

Monodic Fragments of Probabilistic First-order Logic

Jean Christoph Jung¹, Carsten Lutz¹, Sergey Goncharov², and Lutz Schröder²

¹ Universität Bremen, {jeanjung,clu}@informatik.uni-bremen.de,

² Friedrich-Alexander-Universität Erlangen-Nürnberg,
{Sergey.Goncharov,Lutz.Schroeder}@fau.de

Abstract. By classical results of Abadi and Halpern, validity for probabilistic first-order logic of type 2 (ProbFO) is Π_1^2 -complete and thus not recursively enumerable, and even small fragments of ProbFO are undecidable. In temporal first-order logic, which has similar computational properties, these problems have been addressed by imposing *monodicity*, that is, by allowing temporal operators to be applied only to formulas with at most one free variable. In this paper, we identify a monodic fragment of ProbFO and show that it enjoys favorable computational properties. Specifically, the valid sentences of monodic ProbFO are recursively enumerable and a slight variation of Halpern’s axiom system for type-2 ProbFO on bounded domains is sound and complete for monodic ProbFO. Moreover, decidability can be obtained by restricting the FO part of monodic ProbFO to any decidable FO fragment. In some cases, which notably include the guarded fragment, our general constructions result in tight complexity bounds.

1 Introduction

Both logic and probability theory are fundamental to the formalization and solution of many important problems in computer science. While logic is a way to address the combinatorics hidden in such problems, the main use of probabilities is to capture uncertainty that arises from many different sources such as noisy or untrusted data (in database systems), a high level of abstraction (in verification), or incomplete training data (in machine learning). Unfortunately, the combination of logic and probability is notoriously difficult and involves a large number of choices and trade-offs, which has resulted in a broad spectrum of probabilistic logic formalisms to be proposed that vary greatly in spirit, semantics, and expressive power.

A natural and fundamental way to combine logic and probabilities is to enrich classical first-order logic (FO) with a probabilistic component [12, 4, 5]. Although reasoning in the resulting probabilistic FO logics is, of course, undecidable, they are still useful as a general and uniform ‘baseline formalism’ that encompasses many other probabilistic logics, much in the same way that FO provides a baseline formalism for many other logics used in computer science. However, it turns out that probabilistic FO logics are not only undecidable, but tend to be computationally even less well-behaved than classical FO. They come in essentially two versions, called type-1 and type-2 [12]. While type-1 is for reasoning about statistical probabilities, reflected in the semantics by a probability distribution over the domain of the FO structure, the purpose of type-2 is reasoning about subjective probabilities by adopting a possible worlds semantics. In

this paper, we concentrate on the latter and use ‘ProbFO’ to refer to Halpern’s probabilistic FO logic of type-2 [12]. The disastrous computational behaviour of ProbFO was analyzed by Abadi and Halpern, who show that validity is Π_1^2 -complete [1], thus outside the arithmetic and analytic hierarchies and, in particular, far from being recursively enumerable. This result holds up even when only unary predicates are admitted, and we add in this work the observation that ProbFO is still Π_1^1 -hard even with only two (object) variables (and no quantification over real-valued variables, see below).

Our aim in the current work is to analyze how and how far the problematic computational properties of ProbFO can be improved. We start by observing that there is a clear semantic and computational similarity between ProbFO and temporal first-order logic (TFO). Both logics use a possible worlds semantics, and although TFO is only Π_1^1 -complete, just like ProbFO it is not recursively enumerable. In the case of TFO, Hodkinson, Wolter and Zakharyashev have given an elegant explanation of why this is the case and how better computational properties can be recovered, by introducing the *monodic* fragment of TFO that restricts temporal operators to be applied only to formulas with at most one free variable [17]. In fact, monodic TFO turns out to be recursively enumerable [23] and decidable fragments of monodic TFO can often be obtained by restricting the FO part of monodic TFO to a decidable FO fragment [18, 14, 16, 15]. In the present work, we identify a monodic fragment of ProbFO and show that, as in the case of TFO, this recovers good computational properties. Note that the formulas of unrestricted ProbFO are obtained by combining classical FO with the language of real closed fields (including quantification over real numbers) via real-valued terms of the form $w(\varphi)$ denoting the probability that the formula φ is true. Atomic formulas in this extended language of real closed fields are called *weight formulas*. In analogy to TFO, a natural candidate for monodicity in ProbFO is to admit only weight terms $w(\varphi)$ in which φ has at most one free first-order variable. We show that this is not an effective choice since the resulting fragment of ProbFO still fails to be recursively enumerable.

