corresponding polytopes have no equivalents in the context of computational geometry.

² $\#\text{P}$ is the class of counting problems that return the number of solutions of NP problems. $\text{FP}^{\#\text{P}}$ is the class of function problems that can be computed in PTIME using a $\#\text{P}$ oracle.

Our results also include a review of existing definitions for top- k over uncertain data, which motivates our particular choice of definition (in Section 6). We survey related work in more depth in Section 7 and conclude in Section 8. Our results are provided with complete proofs, given in appendix for lack of space.

2 Preliminaries and Problem Statement

This section introduces the formal definitions for the problem that we study in this paper. We model known and unknown item values as *variables*, and order constraints as *equalities* and *inequalities* over them. Then we define the possible valuations for the variables via possible-world semantics, and use this semantics to define a uniform distribution where all worlds are equally likely. The problem of top- k querying over unknown values can then be formally defined with respect to the expected values of variables in the resulting distribution.

2.1 Unknown Data Values under Constraints

Our input includes a set $\mathcal{X} = \{x_1, \dots, x_n\}$ of variables with unknown values $v(x_1), \dots, v(x_n)$, which we assume³ to be in the range $[0, 1]$. We consider two kinds of constraints over them:

- **order constraints**, written $x_i \leq x_j$ for $x_i, x_j \in \mathcal{X}$, encoding that $v(x_i) \leq v(x_j)$;
- **exact-value constraints** to represent variables with known values, written⁴ $x_i = \alpha$ for $0 \leq \alpha \leq 1$ and for $x_i \in \mathcal{X}$, encoding that $v(x_i) = \alpha$.

In what follows, a *constraint set* with constraints of both types is typically denoted \mathcal{C} . We assume that constraints in \mathcal{C} are *not contradictory* (e.g., we forbid $x = 0.1$, $y = 0.2$, $y \leq z$, and $z \leq x$), and that they are *closed under implication*: e.g., if $x = \alpha$, $y = \beta$ are given, and $\alpha \leq \beta$, then $x \leq y$ is implied and thus should also be in \mathcal{C} . We can check in PTIME that \mathcal{C} is non-contradictory by simply verifying that it does not entail a false inequality on exact values (e.g., $0.2 \leq 0.1$ as in our previous example). The closure of \mathcal{C} can be found in PTIME as a transitive closure computation [26] that also considers exact-value constraints. We denote by $\mathcal{X}_{\text{exact}}$ the subset of \mathcal{X} formed of variables with exact-value constraints.

► **Example 1.** In the product classification example from the Introduction, a variable $x_i \in \mathcal{X}$ would represent the compatibility score of the product to the i -th category. If the score is known, we would encode it as a constraint $x_i = \alpha$. In addition, \mathcal{C} would contain the order constraint $x_i \leq x_j$ whenever category i is a sub-category of j (recall that the score of a sub-category cannot be higher than that of an ancestor category).

2.2 Possible World Semantics

The unknown data captured by \mathcal{X} and \mathcal{C} makes infinitely many valuations of \mathcal{X} possible (including the true one). We model these options via possible world semantics: a *possible world* w for a constraint set \mathcal{C} over $\mathcal{X} = \{x_1, \dots, x_n\}$ is a vector of values $w = (v_1, \dots, v_n) \in [0, 1]^n$, corresponding to setting $v(x_i) := v_i$ for all i , such that all the constraints of \mathcal{C} hold under this valuation. The set of all possible worlds is denoted by $\text{pw}_{\mathcal{X}}(\mathcal{C})$, or by $\text{pw}(\mathcal{C})$ when \mathcal{X} is clear from context.

³ Our results extend to other bounded, continuous ranges, because we can rescale them to fall in $[0, 1]$.

⁴ The number α is written as a rational number, represented by its numerator and denominator.

Notice that \mathcal{C} can be encoded as a set of *linear constraints*, i.e., a set of inequalities between linear expressions on \mathcal{X} and constants in $[0, 1]$. Thus, following common practice in linear programming, the feasible region of a set of linear constraints ($\text{pw}(\mathcal{C})$ in our setting) can be characterized geometrically as a convex polytope, termed the *admissible polytope*: writing $n := |\mathcal{X}|$, each linear constraint defines a feasible half-space of \mathbb{R}^n (e.g., the half-space where $x \leq y$), and the convex polytope $\text{pw}(\mathcal{C})$ is the intersection of all half-spaces. In our setting the polytope $\text{pw}(\mathcal{C})$ is bounded within $[0, 1]^n$, and it is non-empty by our assumption that \mathcal{C} is not contradictory. With exact-value constraints, or order constraints such as $x_i \leq x_j$ and $x_j \leq x_i$, it may be the case that the dimension of this admissible polytope is less than $|\mathcal{X}|$. Computing this dimension can easily be done in PTIME (see, e.g., [43]).

► **Example 2.** Let $\mathcal{X} = \{x, y, z\}$. If $\mathcal{C} = \{x \leq y\}$, the admissible polytope has dimension 3 and is bounded by the planes defined by $x = y$, $x = 0$, $y = 1$, $z = 0$ and $z = 1$. If we add to \mathcal{C} the constraint $y = 0.3$, the admissible polytope is a 2-dimensional rectangle bounded by $0 \leq x \leq 0.3$ and $0 \leq z \leq 1$ on the $y = 0.3$ plane. We cannot add, for example, the constraint $x = 0.5$, because \mathcal{C} would become contradictory.

2.3 Probability Distribution

Having characterized the possible worlds of $\text{pw}(\mathcal{C})$, we assume a *uniform* probability distribution over $\text{pw}(\mathcal{C})$, as indicated in the Introduction. This captures the case when all possible worlds are equally likely, and is a natural choice when we have no information about which valuations are more probable.

Since the space of possible worlds is continuous, we formally define this distribution via a probability density function (pdf), as follows. Let \mathcal{X} and \mathcal{C} define a d -dimensional polytope $\text{pw}_{\mathcal{X}}(\mathcal{C})$ for some integer d . The *d -volume* (also called the Lebesgue measure [28] on \mathbb{R}^d) is a measure for continuous subsets of d -dimensional space, which coincides with length, area, and volume for dimensions 1, 2, and 3, respectively. We denote by $V_d(\mathcal{C})$ the d -volume of the admissible polytope, or simply $V(\mathcal{C})$ when d is the dimension of $\text{pw}(\mathcal{C})$.

► **Definition 3.** The *uniform pdf* p maps each possible world $w \in \text{pw}(\mathcal{C})$ to the constant $p(w) := 1/V(\mathcal{C})$.

2.4 Top-k Queries

We are now ready to formally define the main problem studied in this paper, namely, the evaluation of *top-k queries* over unknown data values. The queries that we consider retrieve the k items that are estimated to have the highest values, *along with their estimated values*, with ties broken arbitrarily. We further allow queries to apply a *selection operator* σ on the items before performing the top- k computation. In our example from the Introduction, this is what allows us to select the top- k categories among only the end categories. We denote the subset of \mathcal{X} selected by σ as \mathcal{X}_{σ} .

If all item values are known, the semantics of top- k queries is clear. In presence of unknown values, however, the semantics must be redefined to determine how the top- k items and their values are estimated. In this paper, we estimate unknown items by their *expected value over all possible worlds*, i.e., their expected value according to the uniform pdf p defined above on $\text{pw}(\mathcal{C})$. This corresponds to *interpolating* the unknown values from the known ones, and then querying the result. We use these interpolated values to define the top- k problem as computing the k variables with the highest expected values, but we also study on its own the interpolation problem of computing the expected values.

To summarize, the two formal problems that we study on constraint sets are:

Interpolation. Given a constraint set \mathcal{C} over \mathcal{X} and variable $x \in \mathcal{X}$, the *interpolation problem* for x is to compute the expected value of x in the uniform distribution over $\text{pw}_{\mathcal{X}}(\mathcal{C})$.

Top-k. Given a constraint set \mathcal{C} over \mathcal{X} , a selection predicate σ , and an integer k , the *top-k computation problem* is to compute the ordered list of the k maximal expected values of variables in \mathcal{X}_{σ} (or less if $|\mathcal{X}_{\sigma}| \leq k$), with ties broken arbitrarily.

We review other definitions of top-k on uncertain data in Section 6, where we justify our choice of semantics.

Alternate phrasing. The Interpolation problem can also be defined geometrically, as the computation of the *centroid* (or *center of mass*) of the admissible polytope: the point G such that all vectors relative to G originating at points within the polytope sum to zero. The constraints that we study correspond to a special kind of polytopes, for which we will design a specific algorithm in the next section, and derive an $\text{FP}^{\#P}$ membership bound which does not hold for general polytopes (as explained in the Introduction). However, the geometric connection will become useful when we study the complexity of our problem in Section 4.1.

3 An Algorithm for Interpolation and Top-k

Having defined formally the problems that we study, we begin our complexity analysis by designing an algorithm that computes the expected value of variables.

The algorithm enumerates all possible orderings of the variables (to be defined formally below), but it is still nontrivial: we must handle exact-value constraints specifically, and we must compute the probability of each ordering to determine its weight in the overall expected value computation. From the algorithm, we will deduce that our interpolation and top-k problems are in $\text{FP}^{\#P}$.

Eliminating ties. To simplify our study, we will eliminate from the start the problem of *ties*, which will allow us to assume that values in all worlds are totally ordered. We say that a possible world $w = (v_1, \dots, v_n)$ of \mathcal{C} has a *tie* if $v_i = v_j$ for some i, j . Note that occasional ties, not enforced by \mathcal{C} , have an *overall probability of 0*: intuitively, if the admissible polytope is d -dimensional, then all the worlds where $v_i = v_j$ correspond to a $(d - 1)$ -dimensional hyperplane bounding or intersecting the polytope. A finite set of such hyperplanes (for every pair of variables) has total d -volume 0. Since our computations (volume, expected value) involve integrating over possible worlds, a set of worlds with total probability 0 does not affect the result.

What is left is to consider ties enforced by \mathcal{C} (and thus having probability 1). In such situations, we can rewrite \mathcal{C} by merging these variables to obtain an equivalent constraint set where ties have probability 0. Formally:

► **Lemma 4.** *For any constraint set \mathcal{C} , we can construct in PTIME a constraint set \mathcal{C}' such that the probability that the possible worlds of \mathcal{C}' have a tie (under the uniform distribution) is zero, and such that any interpolation or top-k computation problem on \mathcal{C} can be reduced in PTIME to the same type of problem on \mathcal{C}' .*

Hence, we assume from now on that ties have zero probability in \mathcal{C} , so that we can ignore possible worlds with ties without affecting the correctness of our analysis. Note that this

implies that all of our results also hold for *strict inequality constraints*, of the forms $x < y$ and $x \neq y$.

Under this assumption, we first study in Section 3.1 the case where \mathcal{C} is a total order. We then handle arbitrary \mathcal{C} by aggregating over possible variable orderings, in Section 3.2.

3.1 Total Orders

In this section we assume \mathcal{C} is a total order $\mathcal{C}_1^n(\alpha, \beta)$ defined as $x_0 \leq x_1 \leq \dots \leq x_n \leq x_{n+1}$, where $x_0 = \alpha$ and $x_{n+1} = \beta$ are variables with exact-value constraints in $\mathcal{X}_{\text{exact}}$.

We first consider *unfragmented* total orders, where $x_1, \dots, x_n \notin \mathcal{X}_{\text{exact}}$. In this case, we can show that the expected value of x_i , for $1 \leq i \leq n$, corresponds to a *linear interpolation* of the unknown variables between α and β , namely: $\frac{i}{n+1} \cdot (\beta - \alpha) + \alpha$. This can be shown formally via a connection to the expected value of the order statistics of samples from a uniform distribution, which follows a Beta distribution [21].

Now consider the case of *fragmented* total orders, where \mathcal{C} is allowed to contain more exact-value constraints than the ones on α and β . We observe that we can *split* the total order into *fragments*: by cutting at each variable that has an exact-value constraint, we obtain sub-sequences of variables which follow an unfragmented total order. We can then compute the expected values of each fragment independently, and compute the total order volume as the product of the fragment volumes. The correctness of this computation follows from a more general result (Lemma 12) stated and proven in Section 5.

Hence, given a constraint set \mathcal{C} imposing a (possibly fragmented) total order, the expected value of x_i can be computed as follows. If $x_i \in \mathcal{X}_{\text{exact}}$, analysis is trivial. Otherwise, we consider the fragment $\mathcal{C}_{p+1}^{q-1}(v_p, v_q)$ that contains x_i ; namely, p is the maximal index such that $0 \leq p < i$ and $x_p \in \mathcal{X}_{\text{exact}}$, and q is the minimal index such that $i < q \leq n+1$ and $x_q \in \mathcal{X}_{\text{exact}}$. The expected value of x_i can then be computed within $\mathcal{C}_{p+1}^{q-1}(v_p, v_q)$ using linear interpolation.

The following proposition summarizes our findings:

► **Proposition 5.** *Given a constraint set \mathcal{C} implying a total order, the expected value of any variable $x_i \in \mathcal{X}$ can be computed in PTIME.*

3.2 General Constraint Sets

We can now extend the result for total orders to an expression of the expected value for a general constraint set \mathcal{C} . We apply the previous process to each possible total ordering of the variables, and aggregate the results. To do this, we define the notion of *linear extensions*, inspired by partial order theory:

► **Definition 6.** Given a constraint set \mathcal{C} over \mathcal{X} , we say that a constraint set \mathcal{T} is a *linear extension* of \mathcal{C} if (i) \mathcal{T} is a total order; (ii) the exact-value constraints of \mathcal{T} are exactly those of \mathcal{C} ; and (iii) $\mathcal{C} \subseteq \mathcal{T}$, namely every constraint $x \leq y$ in \mathcal{C} also holds⁵ in \mathcal{T} .

Algorithm 1 presents our general scheme to compute the expected value of a variable $x \in \mathcal{X}$ under an arbitrary constraint set \mathcal{C} , assuming the uniform distribution on $\text{pw}(\mathcal{C})$.

The algorithm iterates over each linear extension \mathcal{T} of \mathcal{C} , and computes the expected value of x in \mathcal{T} and the overall probability of \mathcal{T} in $\text{pw}(\mathcal{C})$. A linear extension is a total order,

⁵ The linear extensions of \mathcal{C} in this sense are thus exactly the linear extensions of the partial order on \mathcal{X} imposed by \mathcal{C} : this partial order is indeed antisymmetric because \mathcal{C} has no ties.

Algorithm 1: Compute the expected value of a variable

Input: Constraint set \mathcal{C} on variables \mathcal{X} with $n := |\mathcal{X}|$ where ties have probability 0 and where value range $[0, 1]$ is enforced by constraints; variable $x \in \mathcal{X}$

Output: Expected value of x

- 1 **if** x is in $\mathcal{X}_{\text{exact}}$, i.e., has an exact-value constraint in \mathcal{C} to some v **then return** v ;
- 2 $\mathbb{E}_{\mathcal{C}}[x] \leftarrow 0$; $V(\mathcal{C}) \leftarrow 0$;
- 3 $m \leftarrow |\mathcal{X}_{\text{exact}}|$;
- 4 Write $x_0 < \dots < x_{m-1}$ the variables of $\mathcal{X}_{\text{exact}}$ and $v_0 < \dots < v_{m-1}$ their values;
- 5 **foreach** linear extension \mathcal{T} of \mathcal{C} **do**
- 6 Write $i_0 < \dots < i_{m-1}$ the indices in \mathcal{T} of variables x_0, \dots, x_{m-1} in $\mathcal{X}_{\text{exact}}$;
- 7 Write k the index in \mathcal{T} of the variable x ;
- 8 Write i_j, i_{j+1} the indices of variables from $\mathcal{X}_{\text{exact}}$ s.t. $i_j < k < i_{j+1}$;
- 9 $\mathbb{E}_{\mathcal{T}}[x] \leftarrow \text{ExpectedValFrag}(i_j, i_{j+1}, k, v_j, v_{j+1})$; // Expected value of x in \mathcal{T}
- 10 $V(\mathcal{T}) \leftarrow \prod_{l=0}^{m-2} \text{VolumeFrag}(i_l, i_{l+1}, v_l, v_{l+1})$; // Volume of \mathcal{T}
- 11 $\mathbb{E}_{\mathcal{C}}[x] \leftarrow \mathbb{E}_{\mathcal{C}}[x] + V(\mathcal{T}) \times \mathbb{E}_{\mathcal{T}}[x]$; // Sum exp. value of x weighted by $V(\mathcal{T})$
- 12 $V(\mathcal{C}) \leftarrow V(\mathcal{C}) + V(\mathcal{T})$; // Sum of total order volumes
- 13 **return** $\frac{\mathbb{E}_{\mathcal{C}}[x]}{V(\mathcal{C})}$;
- 14 **Function** $\text{ExpectedValFrag}(p, q, k, \alpha, \beta)$
- 15 **Input:** $p < q$: indices; k : requested variable index; $\alpha < \beta$: exact values at p, q resp.
- 16 **Output:** Expected value of variable x_k in the fragment $\mathcal{C}_{p+1}^{q-1}(\alpha, \beta)$
- 17 $n \leftarrow q - p - 1$; // num. of variables in the fragment which are $\notin \mathcal{X}_{\text{exact}}$
- 18 **return** $\frac{k-p}{n+1} \cdot (\beta - \alpha) + \alpha$; // linear interpolation (see Section 3.1)
- 19 **Function** $\text{VolumeFrag}(p, q, \alpha, \beta)$
- 20 **Input:** $p < q$: indices; $\alpha < \beta$: exact values at p, q resp.
- 21 **Output:** $V(\mathcal{C}_{p+1}^{q-1}(\alpha, \beta))$: volume of the total order fragment between indices p, q
- 22 $n \leftarrow q - p - 1$; // num. of variables in the fragment which are $\notin \mathcal{X}_{\text{exact}}$
- 23 **return** $\frac{(\beta - \alpha)^n}{n!}$; // Volume of $[\alpha, \beta]^n$ divided by num. of total orders

so x is within a particular fragment of it, namely, between the indices of two consecutive variables with exact-value constraints, i_j and i_{j+1} . The *expected value of x in \mathcal{T}* , denoted by $\mathbb{E}_{\mathcal{T}}[x]$, is then affected only by the constraints and variables of this fragment, and can be computed using linear interpolation by the function ExpectedValFrag (line 9).

Now, the final expected value of x in \mathcal{C} is the average of all $\mathbb{E}_{\mathcal{T}}[x]$ weighted by the probability of each linear extension \mathcal{T} , i.e., the volume of $\text{pw}(\mathcal{T})$ divided by the volume of $\text{pw}(\mathcal{C})$. Recall that, by Lemma 4, worlds with ties have total volume 0 and do not affect this expected value. We compute the volume of \mathcal{T} as the *product of volumes of its fragments* (line 10). The volume of a fragment, computed by function VolumeFrag , is the volume of $[\alpha, \beta]^n$, i.e., all assignments to the n variables of the fragment in $[\alpha, \beta]$, divided by the number of orderings of these variables, to obtain the volume of one specific order (line 19).

The complexity of Algorithm 1 is polynomial in the number of linear extensions of \mathcal{C} , as we can enumerate them in constant amortized time [40]. However, in the general case, there may be up to $|\mathcal{X}|!$ linear extensions. To obtain an upper bound in the general case, we note that we can rescale all constraints so that all numbers are integers, and then nondeterministically sum over the linear extensions. This yields our $\text{FP}^{\#P}$ upper bound:


► **Theorem 7.** *Given a constraint set \mathcal{C} over \mathcal{X} and $x \in \mathcal{X}$ (resp., and a selection predicate σ ,*

and an integer k), determining the expected value of x in $\text{pw}(\mathcal{C})$ under the uniform distribution (resp., the top- k computation problem over \mathcal{X} , \mathcal{C} , and σ) is in $\text{FP}^{\#P}$.

The $\text{FP}^{\#P}$ membership for interpolation does not extend to centroid computation in general convex polytopes, which is not in $\text{FP}^{\#P}$ [31,41]. Our algorithm thus relies on the fact that the polytope $\text{pw}(\mathcal{C})$ is of a specific form, defined with order and exact-value constraints. The same upper bound for the top- k problem immediately follows. We will show in Section 4 that this $\text{FP}^{\#P}$ upper bound is tight.

We also provide a complete example to illustrate the constructions of this section.

Full Example. We exemplify our scheme on variables $\mathcal{X} = \{x, y, y', z\}$ and on the constraint set \mathcal{C} generated by the order constraints $x \leq y$, $y \leq z$, $x \leq y'$, $y' \leq z$ and the exact-value constraint $y' = \gamma$ for some fixed $0 < \gamma < 1$. Remember that we necessarily have $0 \leq x$ and $z \leq 1$ as well. The constraints of \mathcal{C} are closed under implication, so they also include $x \leq z$. The figure shows the Hasse diagram of the partial order defined by \mathcal{C} on \mathcal{X} . Note that ties have a probability of zero in $\text{pw}(\mathcal{C})$.



The two linear extensions of \mathcal{C} are $\mathcal{T}_1 : x \leq y \leq y' \leq z$ and $\mathcal{T}_2 : x \leq y' \leq y \leq z$. Now, \mathcal{T}_1 is a fragmented total order, and we have $\text{pw}(\mathcal{T}_1) = \text{pw}_{\{x,y\}}(\mathcal{C}') \times \{\gamma\} \times [\gamma, 1]$ where \mathcal{C}' is defined on variables $\{x, y\}$ by $0 \leq x \leq y \leq \gamma$. We can compute the volume of $\text{pw}(\mathcal{T}_1)$ as $\alpha_1 = \frac{\gamma^2}{2} \times (1 - \gamma)$. Similarly the volume of $\text{pw}(\mathcal{T}_2)$ is $\alpha_2 = \gamma \times \frac{(1-\gamma)^2}{2}$.

Let us compute the expected value of y for \mathcal{C} . In \mathcal{T}_1 its expected value is $\mu_1 = \mathbb{E}_{\mathcal{T}_1}[y] = \frac{2}{3} \cdot (\gamma - 0) + 0 = \frac{2}{3}\gamma$. In \mathcal{T}_2 its expected value is $\mu_2 = \mathbb{E}_{\mathcal{T}_2}[y] = \frac{1}{3} \cdot (1 - \gamma) + \gamma = \frac{1+2\gamma}{3}$. The overall expected value of y is the average of these expected values weighted by total order probabilities (volumes fractions), namely $\mathbb{E}_{\mathcal{C}}[y] = \frac{\alpha_1 \mu_1 + \alpha_2 \mu_2}{\alpha_1 + \alpha_2}$.

4 Hardness and Approximations

We next show that the intractability of Algorithm 1 in Section 3 is probably unavoidable. We first show matching lower bounds for interpolation and top- k in Section 4.1. We then turn in Section 4.2 to the problem of approximating expected values.

4.1 Hardness of Exact Computation

We now analyze the complexity of computing an exact solution to our two main problems. We show below a new result for the hardness of top- k . But first, we state the lower bound for the interpolation problem, which is obtained via the geometric characterization of the problem. In previous work, centroid computation is proven to be hard for *order polytopes*, namely, polytopes without exact-value constraints, which are a particular case of our setting:

► **Theorem 8. ([41], Theorem 1).** *Given a set \mathcal{C} of order constraints and $x \in \mathcal{X}$, determining the expected value of x in $\text{pw}(\mathcal{C})$ under the uniform distribution is $\text{FP}^{\#P}$ -hard.*

We now show a new lower bound for top- k queries: interestingly, these queries are $\text{FP}^{\#P}$ -hard *even if they do not have to return the expected values*. Recall that σ is the selection operator (see Section 2.4), which we use to compute top- k among a restricted subset of variables. We can show hardness even for top-1 queries, and even when σ only selects two variables:

► **Theorem 9.** *Given a constraint set \mathcal{C} over \mathcal{X} , a selection predicate σ , and an integer k , the top- k computation problem over \mathcal{X} , \mathcal{C} and σ is $\text{FP}^{\#\text{P}}$ -hard even if k is fixed to be 1, $|\mathcal{X}_\sigma|$ is 2, and the top- k answer does not include the expected value of the variables.*

Proof sketch. To prove hardness in this case, we reduce from interpolation. We show that a top-1 computation oracle can be used as a comparison oracle to compare the expected value of a variable x to any other rational value α , by adding a fresh element x' with an exact-value constraint to α and using σ to compute the top-1 among $\{x, x'\}$. What is more technical is to show that, given such a comparison oracle, we can perform the reduction and determine *exactly* the expected value v of x (a rational number) using only a *polynomial* number of comparisons to other rationals. This follows from a bound on the denominator of v , and by applying the rational number identification scheme of [38]. See the Appendix. ◀

In settings where we do not have a selection operator (i.e., $\mathcal{X}_\sigma = \mathcal{X}$), we can similarly show the hardness of top- k (rather than top-1). See Appendix B.1 for details.

4.2 Complexity of Approximate Computation

In light of the previous hardness results, we now review approximation algorithms, again via the geometric characterization of our setting. In Section 5, we will show a novel exact solution in PTIME for specific cases.

The *interpolation* problem can be shown to admit a fully polynomial-time randomized approximation scheme (FPRAS). This result follows from existing work [6, 30], using a tractable almost uniform sampling scheme for convex bodies.

► **Proposition 10. ([30], Algorithm 5.8).** *Let \mathcal{C} be a set of constraints with variable set \mathcal{X} and $x \in \mathcal{X}$. There is an FPRAS that determines an estimate $\hat{\mathbb{E}}_{\mathcal{C}}[x]$ of the expected value $\mathbb{E}_{\mathcal{C}}[x]$ of x in $\text{pw}(\mathcal{C})$ under the uniform distribution.*

This result is mostly of theoretical interest, as the polynomial is in $|\mathcal{X}|^7$ (see [6], Table 1), but recent improved sampling algorithms [36] may ultimately yield a practical approximate interpolation technique for general constraint sets (see [20, 35]).

For completeness, we mention two natural ways to define randomized approximations for *top- k computation*:

- We can define the *approximate top- k* as an ordered list of k items whose expected value does not differ by more than some $\varepsilon > 0$ from that of the item in the actual top- k at the same rank. An FPRAS for this definition of approximate top- k can be obtained from that of Proposition 10.
- It is highly unlikely that there exists a PTIME algorithm to return the *actual top- k* with high probability, even without requiring it to return the expected values. Indeed, such an algorithm would be in the BPP (bounded-error probabilistic time) complexity class; yet it follows from Theorem 9 above that deciding whether a set of variables is the top- k is NP-hard, so the existence of the algorithm would entail that $\text{NP} \subseteq \text{BPP}$.

5 Tractable Cases

Given the hardness results in the previous section and the impracticality of approximation, we now study whether *exact* interpolation and top- k computation can be tractable on restricted classes of constraint sets. We consider tree-shaped constraints (defined formally below) and generalizations thereof: they are relevant for practical applications (e.g., classifying items

into tree- or forest-shaped taxonomies), and we will show that our problems are tractable on them. We start by a splitting lemma to decompose constraint sets into “independent” subsets of variables, and then define and study our tractable class.