We thus have to adopt stronger restrictions and define a ProbFO formula to be monodic if it contains no variables for real numbers (thus no quantification over the reals) and every weight formula in it contains at most one free (object) variable. Under this definition, we can establish a useful abstract representation of models of monodic ProbFO formulas – so-called *quasi-models* – which are essentially a collection of monadic formula types that satisfy certain integrity conditions and are associated with a system of polynomial inequalities over the reals to capture probabilities. This representation yields in a rather direct way that monodic ProbFO is recursively enumerable. Moreover, we exploit quasi-models to establish a concrete axiomatization of monodic ProbFO, a variation of a complete axiomatization of unrestricted ProbFO on finite domains of fixed size by Halpern [12] (we use unrestricted domains). Finally, quasi-models can be used to identify decidable fragments of monodic ProbFO. We show that for any FO-fragment \mathcal{L} such that a slightly generalized version of satisfiability in \mathcal{L} called *realizability* is decidable, monodic Prob \mathcal{L} is decidable, too. In particular, we thus obtain decidability for the case where \mathcal{L} is the monadic fragment of FO, the guarded fragment (GF), the two-variable fragment, and the guarded negation fragment. The finite model property transfers in the same way.

We also analyze the computational complexity of some important decidable fragments of monodic ProbFO. The naive version of our general algorithm yields a $2\text{NEXPTIME}^{\exists\mathbb{R},\mathcal{C}}$ upper bound where superscripts denote access to oracles, $\exists\mathbb{R}$ is the class of problems that reduce in polynomial time to solving systems of polynomial inequalities over the reals (recall $\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}$), and \mathcal{C} is the complexity of deciding realizability in the underlying FO fragment \mathcal{L} . We then propose two improvements to our algorithm. The first one consists of a more careful realizability check as known from monodic TFO, and this modification sometimes allows removing the oracle for \mathcal{C} . For monodic ProbGF, in particular, we obtain in this way an improved $2\text{NEXPTIME}^{\exists\mathbb{R}}$ upper bound. The second improvement applies only when \mathcal{L} satisfies a certain model-theoretic property that we call *closedness under unions of types*, and it allows improving the runtime by one exponential by reducing the size of quasi-models. GF satisfies the mentioned property, and thus we obtain a tight 2EXPTIME upper bound for monodic ProbGF. We also obtain a $\text{NEXPTIME}^{\exists\mathbb{R}}$ upper bound when the arity of predicates is bounded, and a tight NEXPTIME upper bound for the case where only linear weight formulas are admitted, that is, multiplication of probabilities is disallowed. Note that the relatively high computational complexities are partly due to the fact that we aim at identifying *maximal* decidable fragments of monodic ProbFO. In fact, monodic ProbFO can be viewed as a natural generalization of the family of probabilistic description logics introduced in [20, 11] and provides a principled explanation for why these logics are computationally much more well-behaved than traditional ProbFO. Conversely, the mentioned description logics can be viewed as fragments of monodic ProbFO with lower computational complexity, typically EXPTIME -complete.

Proofs are generally omitted or only sketched; full proofs can be found in the appendix.

2 Preliminaries

Type-2 probabilistic first-order logic (ProbFO) [12] comprises two sorts: objects of the domain of discourse and the real numbers \mathbb{R} . Accordingly, there are *object variables* and *field variables*, the latter being used to represent probabilities. *Object terms* are object variables or *object constants*. *ProbFO-formulas* and *field terms* are defined by mutual recursion:

$$\begin{aligned} \varphi, \psi &::= R(t_1, \dots, t_k) \mid \varphi \wedge \psi \mid \neg\varphi \mid \forall x \varphi(x) \mid f_1 \leq f_2 \\ f_1, f_2 &::= 0 \mid 1 \mid \mathbf{w}(\varphi) \mid f_1 + f_2 \mid f_1 \times f_2 \end{aligned}$$

where R is a k -ary predicate symbol, t_1, \dots, t_k are object terms, and formulas of the form $f_1 \leq f_2$ are called *weight formulas*. Quantification is possible both over object and field variables, with field variables ranging over \mathbb{R} . We use $\text{ProbFO}^=$ to denote the extension of ProbFO with equality on object terms. We could admit rational constants in field terms to represent concrete probabilities, but as usual we refrain from doing so because rational constants can be eliminated by clearing denominators.