5.1 Splitting Lemma

We will formalize the cases in which the valuations of two variables in \mathcal{X} are probabilistically dependent (the variables *influence* each other), according to \mathcal{C} . This, in turn, will enable us to define *independent subsets of the variables* and thus *independent subsets of the constraints* over these variables. This abstract result will generalize the notion of fragments from total orders (see Section 3.1) to general constraint sets. In what follows, we use $x_i \prec x_j$ to denote the *covering relation* of the partial order \leq , i.e., $x_i \leq x_j$ is in \mathcal{C} but there exists no $x_k \notin \{x_i, x_j\}$ such that $x_i \leq x_k$ and $x_k \leq x_j$ are in \mathcal{C} .

► **Definition 11.** We define the *influence relation* $x \leftrightarrow y$ between variables of $\mathcal{X} \setminus \mathcal{X}_{\text{exact}}$ as the equivalence relation obtained by the symmetric, reflexive, and transitive closure of the \prec relation on $\mathcal{X} \setminus \mathcal{X}_{\text{exact}}$.

The *uninfluenced classes* of \mathcal{X} under \mathcal{C} is the partition of $\mathcal{X} \setminus \mathcal{X}_{\text{exact}}$ as the subsets X_1, \dots, X_m given by the equivalence classes of the influence relation.

The *uninfluence decomposition* of \mathcal{C} is the collection of constraint sets $\mathcal{C}_1, \dots, \mathcal{C}_m$ of \mathcal{C} where each \mathcal{C}_i has as variables $X_i \sqcup \mathcal{X}_{\text{exact}}$ and contains all exact-value constraints of \mathcal{C} and all order constraints between variables of $X_i \sqcup \mathcal{X}_{\text{exact}}$.

We assume w.l.o.g. that $m > 0$, i.e., there are unknown variables in $\mathcal{X} \setminus \mathcal{X}_{\text{exact}}$; otherwise the uninfluence decomposition is meaningless but any analysis is trivial. Intuitively, two unknown variables x, x' are in different uninfluenced classes if in *every* linear extension there is *some* variable from $\mathcal{X}_{\text{exact}}$ between them, or if they belong to disconnected (and thus incomparable) parts of the partial order. In particular, uninfluenced classes correspond to the fragments of a total order: this is used in Section 3.1. The uninfluence decomposition captures only constraints between *variables that influence each other*, and constraints that can *bound the range of a variable* by making it comparable to variables from $\mathcal{X}_{\text{exact}}$. We formally prove the independence of $\mathcal{C}_1, \dots, \mathcal{C}_m$ via possible-world semantics: every possible world of \mathcal{C} can be decomposed to possible worlds of $\mathcal{C}_1, \dots, \mathcal{C}_m$, and vice versa.

► **Lemma 12.** Let $\mathcal{C}_1, \dots, \mathcal{C}_m$ be the uninfluence decomposition of \mathcal{C} . There exists a bijective correspondence between $\text{pw}(\mathcal{C})$ and $\text{pw}(\mathcal{C}_1) \times \dots \times \text{pw}(\mathcal{C}_m)$.

► **Example 13.** Let \mathcal{X} be $\{x, y, y', z, w\}$, and let \mathcal{C} be defined by $y' = 0.5$ and $x \leq y \leq y' \leq z$. The uninfluence classes are $X_1 = \{x, y\}$, $X_2 = \{z\}$, and $X_3 = \{w\}$. The uninfluence decomposition thus consists of \mathcal{C}_1 , with variables $X_1 \sqcup \{y'\}$, and constraints $x \leq y \leq y'$ and $y' = 0.5$; \mathcal{C}_2 , with variables $X_2 \sqcup \{y'\}$, and constraints $y' \leq z$ and $y' = 0.5$; and \mathcal{C}_3 , with variables $X_3 \sqcup \{y'\}$, and constraint $y' = 0.5$.

We next use this independence property to analyse restricted classes of constraint sets.

5.2 Tree-Shaped Constraints

We define the first restricted class of constraints that we consider: *tree-shaped* constraints. Recall that a *Hasse diagram* is a representation of a partial order as a directed acyclic graph, whose nodes correspond to \mathcal{X} and where there is an edge (x, y) if $x \prec y$. An example of such a diagram is the one used in Section 3.2.

► **Definition 14.** A constraint set \mathcal{C} over \mathcal{X} is *tree-shaped* if the probability of ties is zero, the Hasse diagram of the partial order induced on \mathcal{X} by \mathcal{C} is a directed tree, the root has exactly one child, and exactly the root and leaves are in $\mathcal{X}_{\text{exact}}$. Thus, \mathcal{C} imposes a global minimal value, and maximal values at each leaf, and no other exact-value constraint.

We call \mathcal{C} *reverse-tree-shaped* if the reverse of the Hasse diagram (obtained by reversing the direction of the edges) is tree-shaped.

Tree-shaped constraints are often encountered in practice, in particular in the context of product taxonomies. Indeed, while our example from Figure 1 is a DAG, many real-life taxonomies are trees: in particular, the Google Product Taxonomy [22] and ACM CCS [2].

We now show that for a tree-shaped constraint set \mathcal{C} , unlike the general case, we can tractably compute exact expressions of the expected values of variables. In the next two results, we assume arithmetic operations on rationals to have unit cost, e.g., they are performed up to a fixed numerical precision. Otherwise, the complexities remain polynomial but the degrees may be larger. We first show:

► **Theorem 15.** *For any tree-shaped constraint set \mathcal{C} over \mathcal{X} , we can compute its volume $V(\mathcal{C})$ in time $O(|\mathcal{X}|^2)$.*

Proof sketch. We process the tree bottom-up, propagating a piecewise-polynomial function expressing the volume of the subpolytope on the subtree rooted at each node as a function of the value of the parent node: we compute it using Lemma 12 from the child nodes. ◀

See Appendix C.2 for the complete proof. This result can be applied to prove the tractability of computing the marginal distribution of any variable $x \in \mathcal{X} \setminus \mathcal{X}_{\text{exact}}$ in a tree-shaped constraint set, which is defined as the pdf $p_x(v) := V_{d-1}(\mathcal{C} \cup \{x = v\}) / V_d(\mathcal{C})$, where d is the dimension of $\text{pw}(\mathcal{C})$:

► **Theorem 16.** *For any tree-shaped constraint set \mathcal{C} on variable set \mathcal{X} , for any variable $x \in \mathcal{X} \setminus \mathcal{X}_{\text{exact}}$, the marginal distribution for x is piecewise polynomial and can be computed in time $O(|\mathcal{X}_{\text{exact}}| \times |\mathcal{X}|^2)$.*

Proof sketch. We proceed similarly to the proof of Theorem 15 but with two functions: one for x and its descendants, and one for all other nodes. The additional $|\mathcal{X}_{\text{exact}}|$ factor is because the second function depends on how the value given to x compares to the leaves. ◀

We last deduce that our results for tree-shaped constraints extend to a more general tractable case: constraint sets \mathcal{C} whose uninfluence decomposition $\mathcal{C}_1, \dots, \mathcal{C}_m$ is such that every \mathcal{C}_i is (reverse-)tree-shaped. By Lemma 12, each \mathcal{C}_i (and its variables) can be considered independently, and reverse-tree-shaped trees can be easily transformed into tree-shaped ones. Our previous algorithms thus apply to this general case, by executing them on each constraint set of the uninfluence decomposition that is relevant to the task (namely, containing the variable x to interpolate, or top- k candidates from the selected variables \mathcal{X}_σ):

► **Corollary 17.** *Given any constraint set \mathcal{C} and its uninfluence decomposition $\mathcal{C}_1, \dots, \mathcal{C}_m$, assuming that each \mathcal{C}_i is a (reverse-)tree-shaped constraint set, we can solve the interpolation problem in time $O(\max_i |\mathcal{X}_i|^3)$ and the top- k problem in PTIME.*

On large tree-shaped taxonomies (e.g., the Google Product Taxonomy [22]), in an interactive setting where we may ask user queries (e.g., the one in the Introduction), we can improve running times by asking more queries. Indeed, each answer about a category

adds an exact-value constraint, and reduces the size of the constraint sets of the uninfluence decomposition, which decreases the overall running time, thanks to the superadditivity of $x \mapsto x^3$. We do not study which variables should be queried in order to reduce the running time of the algorithm; see, e.g., [39] for tree-partitioning algorithms.

6 Other Variants

We have defined top- k computation on constraint sets by considering the *expected value* of each variable under the uniform distribution. Comparing to different definitions of top- k on unknown values that have been studied in previous work, our definition has some important properties [14]: it provides a ready estimation for unknown values (namely, their expected value) and guarantees an output of size k . Moreover, it satisfies the *containment property* of [14], defined in our setting as follows:

► **Definition 18.** A top- k definition satisfies the *containment property* if for any constraint set \mathcal{C} on variables \mathcal{X} , for any predicate σ (where we write \mathcal{X}_σ the selected variables), and for any $k < |\mathcal{X}_\sigma|$, letting S_k and S_{k+1} be the ordered lists of top- k and top- $(k+1)$ variables, S_k is a strict prefix of S_{k+1} .

The containment property is a natural desideratum: computing the top- k for some $k \in \mathbb{N}$ should not give different variables or order for the top- k' with $k' < k$. Our definition clearly satisfies the containment property (except in the case of ties). By contrast, we will now review prominent definitions of top- k on uncertain data from related work [14, 46, 51], and show that they do not satisfy the containment property when we apply them to the possible world distributions studied in our setting. We focus on two prominent definitions, *U-top- k* and *global-top- k* and call our own definition *local-top- k* when comparing to them; we also discuss other variants in Appendix D.2.

U-top- k . The *U-top- k* variant does not study *individual* variables but defines the output as the *sequence* of k variables most likely to be the top- k (in that order), for the uniform distribution on $\text{pw}(\mathcal{C})$. We call this alternative definition *U-top- k* by analogy with [14, 46]. Interestingly, the U-top- k and local-top- k definitions sometimes disagree in our setting:

► **Lemma 19.** *There is a constraint set \mathcal{C} and selection predicate σ such that local-top- k and U-top- k do not match, even for $k = 1$ and without returning expected values or probabilities.*

We can easily design an algorithm to compute U-top- k in PSPACE and in polynomial time in the number of linear extensions of \mathcal{C} : compute the probability of each linear extension as in Algorithm 1, and then sum on linear extensions depending on which top- k sequence they realize (on the variables selected by σ), to obtain the probability of each answer. Hence:

► **Proposition 20.** *For any constraint set \mathcal{C} over \mathcal{X} , integer k and selection predicate σ , the U-top- k query for \mathcal{C} and σ can be computed in PSPACE and in time $O(\text{poly}(N))$, where N is the number of linear extensions of \mathcal{C} .*

Unlike Theorem 7, however, this does not imply $\text{FP}^{\#P}$ -membership: when selecting the most probable sequence, the number of candidate sequences may not be polynomial (as k is not fixed). We leave to future work an investigation of the precise complexity of U-top- k .

We show that in our setting U-top- k does not satisfy the *containment property* of [14].

► **Lemma 21.** *There is a constraint set \mathcal{C} without ties such that U-top- k does not satisfy the containment property for the uniform distribution on $\text{pw}(\mathcal{C})$.*

Global-top- k . We now study the *global-top- k* definition [51], and show that it does not respect the containment property either, even though it is defined on individual variables:

► **Definition 22.** The *global-top- k query*, for a constraint set \mathcal{C} , selection predicate σ , and integer k , returns the k variables that have the highest probability in the uniform distribution on $\text{pw}(\mathcal{C})$ to be among the variables with the k highest values, sorted by decreasing probability.

► **Lemma 23.** *There is a constraint set \mathcal{C} without ties such that global-top- k does not satisfy the containment property for the uniform distribution on $\text{pw}(\mathcal{C})$.*

7 Related Work

We extend the discussion about related work from the Introduction.

Ranking queries over uncertain databases. A vast body of work has focused on providing semantics and evaluation methods for order queries over uncertain databases, including top- k and ranking queries (e.g., [14, 18, 24, 25, 27, 32, 42, 45, 49, 50]). Such works consider two main uncertainty types: *tuple-level uncertainty*, where the existence of tuples (i.e., variables) is uncertain, and hence affects the query results [14, 18, 25, 27, 32, 42, 49, 50]; and *attribute-level uncertainty*, more relevant to our problem, where the data tuples are known but some of their values are unknown or uncertain [14, 24, 27, 45]. Top- k queries over uncertain data following [45] was recently applied to crowdsourcing applications in [12]. These studies are relevant to our work as they identify multiple possible semantics for order queries in presence of uncertainty, and specify desired properties for such semantics [14, 27]; our definition of top- k satisfies the desiderata that are relevant to attribute-level uncertainty [27].

We depart from this existing work in two main respects. First, existing work assumes that each variable is given with an *independent function* that describes its probability distribution. We do not assume this, and instead *derive* expressions for the expected values of variables in a principled way from a uniform prior on the possible worlds. Our work is thus well-suited to the many situations where probability distributions on variables are not known, or where they are not independent (e.g., when order constraints are imposed on them). For this reason, the problems that we consider are generally computationally harder. For instance, [45] is perhaps the closest to our work, since they consider the total orders compatible with given partial order constraints. However, they assume independent marginal distributions, so they can evaluate top- k queries by only considering k -sized prefixes of the linear extensions; in our setting even computing the top-1 element is hard (Theorem 9).

The second key difference is that other works *do not try to estimate the top- k values*, because they assume that the marginal distribution is given: they only focus on ranks. In our context, we need to compute missing values, and need to account, e.g., for exact-value constraints and their effect on the probability of possible worlds and on expected values (Section 3).

We also mention our previous work [4] which considers the estimation of uncertain values (expectation and variance), but only in a *total order*, and did not consider complexity issues.

Partial order search. Another relevant research topic, *partial order search*, considers queries over elements in a partially ordered set to find a subset of elements with a certain property [3, 16, 19, 23, 39]. This relates to many applications, e.g., crowd-assisted graph search [39], frequent itemset mining with the crowd [3], and knowledge discovery, where the unknown data is queried via oracle calls [23]. These studies are *complementary to ours*: when the

target function can be phrased as a top- k or interpolation problem, if the search is stopped before all values are known, we can use our method to estimate the complete output.

Computational geometry. Our work reformulates the interpolation problem as a centroid computation problem in the polytope of possible worlds defined by the constraint set. This problem has been studied independently by computational geometry work [30, 37, 41].

Computational geometry mostly studies *arbitrary* convex polytopes (corresponding to polytopes defined by arbitrary linear constraint sets), and often considers the task of *volume computation*, which is related to the problem of computing the centroid [41]. In this context, it is known that computing the exact volume of a polytope is not in $\text{FP}^{\#\text{P}}$ because the output is generally not of polynomial size [31]. Nevertheless, several (generally exponential) methods for exact volume computation [10] have been developed. The problem of approximation has also been studied, both theoretically and practically [15, 17, 20, 30, 35, 44]. Our problem of *centroid* computation is studied in [37], whose algorithm is based on the idea of computing the volume of a polytope by computing the lower-dimensional volume of its facets. This is different from our algorithm, which divides the polytope along linear extensions into subpolytopes, for which we apply a specific volume and centroid computation method.

Some works in computational geometry specifically study *order polytopes*, i.e., the polytopes defined by constraint sets with only order constraints and no exact-value constraints. For such polytopes, volume computation is known to be $\text{FP}^{\#\text{P}}$ -complete [9], leading to a $\text{FP}^{\#\text{P}}$ -hardness result for centroid computation [41]. However, these results do not apply to *exact-value constraints*, i.e., when order polytopes can only express order relations, between variables which are in $[0, 1]$. Exact-value constraints are both highly relevant in practice (to represent numerical bounds, or known information, e.g., for crowdsourcing), allow for more general polytopes, and complicate the design of Algorithm 1, which must perform volume computation and interpolation in each *fragmented* linear order.

Furthermore, to our knowledge, computational geometry works do not study the top- k problem, or polytopes that correspond to tree-shaped constraint sets, since these have no clear geometric interpretation.

Tree-shaped partial orders. Our analysis of tractable schemes for tree-shaped partial orders is reminiscent of the well-known tractability of probabilistic inference in tree-shaped graphical models [7], and of the tractability of probabilistic query evaluation on trees [13] and treelike instances [5]. However, we study continuous distributions on numerical values, and the influence between variables when we interpolate does not simply follow the tree structure; so our results do not seem to follow from these settings.

8 Conclusion

In this paper, we have studied the problems of top- k computation and interpolation for data with unknown values and order constraints. We have provided foundational solutions, including a general computation scheme, complexity bounds, and analysis of tractable cases.

One natural direction for future work is to study whether our tractable cases (tree-shaped orders, sampling) can be covered by more efficient PTIME algorithms, or whether more general tractable cases can be identified: for instance, a natural direction to study would be partial orders with a *bounded-treewidth* Hasse diagram, following recent tractability results for the related problem of linear extension counting [29]. Another question is to extend our scheme to request additional values from the crowd, as in [3, 12], and reduce the expected

error on the interpolated values or top- k query, relative to a user goal. In such a setting, how should we choose which values to retrieve, and could we update incrementally the results of interpolation when we receive new exact-value constraints? Finally, it would be interesting to study whether our results generalize to different prior distributions on the polytope.

Acknowledgements

This work is partially supported by the European Research Council under the FP7, ERC grant MoDaS, agreement 291071, by a grant from the Blavatnik Interdisciplinary Cyber Research Center, by the Israel Science Foundation (grant No. 1157/16), and by the Télécom ParisTech Research Chair on Big Data and Market Insights.

References

- 1 S. Abiteboul, T.-H. H. Chan, E. Kharlamov, W. Nutt, and P. Senellart. Capturing continuous data and answering aggregate queries in probabilistic XML. *TODS*, 36(4), 2011.
- 2 ACM Computing Classification System, 2012. <https://www.acm.org/about/class/class/2012>.
- 3 A. Amarilli, Y. Amsterdamer, and T. Milo. On the complexity of mining itemsets from the crowd using taxonomies. In *ICDT*, 2014.
- 4 A. Amarilli, Y. Amsterdamer, and T. Milo. Uncertainty in crowd data sourcing under structural constraints. In *UnCrowd*, 2014.
- 5 A. Amarilli, P. Bourhis, and P. Senellart. Provenance circuits for trees and treelike instances. In *ICALP*, 2015.
- 6 D. Bertsimas and S. Vempala. Solving convex programs by random walks. *JACM*, 51(4), 2004.
- 7 C. M. Bishop. Graphical models. In *Pattern Recognition and Machine Learning*, chapter 8. Springer, 2006.
- 8 J. Bragg, Mausam, and D. S. Weld. Crowdsourcing multi-label classification for taxonomy creation. In *HCOMP*, 2013.
- 9 G. Brightwell and P. Winkler. Counting linear extensions. *Order*, 8(3), 1991.
- 10 B. Büeler, A. Enge, and K. Fukuda. Exact volume computation for polytopes: a practical study. In *Polytopes – combinatorics and Computation*, 2000.
- 11 R. Cheng, D. V. Kalashnikov, and S. Prabhakar. Evaluating probabilistic queries over imprecise data. In *SIGMOD*, 2003.
- 12 E. Ciceri, P. Fraternali, D. Martinenghi, and M. Tagliasacchi. Crowdsourcing for top-k query processing over uncertain data. *IEEE TKDE*, 28(1), 2016.
- 13 S. Cohen, B. Kimelfeld, and Y. Sagiv. Running tree automata on probabilistic XML. In *PODS*, 2009.
- 14 G. Cormode, F. Li, and K. Yi. Semantics of ranking queries for probabilistic data and expected ranks. In *ICDE*, 2009.
- 15 B. Cousins and S. Vempala. A practical volume algorithm. *Mathematical Programming Computation*, 8(2), 2016.
- 16 S. B. Davidson, S. Khanna, T. Milo, and S. Roy. Using the crowd for top-k and group-by queries. In *ICDT*, 2013.
- 17 J. A. De Loera, B. Dutra, M. Köppe, S. Moreinis, G. Pinto, and J. Wu. Software for exact integration of polynomials over polyhedra. *Computational Geometry*, 46(3), 2013.
- 18 L. Detwiler, W. Gatterbauer, B. Louie, D. Suciu, and P. Tarczy-Hornoch. Integrating and ranking uncertain scientific data. In *ICDE*, 2009.
- 19 U. Faigle, L. Lovász, R. Schrader, and G. Turán. Searching in trees, series-parallel and interval orders. *SIAM J. Comput.*, 15(4), 1986.

- 20 C. Ge and F. Ma. A fast and practical method to estimate volumes of convex polytopes. In *FAW*, 2015.
- 21 J. E. Gentle. *Computational Statistics*. Springer, 2009.
- 22 Google Product Taxonomy, 2016. <https://support.google.com/merchants/answer/1705911?hl=en>.
- 23 D. Gunopulos, R. Khardon, H. Mannila, S. Saluja, H. Toivonen, and R. S. Sharma. Discovering all most specific sentences. *TODS*, 28(2), 2003.
- 24 P. Haghani, S. Michel, and K. Aberer. Evaluating top-k queries over incomplete data streams. In *CIKM*, 2009.
- 25 M. Hua, J. Pei, and X. Lin. Ranking queries on uncertain data. *VLDB J.*, 20(1), 2011.
- 26 Y. E. Ioannidis and R. Ramakrishnan. Efficient transitive closure algorithms. In *VLDB*, 1988.
- 27 J. Jests, G. Cormode, F. Li, and K. Yi. Semantics of ranking queries for probabilistic data. *IEEE TKDE*, 23(12), 2011.
- 28 F. Jones. *Lebesgue Integration on Euclidean Space*. Jones & Bartlett Learning, 2001.
- 29 K. Kangas, T. Hankala, T. Niinimäki, and M. Koivisto. Counting linear extensions of sparse posets. In *IJCAI*, 2016.
- 30 R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algorithms*, 11(1), 1997.
- 31 J. Lawrence. Polytope volume computation. *Mathematics of Computation*, 57(195), 1991.
- 32 J. Li, B. Saha, and A. Deshpande. A unified approach to ranking in probabilistic databases. *PVLDB*, 2(1), 2009.
- 33 X. Lian and L. Chen. Probabilistic ranked queries in uncertain databases. In *EDBT*, 2008.
- 34 X. Lian and L. Chen. A generic framework for handling uncertain data with local correlations. *VLDB*, 4(1), 2010.
- 35 L. Lovász and I. Deák. Computational results of an $O^*(n^4)$ volume algorithm. *European J. Operational Research*, 216(1), 2012.
- 36 L. Lovász and S. Vempala. Hit-and-run from a corner. *SIAM J. Comput.*, 35(4), 2006.
- 37 F. Maire. An algorithm for the exact computation of the centroid of higher dimensional polyhedra and its application to kernel machines. In *ICDM*, 2003.
- 38 C. H. Papadimitriou. Efficient search for rationals. *Information Processing Letters*, 8(1), 1979.
- 39 A. Parameswaran, A. Sarma, H. Garcia-Molina, N. Polyzotis, and J. Widom. Human-assisted graph search: it's okay to ask questions. *PVLDB*, 4(5), 2011.
- 40 G. Pruesse and F. Ruskey. Generating linear extensions fast. *SIAM J. Comput.*, 23(2), 1994.
- 41 L. A. Rademacher. Approximating the centroid is hard. In *SCG*, 2007.
- 42 C. Re, N. N. Dalvi, and D. Suciu. Efficient top-k query evaluation on probabilistic data. In *ICDE*, 2007.
- 43 A. Schrijver. The structure of polyhedra. In *Theory of Linear and Integer Programming*, chapter 8. Wiley-Interscience, 1986.
- 44 M. Simonovits. How to compute the volume in high dimension? *Mathematical programming*, 97(1-2), 2003.
- 45 M. A. Soliman, I. F. Ilyas, and S. Ben-David. Supporting ranking queries on uncertain and incomplete data. *VLDB J.*, 19(4), 2010.
- 46 M. A. Soliman, I. F. Ilyas, and K. C.-C. Chang. Top-k query processing in uncertain databases. In *ICDE*, 2007.
- 47 J. Stoyanovich, S. Amer-Yahia, S. B. Davidson, M. Jacob, T. Milo, et al. Understanding local structure in ranked datasets. In *CIDR*, 2013.

- 48 C. Sun, N. Rampalli, F. Yang, and A. Doan. Chimera: Large-scale classification using machine learning, rules, and crowdsourcing. *PVLDB*, 7(13), 2014.
- 49 C. Wang, L. Yuan, J. You, O. R. Zaïane, and J. Pei. On pruning for top-k ranking in uncertain databases. *PVLDB*, 4(10), 2011.
- 50 K. Yi, F. Li, G. Kollios, and D. Srivastava. Efficient processing of top-k queries in uncertain databases. In *ICDE*, 2008.
- 51 X. Zhang and J. Chomicki. Semantics and evaluation of top-k queries in probabilistic databases. *DAPD*, 26(1), 2009.

A Proofs for Section 3 (An Algorithm for Interpolation and Top-k)

► **Lemma 4.** *For any constraint set \mathcal{C} , we can construct in PTIME a constraint set \mathcal{C}' such that the probability that the possible worlds of \mathcal{C}' have a tie (under the uniform distribution) is zero, and such that any interpolation or top-k computation problem on \mathcal{C} can be reduced in PTIME to the same type of problem on \mathcal{C}' .*

Proof. We consider two types of ties: *persistent ties*, which are enforced by \mathcal{C} and hold in each possible world, and *occasional ties*, which are not enforced by \mathcal{C} and only hold in some possible worlds.

We first formally prove that occasional ties have a total probability 0: let d be the dimension of $\text{pw}(\mathcal{C})$ and $n := |\mathcal{X}|$. Assume that x, y have a tie in some worlds and do not have a tie in other worlds (in particular, $x = y$ is not implied by \mathcal{C}). If we add constraints to enforce a persistent tie, $\text{pw}(\mathcal{C} \cup \{x \leq y, y \leq x\})$ is now at most $(d - 1)$ -dimensional. Geometrically, this is a projection of the admissible polytope on the $x = y$ hyperplane. The d -volume of $\text{pw}(\mathcal{C} \cup \{x \leq y, y \leq x\})$ is thus 0, i.e., this set of worlds has probability 0. Then, taking the union of $O(n^2)$ such sets, to obtain all worlds with occasional ties between any pair of variables, the result has total probability 0.

Next, to handle persistent ties, we define a constraint set where persistently tied variables are replaced by a single variable. Let \mathcal{X}/\sim be the set of equivalence classes of \mathcal{X} . We define \mathcal{C}/\sim to be the constraint set where every occurrence of a variable $x_i \in \mathcal{X}$ is replaced by some representative $s(x_i)$ of its equivalence class. By definition \mathcal{C}/\sim does not have persistent ties; and there is clearly a bijection between $\text{pw}(\mathcal{C})$ and $\text{pw}(\mathcal{C}/\sim)$.