Formulas of ProbFO are interpreted in *probabilistic structures* $\mathfrak{M} = (D, W, \mu, \pi)$ that consist of a non-empty domain D , a set of worlds W , a discrete probability distribution μ over W and a function π that maps each pair (R, w) to a subset of D^k and each pair (c, w) to an element of D for each k -ary predicate symbol R , $w \in W$, and constant symbol c . Intuitively, \mathfrak{M} can be viewed as a set of classical FO structures (over the same

domain) with weights given by μ . A *valuation* for \mathfrak{M} is a function ν that maps object variables to elements of D and field variables to real numbers. Given \mathfrak{M} , ν , and a world $w \in W$, the interpretation $[f]_{(\mathfrak{M}, w, \nu)} \in \mathbb{R}$ of a field term f is defined in the natural way, with terms $w(\varphi)$ interpreted as $[w(\varphi)]_{(\mathfrak{M}, w, \nu)} = \mu(\{w' \in W \mid (\mathfrak{M}, w', \nu) \models \varphi\})$. The semantics of formulas is standard. A ProbFO-sentence φ is *satisfiable* if there is a probabilistic structure $\mathfrak{M} = (D, W, \mu, \pi)$ and a world $w \in W$ such that $(\mathfrak{M}, w) \models \varphi$. A sentence φ is *valid* if $\neg\varphi$ is not satisfiable.

When we speak of (non-probabilistic) first-order logic (FO), we mean the FO fragment of ProbFO as introduced above. In particular, we mean FO without equality unless we write $\text{FO}^=$. A classical FO-structure has the form $\mathfrak{A} = (A, \pi)$ where A is a domain and π is a function as above except that its second argument (the world) is omitted.

3 Monodic ProbFO

Abadi and Halpern have shown that validity in ProbFO is Π_1^2 -complete, and thus highly undecidable and far from being recursively enumerable [1]. They also show that already over vocabularies that contain *only constants*, validity is Π_∞^1 -complete when equality is allowed. The lower bounds of these theorems are proved by reductions of suitable higher-order theories of integer arithmetics. We give additional evidence of the computational difficulty of ProbFO by proving (in the full version) the following orthogonal result by a reduction of recurring domino systems that is rather different in spirit from the mentioned reductions from integer arithmetic.

Theorem 1. *Validity in ProbFO is Π_1^1 -hard even if quantification over field variables is disallowed and only two object variables are admitted.*

The mentioned previous results and Theorem 1 illustrate that several restrictions of ProbFO that might at first sight seem promising fail to improve the computational properties of this logic. Inspired by the good computational properties of *monodic fragments* of temporal first-order logic [17, 23], we aim to define monodic fragments of ProbFO that are computationally well-behaved. In the context of temporal first-order logic, a formula is monodic when temporal operators are applied only to formulas with at most one free variable. We first show that one has to be careful when adapting this notion to ProbFO; the following result is proved by a reduction of finite validity in FO.

Theorem 2. *Validity in ProbFO is Π_1^0 -hard even if only one free object variable is allowed to occur in weight formulas.*

Although a natural candidate for monodicity, the restriction formulated in Theorem 2 is thus not strong enough to regain recursive enumerability. Intuitively, this is because it is still possible to compare the probabilities of different domain elements such as in the formula $\forall x \forall r (w(A(x)) = r \Rightarrow \exists y w(A(y)) = r/2)$, which says that for each object x , there is an object y that has half the probability of satisfying A . To avoid this, we require a weight formula with a free object variable to have *no* other free variables (object or field). This restriction makes field variables and quantification over them mostly useless, so we disallow them altogether.

Definition 3 (Monodic ProbFO formula). A ProbFO formula is *monodic* if it contains no field variables and every weight formula contains at most one free (object) variable.

We will see that the above definition of monotonicity indeed guarantees good computational properties such as recursive enumerability of validity. In the balance, of course, we lose some expressive power, in particular the ability to relate different domain elements in terms of their probabilities. The following proposition gives explicit examples of ProbFO-formulas that cannot be expressed in monodic ProbFO. Its proof relies on Theorem 7 below, which states that every satisfiable monodic ProbFO sentence is satisfiable in a model with only finitely many worlds. In the full version, we show how to enforce infinitely many worlds using the formulas in Proposition 4.

Proposition 4. *The following formulas are not expressible in monodic ProbFO:*

1. $w(P(x, y)) \sim p$ with P binary, $p \in (0, 1)$, and $\sim \in \{<, \leq, =, \geq, >\}$;
2. $w(A(x)) > w(A(y))$ with A unary.