It remains to show that problems on \mathcal{C} can be reduced (in PTIME) to \mathcal{X}/\sim . The interpolation problem for a variable x_i on \mathcal{C} clearly reduces to interpolation for $s(x_i)$ on \mathcal{C}/\sim , since due to the bijection the expected value of x_i on \mathcal{C} equals that of $s(x_i)$ on \mathcal{C}/\sim . If no selection is applied, the reduction for top-k works by selecting the top-k variables in \mathcal{C}/\sim , replacing each $s(x_i)$ by the variables of its equivalence class $[s(x_i)]$, and truncating the obtained ranked list to length k (recall that ties can be broken arbitrarily). If a selection σ is applied, then apply top-k only to representatives of selected variables from \mathcal{X}_σ ; and replace each representative $s(x_i)$ in the top-k result only by the variables from $[s(x_i)] \cap \mathcal{X}_\sigma$. ◀

A.1 Total Orders (Section 3.1)

We now give the omitted formal details about the “Fragmented Total Orders” paragraph. The soundness of splitting can be proved via the possible world semantics: assume that $x_i \in \mathcal{X}_{\text{exact}}$. Then the possible valuations of x_1, \dots, x_{i-1} are affected only by the constraints in \mathcal{C} on x_1, \dots, x_i , and similarly for x_{i+1}, \dots, x_n . The formal claim is as follows, which follows from Lemma 12:

► **Lemma 24.** *Let \mathcal{C} be a constraint set implying a fragmented total order $x_1 \leq \dots \leq x_n$, and $x_i = v$ a constraint in \mathcal{C} . Let $\mathcal{C}_{\leq i}, \mathcal{C}_{\geq i} \subseteq \mathcal{C}$ be the constraints restricted to the variable subsets x_0, \dots, x_i and x_i, \dots, x_{n+1} respectively. Then there is a bijective correspondence between $\text{pw}_{x_1 \dots x_i}(\mathcal{C}_{\leq i}) \times \text{pw}_{x_i \dots x_n}(\mathcal{C}_{\geq i})$ and $\text{pw}(\mathcal{C})$.*

The correspondence is defined from the left-hand side to the right-hand side as merging the duplicated component for x_i to obtain worlds of exactly n values.

A.2 General Constraint Sets (Section 3.2)

► **Theorem 7.** *Given a constraint set \mathcal{C} over \mathcal{X} and $x \in \mathcal{X}$ (resp., and a selection predicate σ , and an integer k), determining the expected value of x in $\text{pw}(\mathcal{C})$ under the uniform distribution (resp., the top- k computation problem over \mathcal{X} , \mathcal{C} , and σ) is in $\text{FP}^{\#\text{P}}$.*

Proof. Let \mathcal{C} be an arbitrary constraint set with order constraints and exact-value constraints on a variable set \mathcal{X} , and let $n := |\mathcal{X}|$. Use Lemma 4 to ensure that there are no ties. To simplify the reasoning, we will make all values occurring in exact-value constraints be integers that are multiples of $(n+1)!$ as follows: let Δ be $(n+1)!$ times the product of the denominators of all exact-value constraints occurring in \mathcal{C} , which can be computed in PTIME, and consider \mathcal{C}' the constraint set defined on $[0, \Delta]^n$ by keeping the same variables and order constraints, and replacing any exact-value constraint $x_i = v$ by $x_i = v\Delta$; the constraint set \mathcal{C}' is computable in PTIME from \mathcal{C} , and the polytope $\text{pw}(\mathcal{C}')$ is obtained by scaling $\text{pw}(\mathcal{C})$ by a factor of Δ along all dimensions; hence, if we can compute the expected value of x_i in \mathcal{C}' (which is the coordinate of the center of mass of $\text{pw}(\mathcal{C}')$ on the component corresponding to x_i), we can compute the expected value of x_i in \mathcal{C} by dividing by Δ . So we can thus assume that $\text{pw}(\mathcal{C}')$ is a polytope of $[0, \Delta]^n$ where all exact-value constraints are integers which are multiples of $(n+1)!$.

We use Lemma 5.2 of [ACK⁺11] to argue that the volume of $\text{pw}(\mathcal{C})$ can be computed in $\#\text{P}$. The PTIME generating Turing machine T , given the constraint set \mathcal{C} , chooses nondeterministically a linear extension of $(\mathcal{X}, \leq_{\mathcal{C}})$, which can clearly be represented in polynomial space and checked in PTIME. The PTIME-computable function g computes, as a rational, the volume of the polytope for that linear extension, and does so according to the scheme of Section 3: the volume is the product of the volumes of each $\mathcal{C}_1^n(\alpha, \beta)$, whose volume is $\frac{(\beta-\alpha)^n}{n!}$. This is clearly PTIME-computable, and as α and β are values occurring in exact-value constraints, they are integers and multiples of $n!$, so the result is an integer, so the overall result is a product of integers, hence an integer. By Lemma 5.2 of [ACK⁺11], $V(\mathcal{C}')$, which is the sum of $V(\mathcal{T})$ across all linear extensions \mathcal{T} of \mathcal{C}' (because there are no ties), is computable in $\#\text{P}$.

We now apply the same reasoning to show that the sum, across all linear extensions \mathcal{T} , of $V(\mathcal{T})$ times the expected value of x_i in \mathcal{T} , is computable in $\#\text{P}$. Again, we use Lemma 5.2 of [ACK⁺11], with T enumerating linear extensions, and with a function g that computes the volume of the linear extension as above, and multiplies it by the expected value of x_i , by linear interpolation in the right \mathcal{C}_{p+1}^{q-1} as in the previous section (it is an integer, as all values of exact-value constraints are multiples of $q-p \leq n+1$). So this concludes, as the expected value v of x_i is $\frac{1}{V(\mathcal{C}')} \sum_{\mathcal{T}} V(\mathcal{T})v_{\mathcal{T}}$ where $v_{\mathcal{T}}$ is the expected value of x_i for \mathcal{T} , and we can compute both the sum and the denominator in $\#\text{P}$. Hence, the result of the division, and reducing back to the answer for the original \mathcal{C} , can be done in $\text{FP}^{\#\text{P}}$. ◀

A.3 Marginal Distributions

Beyond expected values, our results from Section 3 can serve to compute the *marginal distributions* of variables under the uniform distribution; these can be used for other estimation schemes of unknown variables (e.g., other moments of the variables, median values, etc.).

We now formalize this notion. Letting \mathcal{C} be a constraint set, $x \in \mathcal{X} \setminus \mathcal{X}_{\text{exact}}$ a variable with no exact-value constraint, and $v \in [0, 1]$ a value, we write $\mathcal{C}_{|x=v}$ the *marginalized* constraint set $\mathcal{C} \cup \{x = v\}$. Note that, if the admissible polytope $\text{pw}(\mathcal{C})$ is d -dimensional, then $\text{pw}(\mathcal{C}_{|x=v})$ is $(d-1)$ -dimensional. We can now define:

► **Definition 25.** The *marginal distribution* of an unknown variable $x \in \mathcal{X} \setminus \mathcal{X}_{\text{exact}}$ is defined as the pdf $p_x(v) := V_{d-1}(\mathcal{C}_{|x=v})/V_d(\mathcal{C})$.

For unfragmented total orders, the distribution of unknown variables relates to the known problem of computing *order statistics*:

► **Definition 26.** The *i-th order statistic* for n samples of a probability distribution Pr is the distribution Pr_i of the following process: draw n independent values according to Pr , and return the i -th smallest value of the draw.

► **Proposition 27 ([Gen09], p. 63).** *The distribution of the i-th order statistic for n samples of the uniform distribution on $[0, 1]$ is the Beta distribution $B(i, n + 1 - i)$.*

The connection to the marginal distribution of x_i in $\mathcal{C}_1^n(\alpha, \beta)$ is the following:

► **Observation 28.** *The marginal distribution of x_i within the admissible polytope of an unfragmented total order $\mathcal{C}_1^n(\alpha, \beta)$ is the distribution of the i-th order statistic for n samples of the uniform distribution on $[\alpha, \beta]$.*

Proof. The distribution on n uniform independent samples in $[\alpha, \beta]$ can be described as first choosing a total order for the samples, uniformly among all permutations of $\{1, \dots, n\}$ (with ties having a probability of 0 and thus being neglected). Then, the distribution for each total order is exactly that of x_i when variables are relabeled. ◀

Note that the mean of this distribution corresponds to a linear interpolation of the unknown variables between α and β , which we use in section 3.1.

We summarize these findings, as follows:

► **Proposition 29.** *Given \mathcal{C} , a constraint set implying a total order, the marginal distribution of any variable $x_i \in \mathcal{X}$ is a polynomial of degree at most n and can be computed in PTIME.*

Algorithm 1 can be adapted to compute marginal distributions, by replacing line 9 to compute instead the marginal distribution of x in \mathcal{T} on variables $x_{i_j} \dots x_{i_{j+1}}$: according to Observation 28, this marginal distribution is polynomial and defined on the range $[v_{i_j}, v_{i_{j+1}}]$. We modify line 11 to compute the sum of the marginal distributions instead, seeing them as *piecewise-polynomial* functions (which are zero outside of their range), still weighted by the volume of \mathcal{T} . We deduce that the marginal distribution of x in \mathcal{C} (computed by modifying line 13) is a *piecewise-polynomial* function with at most $|\mathcal{X}|$ pieces and degree at most $|\mathcal{X}|$.

B Proofs for Section 4 (Hardness and Approximations)

B.1 Hardness of Exact Computation (Section 4.1)

The following lemma provides a direct connection between the expected value of a variable to its expected rank, for constraint sets that consist only of order constraints. By the hardness of computing expected rank in partial orders [BW91] this provides an alternative proof to the hardness of computing expected value (Theorem 8). We use this connection to obtain an upper bound on the denominator of the expected value, when we reduce to the problem of top- k computation below.

► **Lemma 30.** *Given \mathcal{X} and a set of order constraints \mathcal{C} over \mathcal{X} , if the expected rank of $x \in \mathcal{X}$ is r then its expected value is $\frac{r}{n+1}$.*

Proof. Let N be the number of linear extensions of \mathcal{C} . Since \mathcal{C} only consists of order constraints, by symmetry, the probability of each of its linear extensions is identical, namely $1/N$. Denote by r_i the rank of x in the linear extension \mathcal{T}_i . Then $r = \sum_{i=1}^N \frac{1}{N} r_i$. By Algorithm 1, the expected value of x relative to \mathcal{T}_i is $\frac{r_i}{n+1}$, and thus the overall expected value of x relative to \mathcal{C} is $\mathbb{E}_{\mathcal{C}}[x] = \sum_{i=1}^N \frac{1}{N} \frac{r_i}{n+1} = \frac{r}{n+1}$. ◀

► **Theorem 9.** *Given a constraint set \mathcal{C} over \mathcal{X} , a selection predicate σ , and an integer k , the top- k computation problem over \mathcal{X} , \mathcal{C} and σ is $\text{FP}^{\#\text{P}}$ -hard even if k is fixed to be 1, $|\mathcal{X}_{\sigma}|$ is 2, and the top- k answer does not include the expected value of the variables.*

Proof. We will perform a reduction from the interpolation problem (i.e., expected value computation) on sets of order constraints, which is $\text{FP}^{\#\text{P}}$ -hard by Theorem 8. Write $n := |\mathcal{X}|$. Assume using Lemma 4 that the probability of ties in $\text{pw}(\mathcal{C})$ is zero.

We first observe, using Lemma 30, that the expected value v of x_i can be written as $\frac{1}{N(n+1)} \sum_{\mathcal{T}} r_{\mathcal{T}}$, where the sum is over all the linear extensions \mathcal{T} of \mathcal{C} , $r_{\mathcal{T}} \in \{1, \dots, n\}$ is the rank of x_i in the linear extension \mathcal{T} , and $N \leq n!$ is the number of linear extensions. This implies that v can be written as a rational p/q with $0 \leq p \leq q$ and $0 \leq q \leq M$, where we write $M := (n+1)!$.

We determine this fraction p/q using the algorithm of [Pap79], that proceeds by making queries of the form “is $p/q \leq p'/q'$ ” with $0 \leq p', q' \leq M$, and runs in time logarithmic in the value M , so polynomial in the input \mathcal{C} . To do so, we must describe how to decide in PTIME whether $v \leq p'/q'$ for $0 \leq p', q' \leq M$, using an oracle for the top-1 computation problem that does not return the expected values.

Fix $v' = p'/q'$ the query value and let $v = p/q$ be the unknown target value, the expected value of x_i . We illustrate how to decide whether $v \leq v'$. The general idea is to add a variable with exact-value constraint to v' and compute the top-1 between x_i and the new variable, but we need a slightly more complicated scheme because the top-1 answer variable can be arbitrary in the case where $v = v'$ (i.e., we have a tie in computing the top-1). Let $\varepsilon := 1/(2(M^2 + 1))$, which is computable in PTIME in the value of n (so in PTIME in the size of the input \mathcal{C}). Construct \mathcal{C}' (resp., \mathcal{C}'_{-} , \mathcal{C}'_{+}) by adding an exact-value constraint $x' = v'$ (resp., $x'_{-} = v' - \varepsilon$, $x'_{+} = v' + \varepsilon$) for a fresh variable x' (resp., x'_{-} , x'_{+}). Now use the oracle for \mathcal{C}' (resp., \mathcal{C}'_{-} , \mathcal{C}'_{+}) and the selection predicate that selects x_i and x' (resp., x'_{-} , x'_{+}), taking $k = 1$ in all cases. The additional variables do not affect the expected value of x_i in \mathcal{C}' , \mathcal{C}'_{+} , and \mathcal{C}'_{-} , so it is also v in them. Further, we know that $v = p/q$, $v' = p'/q'$, with $0 \leq p, q, p', q' \leq M$, hence either $v = v'$, or $|v - v'| \geq \frac{1}{M^2}$. Hence, letting $S = \{v', v' - \varepsilon, v' + \varepsilon\}$, there are three possibilities: $v = v'$, $v < v''$ for all $v'' \in S$, or $v > v''$ for all $v'' \in S$. Thus, if the top-1 variable in all oracle calls is always x_i (resp., never x_i), then we are sure that $v > v'$ (resp., $v < v'$). If some oracle calls return x_i but not all of them, we are sure that $v = v'$. Hence, we can find out in PTIME using the oracle whether $v \leq v'$. This concludes the proof, as we then have an overall PTIME reduction from the $\text{FP}^{\#\text{P}}$ -hard problem of interpolation (i.e., expected value computation) to the top-1 computation problem, showing that the latter is also $\text{FP}^{\#\text{P}}$ -hard. ◀

We conclude with a proof of the result variant stated informally after Theorem 9:

► **Theorem 31.** *Given a constraint set \mathcal{C} over \mathcal{X} and an integer k , the top- k computation problem over \mathcal{X} and \mathcal{C} is $\text{FP}^{\#\text{P}}$ -hard even for the trivial selection predicate σ (with $\mathcal{X}_{\sigma} = \mathcal{X}$) and even when the top- k answer does not include the expected value of the variables.*

Proof. We prove the result exactly like Theorem 9, except that, instead of computing the top-1 between two variables x_i and some $x' \in \mathcal{X}_{\text{exact}}$, we compute the top- k over all variables

for different values of k . Note that we can assume that the expected value v of x_i is not equal to the exact value v' of x' (otherwise, we can resolve it using the technique of Theorem 9). Under this assumption, we have $v < v'$ iff the smallest k such that x occurs in the top- k is less than the smallest k' such that x' occurs in the top- k' , and likewise $v > v'$ whenever $k > k'$. k and k' can be found e.g. using binary search. ◀

B.2 Complexity of Approximate Computation (Section 4.2)

► **Proposition 10.** ([KLS97], Algorithm 5.8). *Let \mathcal{C} be a set of constraints with variable set \mathcal{X} and $x \in \mathcal{X}$. There is an FPRAS that determines an estimate $\hat{\mathbb{E}}_{\mathcal{C}}[x]$ of the expected value $\mathbb{E}_{\mathcal{C}}[x]$ of x in $\text{pw}(\mathcal{C})$ under the uniform distribution.*

For clarity and completeness, we give a self-contained proof of this result, using the almost uniform sampling scheme of [KLS97]:

Proof. We use a result by Kannan, Lovász, and Simonovits (Theorem 2.2 of [KLS97]) that shows that sampling a point *almost uniformly* from a convex body in dimension n can be done by $\tilde{O}(n^5)$ calls to an oracle deciding membership in the body, and an additional factor of $\tilde{O}(n^2)$ arithmetic operations (here, $\tilde{O}(\cdot)$ is the *soft- O* notation, that ignores polylogarithmic factors). “Almost uniformly” means that the *total variation distance* between the uniform distribution and the actual distribution realized by the sampling is less than any fixed number $\varepsilon' > 0$. The dependency of the running time in ε' is logarithmic, hence hidden in $\tilde{O}(\cdot)$.

Let $\varepsilon, \delta > 0$ be two reals. For some number N that we define further, we first apply N times consecutively the sampling algorithm of [KLS97] with $\varepsilon' := \frac{\varepsilon}{4}$ to obtain a set $P = \{p_1, \dots, p_N\}$ of N independent samples. The algorithm uses an oracle for membership to the polytope, which in our setting can be determined in time $O(|\mathcal{C}|)$; hence, as the dimension of $\text{pw}(\mathcal{C})$ is bounded by $|\mathcal{X}|$, by the complexity analysis of [KLS97], this process has a running time of $O(N \times |\mathcal{X}|^7 \times |\mathcal{C}|)$.

We then consider the projection of every point $p_i \in P$ to the dimension defined by the variable x , giving N values v_1, \dots, v_N . We then compute the estimate $\hat{\mathbb{E}}_{\mathcal{C}}[x]$ as $\frac{1}{N} \sum_{i=1}^N v_i$. Clearly this has no impact on the running time. It only remains to show that this estimate satisfies the given bounds, i.e., that $|\hat{\mathbb{E}}_{\mathcal{C}}[x] - \mathbb{E}_{\mathcal{C}}[x]| \leq \varepsilon$ with probability at least $1 - \delta$.

Let $f : \text{pw}(\mathcal{C}) \rightarrow [0; 1]$ be the probability density function of the distribution realized by the algorithm of [KLS97], according to which the independent samples were drawn. By definition of the total variation distance, we know that $\frac{1}{2} \int_{\text{pw}(\mathcal{C})} \left| \frac{1}{V(\mathcal{C})} - f(p) \right| dp \leq \frac{\varepsilon}{4}$. If we denote by μ_v the average of the (independent) x -value samples v_1, \dots, v_N , and we denote by $f_{x=t}$ the function obtained from f by setting to t the coordinate corresponding to x , we have:

$$\begin{aligned} |\mu_v - \mathbb{E}_{\mathcal{C}}[x]| &= \left| \int_0^1 t \int_{\text{pw}(\mathcal{C}_{x=t})} f_{x=t}(p) dp dt - \int_0^1 t \int_{\text{pw}(\mathcal{C}_{x=t})} \frac{1}{V(\mathcal{C})} dp dt \right| \\ &= \left| \int_0^1 t \int_{\text{pw}(\mathcal{C}_{x=t})} \left(f_{x=t}(p) - \frac{1}{V(\mathcal{C})} \right) dp dt \right| \\ &\leq \int_0^1 1 \times \int_{\text{pw}(\mathcal{C}_{x=t})} \left| \frac{1}{V(\mathcal{C})} - f_{x=t}(p) \right| dp dt \\ &= \int_{\text{pw}(\mathcal{C})} \left| \frac{1}{V(\mathcal{C})} - f(p) \right| dp \leq \frac{\varepsilon}{2}. \end{aligned}$$

Now, as the v_i are independent and identically distributed, we obtain by Hoeffding's inequality [Hoe63]:

$$\Pr\left(\left|\hat{\mathbb{E}}_{\mathcal{C}}[x] - \mu_v\right| \geq \frac{\varepsilon}{2}\right) \leq 2 \exp\left(\frac{-\varepsilon^2 N}{2}\right).$$

Therefore, setting $N := \frac{2 \ln(\frac{2}{\delta})}{\varepsilon^2}$, we have:

$$\Pr\left(\left|\hat{\mathbb{E}}_{\mathcal{C}}[x] - \mathbb{E}_{\mathcal{C}}[x]\right| \leq \varepsilon\right) \geq \Pr\left(\left|\hat{\mathbb{E}}_{\mathcal{C}}[x] - \mu_v\right| \leq \frac{\varepsilon}{2} \wedge \left|\mu_v - \mathbb{E}_{\mathcal{C}}[x]\right| \leq \frac{\varepsilon}{2}\right) \geq 1 - \delta. \quad \blacktriangleleft$$

C Proofs for Section 5 (Tractable Cases)

C.1 Splitting Lemma (Section 5.1)

► **Lemma 12.** *Let $\mathcal{C}_1, \dots, \mathcal{C}_m$ be the uninfluence decomposition of \mathcal{C} . There exists a bijective correspondence between $\text{pw}(\mathcal{C})$ and $\text{pw}(\mathcal{C}_1) \times \dots \times \text{pw}(\mathcal{C}_m)$.*

Proof. Define the function $M : \text{pw}(\mathcal{C}) \rightarrow \text{pw}_{\mathcal{X}_1 \sqcup \mathcal{X}_{\text{exact}}}(\mathcal{C}_1) \times \dots \times \text{pw}_{\mathcal{X}_m \sqcup \mathcal{X}_{\text{exact}}}(\mathcal{C}_m)$, that maps each possible world w of \mathcal{C} to a tuple $\langle w_1, \dots, w_m \rangle$ of possible worlds of $\mathcal{C}_1, \dots, \mathcal{C}_m$ defined as follows: for each variable $x \in \mathcal{X}$, letting v its value in \mathcal{C} , we give to x the value v in all possible worlds of w_1, \dots, w_m where x appears. Recall that we assumed w.l.o.g. that $m > 0$, so each variable of \mathcal{X} must occur in at least one of w_1, \dots, w_m .

It is immediate that any $w \in \text{pw}(\mathcal{C})$ yields a tuple of possible worlds of $\mathcal{C}_1, \dots, \mathcal{C}_m$ by this definition of M , since they are subsets of \mathcal{C} .

Conversely, consider any tuple of possible worlds $\langle w_1, \dots, w_m \rangle \in \mathcal{C}_1, \dots, \mathcal{C}_m$. Note that each $x_i \in \mathcal{X}_{\text{exact}}$ is always consistently mapped to the same value as each \mathcal{C}_i contains every exact-value constraint. Every other variable $x_i \in \mathcal{X}_j$ appears only in $w_j \in \text{pw}_{\mathcal{X}_j \sqcup \mathcal{X}_{\text{exact}}}(\mathcal{C}_j)$. Thus $w = M^{-1}(\langle w_1, \dots, w_m \rangle)$ is well-defined. Let us show that $w \in \text{pw}(\mathcal{C})$.

First, we observe that exact-value constraints are necessarily respected, because each \mathcal{C}_i contains all such constraints. Second, let us show that order constraints are respected. Consider an order constraint of the form $x \leq x'$ in \mathcal{C} . Clearly, if x and x' belong to the same uninfluenced class \mathcal{X}_i (or if at least one of them has an exact-value constraint), the constraint is reflected in \mathcal{C}_i , so it must be respected in w . Hence, we focus on the case where $x \in \mathcal{X}_i$ and $x' \in \mathcal{X}_j$ with $i \neq j$. Now, as $x \leq x'$, there exists a path \mathcal{P} of the form $x = x_1 \prec \dots \prec x_n = x'$, but as they are not in the same uninfluenced class there must be a variable of $\mathcal{X}_{\text{exact}}$ in the sequence. Let x_k be this variable with exact-value v . By definition of the uninfluence decomposition, $x \leq x_k$ and $x_k = v$ are in \mathcal{C}_i ; similarly, $x_k \leq x'$ and $x_k = v$ are in \mathcal{C}_j . Thus $x \leq x'$ must be respected in w overall. \blacktriangleleft

C.2 Tree-Shaped Constraints (Section 5.2)

► **Theorem 15.** *For any tree-shaped constraint set \mathcal{C} over \mathcal{X} , we can compute its volume $V(\mathcal{C})$ in time $O(|\mathcal{X}|^2)$.*

Proof. Let T be the tree with vertex set \mathcal{X} which is the Hasse diagram of the order constraints imposed by \mathcal{C} . For any variable $x \in \mathcal{X}$ that has no exact-value constraint (so it is not the root of T or a leaf of T), let \mathcal{C}_x be the constraint set obtained as a subset of \mathcal{C} by keeping only constraints between x and its descendants in T , as well as between x and its parent. For $v \in [0, 1]$, we call $V_x(v)$ the d -volume of $\text{pw}(\mathcal{C}_x \cup \{x' = v\})$ where x' is the parent of x and d is the dimension of $\text{pw}(\mathcal{C}_x)$. In other words, $V_x(v)$ is the d -volume of the admissible polytope

for the subtree $T|_x$ of T rooted at x , as a function of the minimum value on x imposed by the exact-value constraint on the parent of x . It is clear that, letting x'_r be the one child of the root x_r of T , we have $V(\mathcal{C}) = V_{x'_r}(v_r)$, where v_r is the exact value imposed on x_r .

We show by induction on T that, for any node x of T , letting m_x be the minimum exact-value among all leaves that are descendants of x , the function V_x is zero in the interval $[m_x, 1]$ and can be expressed in $[0, m_x]$ as a polynomial whose degree is at most the number of nodes in $T|_x$, written $|T|_x|$. Since the probability of ties is 0, we have $m_x > 0$ for all x .

The base case is for a node x of T which has only leaves as children; in this case it is clear that $V_x(v)$ is $m_x - v$ for $v \in [0, m_x]$, and is zero otherwise. For the inductive case, let x be a variable. It is clear that $V_x(v)$ is 0 for $v \in [m_x, 1]$. Otherwise, let $v' \in [0, m_x]$ be the value of the parent x' of x . For every value $v' \leq v \leq m_x$ of x , consider the constraint set $\mathcal{C}_{x,v,v'} = \mathcal{C}_x \cup \{x' = v', x = v\}$. By Lemma 12, we have $V(\mathcal{C}_{x,v,v'}) = \prod_i V_{x_i}(v)$ where x_1, \dots, x_l are the children of x . Hence, by definition of the volume, we know that $V_x(v') = \int_{v'}^{m_x} \prod_i V_{x_i}(v) dv$. Now, we use the induction hypothesis to deduce that $V_{x_i}(v)$, for all i , in the interval $[0, m_x]$, is a polynomial whose degree is at most $|T|_{x_i}|$. Hence, as the product of polynomials is a polynomial whose degree is the sum of the input polynomials, and integrating a polynomial yields a polynomial whose degree is one plus that of the input polynomial, V_x in the interval $[0, m_x]$ is a polynomial whose degree is at most $|T|_x|$.

Hence, we have proved the claim by induction, and we use it to determine $V(\mathcal{C})$ as explained in the first paragraph.