Formulas as in Item 1 can be used to express that any two persons who show up at a party together and both wear rings are probably married; with a formula as in Item 2, we could say that children are more likely to use a smartphone than their parents. Note that formulas such as $w(\exists y P(x, y)) \sim p$ and $w(\forall y P(x, y)) \sim p$, which are similar to the formulas in Item 1 but only have one free variable, *do* fall within monodic ProbFO.

The following theorem illustrates that the positive results for monodic ProbFO rely on disallowing equality; it is again proved by reduction of finite validity in FO.

Theorem 5. *Validity in monodic ProbFO⁼ is Π_1^0 -hard.*

4 The Quasi-Model Machinery

We introduce quasi-models, an abstraction of probabilistic structures that underlies the proofs of all positive results established in this paper. This requires some preliminaries. In the following, fix a monodic ProbFO-sentence φ_0 . We denote by $\text{sub}(\varphi_0)$ the set of all subformulas of φ_0 and their negations, and by $\text{sub}_n(\varphi_0)$ the formulas from $\text{sub}(\varphi_0)$ with precisely n free variables, for $n \in \{0, 1\}$. By $\text{con}(\varphi_0)$, we denote the set of all constant symbols that occur in φ_0 . Reflecting monotonicity, we concentrate on formulas with at most one free variable when defining quasi-models. In particular, these formulas are from the following set, where x is a distinguished variable:

$$\text{sub}_x(\varphi_0) = \text{sub}_0(\varphi_0) \cup \{\psi(x), \psi(c) \mid \psi(y) \in \text{sub}_1(\varphi_0), c \in \text{con}(\varphi_0)\}.$$

We introduce a way to represent ProbFO formulas as FO formulas by replacing weight formulas with new predicates. Introduce a fresh nullary predicate symbol P_ψ for every weight formula $\psi \in \text{sub}_0(\varphi_0)$ and a fresh unary predicate symbol P_ψ for every weight formula $\psi \in \text{sub}_1(\varphi_0)$. Denote by $\bar{\varphi}$ the FO formula that is obtained from the ProbFO formula φ by replacing each weight formula $\psi()$ (resp. $\psi(x)$) that is not within the scope of another weight formula with $P_\psi()$ (resp. $P_\psi(x)$). This notation is lifted to sets of formulas in the obvious way.

A *type* is a subset t of $\text{sub}_x(\varphi_0)$ such that the set of FO formulas \bar{t} is a maximal satisfiable subset of $\text{sub}_x(\varphi_0)$. Intuitively, a type is a set of FO formulas with one free variable that are satisfied by a domain element in a world of a probabilistic structure; it also records the sentences true in that world, including the FO formulas with one

free variable that are satisfied by constants. Two types t_1, t_2 *agree on sentences*, written $t_1 \equiv_0 t_2$, if for all sentences $\psi \in \text{sub}_x(\varphi_0)$, we have $\psi \in t_1$ iff $\psi \in t_2$.

A *world type* is a set of types that agree on sentences; it can be viewed as an abstract representation of a world in a probabilistic structure, that is, of an FO structure. For an FO structure $\mathfrak{A} = (A, \pi)$ and an element $d \in A$, define

$$\text{tp}(\mathfrak{A}, d) = \{\psi \in \text{sub}_x(\varphi_0) \mid \mathfrak{A} \models \overline{\psi}[d]\} \quad \text{and} \quad \text{tp}(\mathfrak{A}) = \{\text{tp}(\mathfrak{A}, d) \mid d \in A\}.$$

Note that $\text{tp}(\mathfrak{A}, d)$ is a type and $\text{tp}(\mathfrak{A})$ is a world type. A world type T is *realizable* if there is an FO structure \mathfrak{A} such that $\text{tp}(\mathfrak{A}) = T$, that is, if the FO formula $\overline{\chi}(T)$ is satisfiable, where we define

$$\chi(T) = \bigwedge_{t \in T} \exists x \bigwedge t(x) \wedge \forall x \bigvee_{t \in T} \bigwedge t(x).$$