We now prove that the computation is quadratic. We first assume that the tree T is binary. We show by induction that there exists a constant $\alpha \geq 0$ such that the computation of the polynomial V_{x_i} in expanded form has cost less than αn_i^2 , where n_i is $|T|_{x_i}|$. The claim is clearly true for nodes where all children are leaves, because the cost is linear in the number of child nodes as long as α is at least the number of operations per node α_0 . For the induction step, if x_i is an internal node, let x_p and x_q be the two children. By induction hypothesis, computing V_{x_p} and V_{x_q} in expanded form has cost $\leq \alpha(n_p^2 + n_q^2)$. Remembering that arithmetic operations on rationals are assumed to take unit time, computing the product of V_{x_p} and V_{x_q} in expanded form has cost linear in the product of the degrees⁶ of V_{x_p} and V_{x_q} which are less than n_p and n_q , so the cost of computing the product is $\leq \alpha_1 n_p n_q$ for some constant α_1 . Integrating has cost linear in the degree of the resulting polynomial, that is, $n_p + n_q$. So the total cost of computing V_{x_i} is $\leq \alpha(n_p^2 + n_q^2) + \alpha_1 n_p n_q + \alpha_2(n_p + n_q) + \alpha_3$ for some constants α_2, α_3 . Now, as $n_q = n_i - n_p - 1$, computing V_{x_i} costs less than:

$$\begin{aligned} & \alpha n_i^2 + 2\alpha n_p^2 + \alpha - 2\alpha n_i n_p - 2\alpha n_i + 2\alpha n_p + \alpha_1 n_p n_q + \alpha_2 n_i - \alpha_2 + \alpha_3 \\ &= \alpha n_i^2 + (\alpha_1 - 2\alpha) n_p n_q + (\alpha_2 - 2\alpha) n_i + \alpha - \alpha_2 + \alpha_3 \end{aligned}$$

As long as α is set to be $\geq \max(\frac{\alpha_1}{2}, \frac{\alpha_2}{2})$, the second and third terms are negative, which means (since $n_p n_q$ and n_i are both ≥ 1) that V_{x_i} costs less than:

$$\begin{aligned} & \alpha n_i^2 + \alpha_1 - 2\alpha + \alpha_2 - 2\alpha + \alpha - \alpha_2 + \alpha_3 \\ &= \alpha n_i^2 - 3\alpha + \alpha_1 + \alpha_3 \leq \alpha n_i^2 \end{aligned}$$

if $\alpha \geq \frac{\alpha_1 + \alpha_3}{3}$. This concludes the induction case, by setting α to any arbitrary value which is greater than $\max(\alpha_0, \frac{\alpha_1}{2}, \frac{\alpha_2}{2}, \frac{\alpha_1 + \alpha_3}{3})$. Hence the claim is proven if T is binary.

⁶ We could compute the product of the polynomials more efficiently using an FFT, but this would not improve the overall complexity bound.

If T is not binary, we use the associativity of product to make T binary, by adding virtual nodes that represent the computation of the product. In so doing, the size of T increases only by a constant multiplicative factor (recall that the number of internal nodes in a full binary tree is one less than the number of leaves, meaning that the total number of nodes in a binary expansion of a n -ary product is less than twice the number of operands of the product). So the claim also holds for arbitrary T . ◀

► **Theorem 16.** *For any tree-shaped constraint set \mathcal{C} on variable set \mathcal{X} , for any variable $x \in \mathcal{X} \setminus \mathcal{X}_{\text{exact}}$, the marginal distribution for x is piecewise polynomial and can be computed in time $O(|\mathcal{X}_{\text{exact}}| \times |\mathcal{X}|^2)$.*

Proof. Recall that $\mathcal{C}_{|x=v}$ is \mathcal{C} plus the exact-value constraint $x = v$. For any variable x' , we let $m_{x'}$ be the minimum exact-value among all leaves reachable from x' . By definition, the marginal distribution for x is $v \mapsto \frac{1}{V(\mathcal{C})} V(\mathcal{C}_{|x=v})$. We have seen in Theorem 15 that $\frac{1}{V(\mathcal{C})}$ can be computed in quadratic time; we now focus on the function $V(\mathcal{C}_{|x=v})$.

By Lemma 12, letting x_1, \dots, x_k be the children of x , D_1, \dots, D_k be their descendants (the x_i included), and D be all variables except x and its descendants, that is, $D := \mathcal{X} \setminus (\{x\} \cup \bigsqcup_i D_i)$, we can express $V(\mathcal{C}_{|x=v})$ as $V'_x(v) \times \prod_i V_{x_i}(v)$, where V_{x_i} is as in the proof of Theorem 15, and $V'_x(v)$ is the volume of the constraint set $\mathcal{C}'_{x,v}$ over D obtained by keeping all constraints in \mathcal{C} about variables of D , plus the exact-value constraint $x = v$. Indeed, the uninfluenced classes of $\mathcal{C}_{|x=v}$ are clearly $D \setminus \mathcal{X}_{\text{exact}}$, $D_1 \setminus \mathcal{X}_{\text{exact}}, \dots, D_k \setminus \mathcal{X}_{\text{exact}}$. We denote by \mathcal{X}' the variables of $\mathcal{C}'_{x,v}$.

We know by the proof of Theorem 15 that V_{x_i} , in the interval $[0, m_{x_i}]$, is a polynomial whose degree is at most $|T_{|x_i}|$, and that it can be computed in $O(|T_{|x_i}|^2)$. Hence, the product of the $V_{x_i}(v)$ can be computed in quadratic time in $|T_{|x}|$ overall (as in the proof of Theorem 15) and it has linear degree. We thus focus on $\mathcal{C}'_{x,v}$, for which it suffices to show that $V(\mathcal{C}'_{x,v})$ is a piecewise-polynomial function with at most $|\mathcal{X}_{\text{exact}}|$ pieces, each piece having a linear degree and being computable in quadratic time in \mathcal{X}' . Indeed, this suffices to justify that computing the product of $V(\mathcal{C}'_{x,v})$ with $\prod_i V_{x_i}(v)$, and integrating to obtain the marginal distribution, can be done in time $O(|\mathcal{X}_{\text{exact}}| \times |\mathcal{X}|^2)$, and that the result is indeed piecewise polynomial.

For any node x_i of D with no exact-value constraint, we let $V'_{x_i,x}(v, v')$ be the volume of the constraint set obtained by restricting $\mathcal{C}'_{x,v}$ to the descendants of the parent x'_i of x_i and adding the exact-value constraint $x'_i = v$. We let (v_1, \dots, v_q) be the values occurring in exact-value constraints in \mathcal{C} , in increasing order, so that $q \leq |\mathcal{X}_{\text{exact}}|$. We show by induction on D the following claim: for any $1 \leq i < q$, for any variable x_i in D with no exact-value constraint, in the intervals $v \in [0, m_{x_i}]$ and $v' \in [v_i, v_{i+1}]$, $V'_{x_i,x}(v, v')$ can be expressed as $P(v) + v'P'(v)$, where P and P' are polynomials of degree at most $|T_{|x_i} \cap D|$ and can be computed in quadratic time in $|T_{|x_i} \cap D|$.

The proof is the same as in Theorem 15: for the base case where all children of x_i have exact-value constraints, $V'_{x_i,x}(v, v')$ is either $m_{x_i} - v$ if x is not reachable from x_i or $v_i \geq m_x$, or $v' - v$ otherwise. For the inductive case, we do the same argument as before, noting that, clearly, taking the product of the $V'_{x'_i}(v, v')$ among the children of x_i , the variable v' occurs in at most one of them, namely the one from which x is reachable. We conclude that $V'_x(v)$ is indeed a piecewise polynomial function with at most $|\mathcal{X}_{\text{exact}}|$ many pieces, all of which have a linear degree, by evaluating $V'_{x'',x}(v'', v)$, where x'' is the one child of the root of T and v'' is the value to which it has an exact-value constraint. The overall computation time is then in $O(|\mathcal{X}_{\text{exact}}| \times |\mathcal{X}|^2)$. ◀

► **Corollary 17.** *Given any constraint set \mathcal{C} and its uninfluence decomposition $\mathcal{C}_1, \dots, \mathcal{C}_m$, assuming that each \mathcal{C}_i is a (reverse-)tree-shaped constraint set, we can solve the interpolation*

problem in time $O(\max_i |X_i|^3)$ and the top- k problem in $PTIME$.

Proof. First, notice that any reverse-tree-shaped constraint set \mathcal{C} can be transformed to a tree-shaped constraint set \mathcal{C}' such that $\mathbb{E}_{\mathcal{C}'}(x) = 1 - \mathbb{E}_{\mathcal{C}}(x)$ for every $x \in \mathcal{X}$, by reversing order constraints and replacing exact-value constraints $x = \alpha$ with $x = 1 - \alpha$. Second, we also observe that, formally, the constraint sets in the uninfluence decomposition may include variables with exact value constraints that are not connected to the tree-shaped structure, but it is obvious that these variables have no impact on the interpolation problem.

Now, we use Lemma 12 to deduce that we can indeed solve the problems by solving them separately in each constraint set of the uninfluence partition. For the interpolation problem, we can compute the interpolated value of a variable by looking only at its uninfluence class. The complexity of top- k follows. ◀

D Proofs for Section 6 (Other Variants)

D.1 Proofs of Comparison Results

► **Lemma 19.** *There is a constraint set \mathcal{C} and selection predicate σ such that local-top- k and U-top- k do not match, even for $k = 1$ and without returning expected values or probabilities.*

Proof. Let $\mu = 2/3$, $m = 1/\sqrt{2}$, and pick any v such that $\mu < v < m$. Consider variables x , x' and y , with the constraint set that imposes $x' \leq x$ and $y = v$. Fix $k = 1$ and consider the predicate σ that selects all variables. It is immediate by linear interpolation that the expected value of x is μ . Hence y , that has a greater expected value, is the local-top-1. However, the marginal distribution of x can easily be computed to be $p_x : t \mapsto 2t$. Intuitively, when x becomes larger, it makes a larger range of values possible for x' , and thus also a larger range of possible worlds. By integrating, we obtain that x has a probability of 0.5 of exceeding m – a value greater than v , y 's value – and hence its probability of being the top-1 is greater than y 's. Namely, x is the U-top-1. Intuitively, the volume of possible worlds where x is the top-1 is greater due to the asymmetry of x 's distribution. However, there is a larger *gap* (on average) between y and x in worlds where y is the top-1, which ultimately leads to a higher expected value for y . ◀

► **Lemma 21.** *There is a constraint set \mathcal{C} without ties such that U-top- k does not satisfy the containment property for the uniform distribution on $\text{pw}(\mathcal{C})$.*

Proof. Consider variables x_1 , x_h , x_f^+ , and x_f^- , and the constraint set \mathcal{C} that imposes $x_1 \leq x_h$, $x_f^+ = .7$, and $x_f^- = .69$. Consider the selection predicate σ that selects all variables. The total volume of the constraint set is clearly $V = 1/2$.

We first set $k = 1$. The first possible answer is (x_f^+) with probability $\frac{.7 \times .7}{2 \cdot V} = .49$, and the second is (x_h) with probability .51, so the U-top-1 is (x_h) .

We then set $k = 2$. There are four possible answers. The first possible answer is (x_f^+, x_f^-) with probability $\frac{.69 \times .69}{2 \cdot V} = .4761$. The second possible answer is (x_h, x_f^+) with probability $\frac{.3 \times .7}{V} = .42$. The third possible answer is (x_h, x_1) with probability .09. The fourth possible answer is (x_f^+, x_h) with probability $\frac{.01 \times .69 + .01 \times .01 \times .5}{V} = .0139$. Hence, the U-top-2 is (x_f^+, x_f^-) .

Hence, the U-top-1 variable does not occur in the U-top-2. ◀

► **Lemma 23.** *There is a constraint set \mathcal{C} without ties such that global-top- k does not satisfy the containment property for the uniform distribution on $\text{pw}(\mathcal{C})$.*

Proof. Consider variables x_s, x_f, x_l and x_h and the constraint set \mathcal{C} that imposes $x_l \leq x_h$, $x_l = .45$ and $x_f = .73$. Consider σ that selects all variables.

Set $k := 1$. Variable x_h has the highest value with probability $\frac{1}{\sqrt{v}}(.73 \times (1 - .73) + (1 - .73)^2/2)$. Variable x_f has the highest value with probability $\frac{1}{\sqrt{v}}(.73 - .45) \times .73$, which is less. The probability for x_s is also less. So the global-top-1 is (x_h) .

Now, set $k := 2$. Variable x_h has one of the two highest values in all cases except for $x_f \geq x_s \geq x_h \geq x_l$ and $x_s \geq x_f \geq x_h \geq x_l$, so it has one of the two highest values with probability $1 - \frac{1}{\sqrt{v}}((.73 - .45) \times (1 - .73) + (.73 - .45)^2/2)$. However, variable x_f has one of the two highest values in all cases except for $x_h \geq x_s \geq x_f \geq x_l$ and $x_s \geq x_h \geq x_f \geq x_l$, so it has one of the two highest values with probability $1 - \frac{1}{\sqrt{v}}(1 - .73)^2$, which is more. Hence the first variable of the global-top-2 is x_f and not x_h . ◀

D.2 Other Variants

Additional variants of top- k have been studied, see [CLY09, ZC09]. However, in the context of [CLY09], these definitions do not satisfy the containment property either, except for two. The first, U-kRanks [SIC07], does not satisfy the natural property that top- k answers always contain k different variables. The second, expected ranks [CLY09], resembles local-top- k but uses ranks instead of values, so the definition is *value-independent*. While this makes sense for top- k queries designed to return tuples, as in [CLY09], we argue it is less sensible when focusing on the numerical value of variables; this justifies our focus on local-top- k .

Another possibility to define top- k in our context, however, is to design it based on different assumptions. One natural choice is to require a *stability* property, namely, adding exact-value constraints to fix some variables to their interpolated values does not change the interpolated values of the other variables. We can show that this property is not respected by our scheme, but that it can be enforced on tree-shaped constraint sets: see Appendix E.

E Alternative Interpolation Scheme

We now consider variants of the interpolation problem, thus far performed by considering the expected value under the uniform distribution. Since we are not aware of candidate variants in previous work, for interpolating over partial orders with exact-value constraints, we propose a new alternative variant. Rather than imposing the connection to the uniform prior, we present natural desiderata for an interpolation scheme on partial orders. We show that they are not respected by our current definition, and show that a definition that respects them can be proposed for tree-shaped partial orders (Definition 14), it is in fact unique, and it can be computed tractably.

We first define the notion of interpolation scheme.

► **Definition 32.** An *interpolation scheme* is a function that maps any constraint set \mathcal{C} on variables \mathcal{X} to a mapping from \mathcal{X} to its interpolated value in $[0, 1]$.

For instance, the interpolation scheme that we have studied thus far maps each variable to its expected value under the uniform distribution on $\text{pw}(\mathcal{C})$. We refer to this scheme as **Uniform** in the sequel.