World types will play a central role in the definition of quasi-models, but need to be suitably enriched with (i) runs that describe the types of a single domain element in *all* worlds of a probabilistic structure and (ii) relevant conditions that have to be satisfied by the probabilities of worlds. Note that runs and world types in a sense represent orthogonal dimensions. Let Q be a set of world types. A *run through Q* is a function r that assigns to each world type $T \in Q$ a non-empty set $r(T) \subseteq T$ and is *coherent*, that is, whenever some $t \in r(T)$ contains a weight formula θ , then for all $T' \in Q$ and $t' \in r(T')$, we have $\theta \in t'$. Coherence allows us to write $\theta \in r$ to denote that for all (equivalently: some) $T \in Q$ and $t \in r(T)$, we have $\theta \in t$. A run selects a set of types for each world type instead of only a single type because each world type can represent several actual worlds, and an element might have different types in each of these worlds. A *quasi-model candidate* is a triple (T_0, Q, R) with T_0 a world type, Q a set of world types, and R a set of runs through $Q \cup \{T_0\}$ such that for all $T \in Q \cup \{T_0\}$ and $t \in T$, there is a run $r \in R$ with $t \in r(T)$. Intuitively, T_0 describes a (single) world of probability 0 while each $T \in Q$ describes worlds of positive probability. To address Point (ii) above and obtain our final quasi-model representation, we augment quasi-model candidates with a system of polynomial inequalities. It uses a variable x_T for each world type T to represent the probability of T (obtained by summing up the probabilities of all worlds of world type T) and a variable $x_{r,t,T}$ for each run r , world type T , and type $t \in T$ to describe the (summed up) probability of those worlds of world type T in which the element described by run r has type t .

Definition 6 (Quasi-Model). A quasi-model candidate (T_0, Q, R) is a *quasi-model* if every $T \in Q \cup \{T_0\}$ is realizable and the following system of polynomial inequalities $\mathcal{E}(Q, R)$ has a *positive* solution over the reals:

1. distribution on world types: $\sum_{T \in Q} x_T = 1$;
2. the probabilities of the types associated by a run $r \in R$ to a quasi-world $T \in Q$ sum up to the probability of T : $x_T = \sum_{t \in r(T)} x_{r,t,T}$;
3. runs respect weight formulas, that is, for all $f_1 \sim f_2 \in r$ with $\sim \in \{\leq, >\}$ ³ we include an equation $[f_1]_r \sim [f_2]_r$ where $[f]_r$ is obtained from f by replacing each outermost term $w(\psi(x))$ with the following expression describing its probability:

$$\sum_{T \in Q} \sum_{t \in r(T), \psi(x) \in t} x_{r,t,T}.$$

³ We write $f_1 > f_2 \in r$ in place of $f_1 \leq f_2 \notin r$.

Note that the field terms f_1, f_2 in Item 3 of Definition 6 can contain addition and multiplication, thus the system $\mathcal{E}(Q, R)$ need not be linear.

We say that a quasi-model candidate (or quasi model) (T_0, Q, R) *satisfies* a ProbFO sentence φ_0 if $\varphi_0 \in t$ for some $t \in T_0$. The following provides the basis for our use of quasi-models in subsequent sections.

Theorem 7. *A monodic ProbFO sentence φ_0 is satisfiable iff it is satisfied in some quasi-model. Moreover, any satisfiable monodic ProbFO sentence is satisfied in a probabilistic structure with finitely many worlds.*

In the “ \Rightarrow ” direction, we read off a quasi-model satisfying φ_0 from a probabilistic structure that satisfies φ_0 . To show that the system $\mathcal{E}(Q, R)$ has a solution, the values for the variables x_T and $x_{r,t,T}$ are also read off in a straightforward way.

The “ \Leftarrow ” direction is more interesting. Let (T_0, Q, R) be a quasi-model that satisfies φ_0 . Hence, every $T \in Q \cup \{T_0\}$ is realizable and $\mathcal{E}(Q, R)$ has a positive solution; we use x_T^* to denote the value of x_T in this solution and likewise for $x_{r,t,T}^*$. To construct a probabilistic structure \mathfrak{M} that satisfies φ_0 , it would be convenient to use the world types in Q as worlds. Since runs can associate more than one type with a world type, though, this is not sufficient. We thus need to subdivide each $T \in Q$ into several worlds, each accommodating a *single* type that a given run assigns to T . This has to be done in a careful way since we have to do this simultaneously for *all runs* while also ensuring that all types in T are realized in each of the worlds that T is subdivided into.

Let $r \in R$ and $T \in Q$. A *subdivision of T for r* is a tuple $s = (b_1, \dots, b_n, \zeta)$ such that $b_1 < b_2 < \dots < b_n = x_T^*$, $n = |r(T)| + 1$, and ζ is a surjective function that assigns to every b_i a type $\zeta(b_i) \in r(T)$ such that for all $t \in r(T)$ we have $\sum_{i \in [1, n], \zeta(b_i) = t} (b_i - b_{i-1}) = x_{r,t,T}^*$ where, here and in what follows, $b_0 := 0$. Intuitively, the interval $[0, x_T^*]$ represents the probability covered by all worlds of type T and we subdivide this range into the intervals $(b_i, b_{i+1}]$, with $i < n$. Elements described by the run r then have type $\zeta(b_{i+1})$ in the interval $(b_i, b_{i+1}]$. For easier reference, we say for all $p \in (0, x_T^*]$ that s has type t at p if $\zeta(b_i) = t$ and $p \in (b_{i-1}, b_i]$. A *subdivided run* is a pair (r, S) with r a run through Q and S a function that assigns to every $T \in Q$ a subdivision $S(T)$ of T for r . If we had only the single run r , we could use the subintervals identified by a subdivided run (r, S) as worlds. Since this is not the case, we first identify a sufficiently rich set of subdivisions which we then combine into a finer ‘overall’ subdivision: in the full version, we show how to define a finite set Γ of subdivided runs such that

- (*) for all $T \in Q$, $t \in T$, and $p \in (0, x_T^*]$, there is some $(r, S) \in \Gamma$ such that $S(T)$ has type t at p .

To define the worlds for a world type T , let $z_1 < \dots < z_m$ be all numbers that occur in a subdivision for T in (a subdivided run from) Γ . We introduce one world of type T for every z_i and assign to it the probability $z_i - z_{i-1}$, with $z_0 := 0$. Note that the probabilities of all worlds for T sum up to x_T^* . Doing this for all world types T (and adding one world with world type T_0 and probability 0) gives us the set of worlds W for the desired probabilistic structure \mathfrak{M} along with their probabilities $\mu(w)$. Since every $T \in Q \cup \{T_0\}$ is realizable, we find for each $w \in W$ an FO structure \mathfrak{A}_w that realizes the world type T associated with w . The domain of \mathfrak{M} is the disjoint union of the domains

of all these \mathfrak{A}_w (recall that we do not allow equality), and the further construction of \mathfrak{M} is detailed in the full version. Notably, (*) guarantees that every $t \in T$ is realized in every world w associated with world type T .

5 Recursive Enumerability and Axiomatization

We now show that the set of valid monodic ProbFO sentences is recursively enumerable and also provide a concrete axiomatization. For the former, it suffices to provide a semi-decision procedure for unsatisfiability, based on Theorem 7. The crucial observation is that, for any input sentence φ_0 , the number of quasi-model candidates (T_0, Q, R) that satisfy φ_0 is bounded. It is thus possible to construct all quasi-model candidates that satisfy φ_0 and then eliminate those that do not satisfy the system of polynomial inequalities $\mathcal{E}(Q, R)$ from Definition 6. Then, enumerate all unsatisfiable FO formulas. For each such formula ψ , eliminate all quasi-model candidates (T_0, Q, R) such that $\bar{\chi}(T) = \psi$ for some $T \in Q \cup \{T_0\}$ (since T is not realizable, (T_0, Q, R) cannot be a quasi-model). Once all quasi-model candidates have been eliminated, return with ‘ φ_0 is unsatisfiable’.

Theorem 8. *The set of valid monodic ProbFO sentences is recursively enumerable.*

Halpern gives an axiomatization of ProbFO for the case where probabilistic structures are restricted to a domain of bounded size [12]. We propose a variation of this axiomatization that is sound and complete for monodic ProbFO (without assuming bounded domains). Let AX_2 be the set of the following axioms:

- *PC*: an axiomatization of FO [7];⁴
- *OF*: all instances of the axioms of ordered fields (formulated in terms of \leq) that are well-formed formulas in monodic ProbFO;
- *PW₁*: $\varphi \Rightarrow (\mathfrak{w}(\varphi) = 1)$ if all occurrences of predicate symbols in φ are inside the scope of $\mathfrak{w}()$;
- *PW₂*: $\mathfrak{w}(\varphi) \geq 0$;
- *PW₃*: $\mathfrak{w}(\varphi \wedge \psi) + \mathfrak{w}(\varphi \wedge \neg\psi) = \mathfrak{w}(\varphi)$;
- *PW₄*: $\mathfrak{w}(\exists x \varphi(x)) > 0 \Rightarrow \exists x \mathfrak{w}(\varphi(x)) > 0$;
- *RPW*: from $\varphi \equiv \psi$ infer $\mathfrak{w}(\varphi) = \mathfrak{w}(\psi)$.