We define the first natural desideratum for interpolation schemes, stability: intuitively, an interpolation scheme is stable if assigning variables to their interpolated value does not change the result of interpolation elsewhere. Formally,

► **Definition 33.** An interpolation scheme \mathcal{S} is *stable* if, for every constraint set \mathcal{C} over \mathcal{X} and every $x \in \mathcal{X}$, \mathcal{S} assigns the same mapping $f : \mathcal{X} \rightarrow [0, 1]$ to both \mathcal{C} and $\mathcal{C} \cup \{x = f(x)\}$.

This property can be shown to be respected, e.g., by linear interpolation on total orders. However, a counterexample shows that the stability property is not respected by the uniform scheme:

► **Lemma 34.** *The Uniform scheme is not stable, even on tree-shaped constraint sets.*

Proof. Consider the set of variables $\{x_r, x_a, x_b, x_c, x_d, x_e\}$ and the constraint set \mathcal{C} formed of the order constraints $x_r \leq x_a$, $x_a \leq x_b \leq x_c$, $x_a \leq x_d \leq x_e$, and the exact-value constraints $x_r = 0$, $x_c = .5$ and $x_e = 1$. We can compute that the interpolated values for x_a and for x_b are $3/20$ and $13/40$ respectively. However, adding the exact-value constraint $x_b = 13/40$, the interpolated value for x_a becomes $611/4020$, which is different from $3/20$. ◀

Let us use stability as a guide to design a different interpolation scheme. We impose another desideratum to act as a *base case*, specifying what one expects from an interpolation scheme when there is only a single unknown variable:

► **Definition 35.** Let \mathcal{C} be a (non-contradictory) constraint set such that x is the only unknown variable; y_1, \dots, y_n are variables with exact-value constraints such that $y_i \leq x$; and z_1, \dots, z_m are variables with exact-value constraints such that $x \leq z_i$ (having $n, m \geq 1$). We say an interpolation scheme is *balanced* if, for each such \mathcal{C} , its interpolated value for x is $\frac{\max_i(v(y_i)) + \min_i(v(z_i))}{2}$.

In particular, the Uniform scheme is balanced in this sense. However, we would like to find a scheme that is *both balanced and stable*. For the case of general constraint sets, this problem remains open, and we leave it for future work. For tree-shaped constraint sets, we next not only show such a scheme, but also prove it is unique, as follows.

► **Proposition 36.** *There is at most one interpolation scheme on trees that is both stable and balanced.*

Proof. We first observe that, because of the balanced requirement, in any constraint set \mathcal{C} , for any unknown variable x whose parent y was interpolated to value v and whose children z_i were interpolated to w_i , x must have been interpolated to $(v + (\min_i w_i))/2$. Indeed, considering the constraint set \mathcal{C}' where y and z_i have been set to those values, by stability, the interpolation value of x does not change. Hence, as the interpolation scheme is balanced, we conclude that the claimed property holds.

We now show that the resulting set of equations always has at most one solution on any tree-shaped constraint sets. Indeed, assume that there are two stable and balanced interpolation schemes f and g which yield different results on a tree-shaped constraint set \mathcal{C} . For all variables x of \mathcal{C} , let $d(x) := g(x) - f(x)$. Calling x_r the root variable of \mathcal{C} , we must have $d(x_r) := 0$ for the root, because it has an exact-value constraint by definition of tree-shaped constraint sets.

Now, as f and g differ on \mathcal{C} , there must be a variable x where $d(x) \neq 0$. Without loss of generality, we have $d(x) > 0$. Hence, let us consider a variable x with parent y so that we have $d(x) > d(y)$: as $d(x_r) = 0$, we can find such a variable x by picking a variable which is as high as possible in the tree, such that $d(x) > 0$ but $d(y) = 0$. Necessarily x is not a leaf (as they have exact-value constraints, so $d(x) = 0$), so x has children. We show that x has a child x_g such that $d(x_g) > d(x)$.

Consider x_f the child of x such that $f(x_f)$ is minimal among children of x , and x_g defined in the analogous manner for g . Now, as f and g are balanced, by our preliminary observation we have $f(x) = (f(y) + f(x_f))/2$ hence $f(x_f) = 2 \cdot f(x) - f(y)$, and likewise $g(x_g) = 2 \cdot g(x) - g(y)$.

But then, by minimality of $g(x_g)$, we have $g(x_g) - f(x_f) \leq g(x_f) - f(x_f)$. Now, we have $g(x_g) - f(x_f) = 2 \cdot d(x) - d(y)$. Now, as we have $d(y) < d(x)$, we have $d(x_f) > d(x)$, which is what we wanted to show.

Now, repeating the argument on x_f , we obtain a child x_f^2 of x_f such that $d(x_f^2) > d(x_f)$. Repeating the argument, we thus build a descending chain of variables x in the tree-shaped constraint set \mathcal{C} along which d is strictly increasing. When we reach the leaves, we obtain a contradiction. This implies that we must have $d(x) = 0$ for all $x \in \mathcal{X}$, so that $f = g$ on \mathcal{C} . Hence, there cannot be two different stable and balanced interpolation schemes on tree-shaped constraint sets. \blacktriangleleft

We now prove the *existence* of a stable and balanced interpolation scheme on trees, which we dub **Stable**, and show that expected values under this scheme can be computed in linear time:

► **Proposition 37.** *There exists a stable and balanced interpolation scheme on tree-shaped constraint sets, and we can compute the interpolated values of all variables according to this scheme in time $O(|\mathcal{X}|^2)$.*

Proof. We compute the interpolation scheme on a tree-shaped constraint set \mathcal{C} top-down. For each variable x which has no exact-value constraint or interpolated value, but whose parent y has an exact value or an interpolated value v_y , we consider all leaves z which are descendants of x (and have an exact-value constraint to some value v_z), and set the interpolated value of x to be the minimum of linear interpolation from y to z ; namely, letting $d_x(z)$ be the depth of leaf z in the subtree rooted at x , we set $v_x := \min_z \left(v_y + \frac{v_z - v_y}{d_x(z) + 1} \right)$. This can clearly be done in the indicated time bound.

We now show that the resulting interpolation scheme is indeed balanced and stable. It is immediate to observe that this scheme is balanced. We now show that it is stable. Towards this, let us first show that for every variable x with parent y , if z is a leaf that achieved the minimum when interpolating x to its value, then for all variables on the path from x to z , z was also a leaf that achieved the minimum when interpolating their value. Indeed, it suffices to show the claim for the first variable x' of this path, a child of x , and then repeat the argument. We know that, from our choice when interpolating x , by minimality of the interpolation result for x using z , we have $\frac{v_z - v_y}{d_x(z) + 1} \leq \frac{v_{z'} - v_y}{d_x(z') + 1}$, where v_y is the value of the parent y of x ; let us call the first quantity δ and the second δ' . By definition of the interpolation of x , we then have $v_x := v_y + \delta$. Now, consider any leaf z' reachable from x' . We must show that z achieves the minimum when interpolating x' ; in other words, we must compare $\eta := \frac{v_z - v_{x'}}{d_{x'}(z) + 1}$ and $\eta' := \frac{v_{z'} - v_{x'}}{d_{x'}(z') + 1}$, and show that $\eta \leq \eta'$; note that $d_{x'}(z) + 1 = d_x(z)$ and $d_{x'}(z') + 1 = d_x(y)$. The quantity η can then be rewritten as $\frac{d_x(z) + 1}{d_x(z)} \times \frac{1}{d_x(z) + 1} (v_z - v_y - \delta)$, i.e., $\frac{d_x(z) + 1}{d_x(z)} \times \left(\delta - \frac{\delta}{d_x(z) + 1} \right)$, which simplifies to δ : hence, $\eta = \delta$. The quantity η' can be written as $\frac{d_x(z') + 1}{d_x(z')} \times \frac{1}{d_x(z') + 1} (v_{z'} - v_y - \delta)$, i.e., $\frac{d_x(z') + 1}{d_x(z')} \times \left(\delta' - \frac{\delta}{d_x(z') + 1} \right)$, which simplifies to $\frac{(d_x(z') + 1)\delta' - \delta}{d_x(z')}$. Now, as $\delta' \geq \delta$, we deduce that $\eta' \geq \frac{(d_x(z') + 1)\delta - \delta}{d_x(z')}$, so that $\eta' \geq \delta$. Hence, we have $\eta' \geq \eta$, so that the leaf z also achieves the minimum for variable x' . Repeating the argument on the path from x to z , we have shown the claim.

From this initial claim, we deduce the following (*): for any variable x , letting z be a leaf that achieves the minimum when interpolating x (once the value of its parent y is known), then the variables in the path from y to z are interpolated according to linear interpolation on that chain. This is immediate by the previous claim, as all variables on

that path are interpolated using linear interpolation from their parent to that same leaf (or another minimal leaf that sets them to the same value).

We now show a similar auxiliary claim. Let us define, once we have interpolated in \mathcal{C} , the function u that maps each non-root variable x to $u(x)$ defined as the interpolated value of x minus that of its parent. We show that (**): for any variables y, x, x' , where y is the parent of x and x is the parent of x' , then $u(x) \leq u(x')$. Indeed, let v_y, v_x and $v_{x'}$ be the interpolated values, and let z' be the witness leaf used to interpolate for x' . By definition of the scheme, we have $u(x') = \frac{v_{z'} - v_x}{d_{x'}(z') + 1}$. Furthermore, letting z be the witness leaf used to interpolate for x , we have $u(x) = \frac{v_z - v_y}{d_x(z) + 1}$. Using the notation above, note that $u(x') = \eta'$ and $u(x) = \delta$. By the same reasoning as for claim (*) to show $\delta = \eta \leq \eta'$, we conclude that $u(x) \leq u(x')$.

We are now ready to show that the scheme is stable. Consider the initial tree-shaped constraint set \mathcal{C} , and let us set a variable x to its interpolated value v_x , yielding \mathcal{C}' . Note that \mathcal{C}' is no longer tree-shaped, but it can be rewritten by Lemma 12 to two tree-shaped constraint sets. It is then clear that all variables that are descendants of x in \mathcal{C} are interpolated in the same manner in \mathcal{C}' and in \mathcal{C} , as the scheme proceeds top-down and the value of x in \mathcal{C}' is by definition the same as its interpolated value in \mathcal{C} . We now show that the ancestors of x in \mathcal{C} are interpolated in the same way in \mathcal{C}' than in \mathcal{C} , which is clearly sufficient to justify the claim that *all* variables in \mathcal{C}' are interpolated in the same way as in \mathcal{C} . Let us therefore pick an ancestor x' of x , which is neither x nor the root variable, otherwise the claim is trivial; we pick it as high as possible in the tree, so the interpolated value v_y of its ancestor y is the same in \mathcal{C} and in \mathcal{C}' .

We first show that the interpolated value for x' in \mathcal{C}' is no higher than in \mathcal{C} . Assuming to the contrary that it is, then it must be the case that x' was interpolated in \mathcal{C} using as minimal leaf z some leaf which is a descendant of x in \mathcal{C} , as otherwise we can still interpolate using z in \mathcal{C}' and obtain the same result. Now, if x' was interpolated in \mathcal{C} using z as minimal leaf, then, by our preliminary claim (*), x was interpolated in \mathcal{C} following linear interpolation between the parent y of x' and the leaf z . Hence, using the new leaf x in \mathcal{C}' to interpolate x' in \mathcal{C}' yields the same result as the interpolation in \mathcal{C} . Contradiction.

Second, we show that the interpolated value for x' in \mathcal{C}' is no lower than in \mathcal{C} . Assuming to the contrary that it is, then, if x' was interpolated in \mathcal{C}' following a leaf z which is not x , then we immediately reach a contradiction as we should have used the same leaf z to interpolate to the same value in \mathcal{C} . Hence, we must have interpolated x' in \mathcal{C}' using the new leaf x , and x' was interpolated in \mathcal{C}' following linear interpolation between y and x . Let γ be the value difference between two consecutive nodes in \mathcal{C}' on this path, and l the length of the path. Calling $u(x)$ for a variable x the difference between its interpolated value and the value of its parent in \mathcal{C} , we must then have $u(x') > \gamma$, because the value of x' in \mathcal{C} is strictly greater than in \mathcal{C}' , and y has same value in \mathcal{C} and \mathcal{C}' . By preliminary claim (**), we have reached a contradiction, because then the function u always takes values which are $> \gamma$ on the path from x' to x in \mathcal{C} , so that when we reach x we know that the value of x in \mathcal{C} is $> l \cdot \gamma$, contradicting the fact that it is $l \cdot \gamma$, as we know from \mathcal{C}' . ◀

References for the Appendix

- ACK⁺11** S. Abiteboul, T.-H. H. Chan, E. Kharlamov, W. Nutt, and P. Senellart. Capturing continuous data and answering aggregate queries in probabilistic XML. *TODS*, 36(4), 2011.
- BW91** G. Brightwell and P. Winkler. Counting linear extensions. *Order*, 8(3), 1991.
- CLY09** G. Cormode, F. Li, and K. Yi. Semantics of ranking queries for probabilistic data and expected ranks. In *ICDE*, 2009.

- Gen09** J. E. Gentle. *Computational Statistics*. Springer, 2009.
- Hoe63** W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), 1963.
- KLS97** R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Struct. Algorithms*, 11(1), 1997.
- Pap79** C. H. Papadimitriou. Efficient search for rationals. *Information Processing Letters*, 8(1), 1979.
- SIC07** M. A. Soliman, I. F. Ilyas, and K. C.-C. Chang. Top-k query processing in uncertain databases. In *ICDE*, 2007.
- ZC09** X. Zhang and J. Chomicki. Semantics and evaluation of top-k queries in probabilistic databases. *DAPD*, 26(1), 2009.