In comparison to Halpern’s axiomatization, we have removed the axiom FIN_N for bounded domains of size N and added axiom PW_4 . This axiom follows from Halpern’s axiomatization, but is independent of the axioms that remain when FIN_N is removed – in a nutshell, its soundness over discrete measures depends on σ -additivity, while PW_3 captures only finite additivity. Moreover, as we exclude field variables, we no longer need the full axiomatization of real-closed fields but, by the Artin-Schreier Theorem [3], can make do with the axioms of ordered fields. These can be phrased as quantifier-free open formulas (e.g. $x \geq 0 \vee -x \geq 0$) and hence can be instantiated to monodic ProbFO formulas (by replacing real variables with weight terms, observing the monodicity restriction).

⁴ Since constants can be interpreted differently in different worlds, a slight adaptation of the definition of when a term t is *substitutable* for x in the axiom $\forall x \varphi \Rightarrow \varphi(x/t)$ is necessary [12].

Theorem 9. AX_2 axiomatizes validity in monodic ProbFO.

Soundness is proved essentially as in [12]. For showing completeness, we make use of Theorem 7. We use $AX_2 \vdash \varphi$ to denote that φ can be derived in AX_2 , and call a sentence φ *consistent* if $AX_2 \vdash \neg\varphi$ does not hold. By Theorem 7, it suffices to show that if a monodic ProbFO sentence φ_0 is consistent, then there is a quasi-model that satisfies φ_0 . The strategy is to use consistency of φ_0 to derive a consistent sentence φ' that describes a quasi-model that satisfies φ_0 . The general structure of φ' is

$$\chi(T_0) \wedge \sum_{T \in Q} w(\chi(T)) = 1 \wedge \bigwedge_{T \in Q} w(\chi(T)) > 0 \wedge \Psi(x_1, \dots, x_k)$$

where T_0 is a world type that contains some t with $\varphi_0 \in t$, Q is a set of world types, $\Psi(x_1, \dots, x_k)$ is a quantifier-free formula in which each free variable x_i identifies a run r_i through $Q \cup \{T_0\}$, and Ψ is a conjunction of weight formulas with weight terms of the form $w(\chi(T))$ and $w(\chi(T) \wedge t(x_i))$ for some $t \in T$. Intuitively, these weight terms correspond to the variables x_T and $x_{r_i, t, T}$, respectively, in $\mathcal{E}(Q, R)$; moreover, $\Psi(x_1, \dots, x_k)$ describes precisely $\mathcal{E}(Q, R)$ under this correspondence. Observe that φ' is consistent relative to OF and thus has a solution which is also a solution to $\mathcal{E}(Q, R)$.

6 Decidability and Complexity

Theorem 7 reduces satisfiability in monodic ProbFO to satisfiability in FO and solvability of systems of polynomial inequalities over the reals. In the following, we use this observation to establish decidability results for fragments of monodic ProbFO that are obtained by restricting its FO part to a decidable FO fragment such as the guarded fragment or the two-variable fragment. We also derive complexity results, which in some cases are tight. For a fragment \mathcal{L} of FO, *monodic Prob \mathcal{L}* is the fragment of monodic ProbFO that consists of all formulas φ such that, for all $\psi \in \text{sub}(\varphi)$, the FO formula $\bar{\psi}$ belongs to \mathcal{L} . To warm up, we start with considering the finite model property (FMP). Recall that, by Theorem 7, even full monodic ProbFO has the FMP regarding the number of worlds. Here, we thus mean the number of domain elements.

Theorem 10. *For an FO fragment \mathcal{L} , monodic Prob \mathcal{L} has the FMP iff \mathcal{L} has the FMP.*

Theorem 10 is a direct consequence of the proof of Theorem 7. In the “if”-direction of that proof, we combine FO structures that witness realizability of world types. If \mathcal{L} has the finite model property, we can choose these structures to be finite. Then, the resulting probabilistic structure is also finite.

Based on quasi-models, transfer of decidability is also easy to establish. We say that *realizability is decidable in \mathcal{L}* if it is decidable whether a given world type T formulated in monodic Prob \mathcal{L} is realizable, that is, whether the \mathcal{L} formula $\bar{\chi}(T)$ is satisfiable.

Theorem 11. *If realizability is decidable in the FO fragment \mathcal{L} , then so is satisfiability in monodic Prob \mathcal{L} .*

Theorem 11 is established by the following algorithm which decides satisfiability of a given Prob \mathcal{L} sentence φ_0 :

1. guess a quasi-model candidate (T_0, Q, R) that satisfies φ_0 ;

2. verify that the system $\mathcal{E}(Q, R)$ has a positive solution in \mathbb{R} ;
3. verify that each world type $T \in Q \cup \{T_0\}$ is realizable.

Step 1 is effective since the size of quasi-model candidates is bounded by a computable function in the size of φ_0 (analyzed in more detail below).

Theorem 11 applies for instance to the monadic fragment of FO (MonaFO), the guarded fragment (GF) [2], the guarded negation fragment (GNFO) [6], and the two-variable fragment FO_2 [10]: In all these cases, the formulas $\bar{\chi}(T)$ for checking realizability remain within the fragment, and satisfiability in all the mentioned fragments is decidable.

Corollary 12. *Let \mathcal{L} be one of MonaFO, GF, GNFO, FO_2 . Then satisfiability in monodic $\text{Prob}\mathcal{L}$ is decidable.*

To analyze the complexity of the algorithm from the proof of Theorem 11, first note that it suffices to guess a quasi-model candidate (T_0, Q, R) of size at most double exponential in the size of φ_0 . In fact, Q contains at most double exponentially many world types T , and each T contains at most exponentially many types. While R can in principle be larger than double exponential, it suffices to include one run r for each $T \in Q \cup \{T_0\}$ and $t \in T$, such that $t \in r(T)$. Considering for example GF in which satisfiability is 2EXPTIME -complete, we thus obtain a $2\text{NEXPTIME}^{\exists\mathbb{R}, 2\text{EXPTIME}}$ upper bound for satisfiability in monodic ProbGF where the superscripts indicate access to two oracles: one for solving systems of polynomial inequalities over the reals and one for realizability in GF. Recall that $\exists\mathbb{R}$ denotes the class of all problems that are reducible in polynomial time to solving the mentioned systems [22], and that $\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}$.

For many FO fragments \mathcal{L} , though, we can improve on the upper bounds obtained in this direct way. First, it is helpful to not consider satisfiability of the exponential size realizability formula $\bar{\chi}(T)$ as a black box. In particular, the regular structure of $\bar{\chi}(T)$ implies that its satisfiability can be decided in time double exponential in the size of φ_0 for GF and in space exponential in the size of φ_0 for both MonaFO and FO_2 [16]. This yields a $2\text{NEXPTIME}^{\exists\mathbb{R}}$ upper bound for monodic ProbGF , monodic ProbMonaFO , and monodic ProbFO_2 . Second, for some FO fragments \mathcal{L} the quasi-model machinery can be refined so that each quasi-model candidate has at most exponential size. The following is a sufficient condition for when this is possible.

Definition 13. An FO fragment \mathcal{L} is *closed under unions of types* if for each \mathcal{L} -sentence ψ and any two structures \mathfrak{A}_1 and \mathfrak{A}_2 that satisfy the same sentences from $\text{sub}_x(\psi)$, there is a structure \mathfrak{B} such that $\text{tp}(\mathfrak{B}) = \text{tp}(\mathfrak{A}_1) \cup \text{tp}(\mathfrak{A}_2)$.

For GF without constant symbols, closure under unions of types can be shown easily by taking disjoint unions. The following result is proved in the full version.

Theorem 14. *If \mathcal{L} is closed under unions of types, then for every satisfiable monodic $\text{Prob}\mathcal{L}$ sentence φ_0 , there is a quasi-model (T_0, Q, R) that satisfies φ_0 and in which no two distinct world types agree on sentences.*

As a consequence of Theorem 14, we obtain the following improved complexity bounds for monodic ProbGF .

Corollary 15. *Satisfiability in monodic ProbGF is*
(a) 2EXPTIME -complete;

- (b) in $\text{NEXPTIME}^{\exists\mathbb{R}}$ when the arity of predicates is bounded;
(c) NEXPTIME -complete when additionally only linear weight formulas are allowed.

For part (a), where our general machinery even yields a tight upper bound, it suffices to guess a quasi-model candidate (T_0, Q, R) of exponential size (Theorem 14); the associated system $\mathcal{E}(Q, R)$ is then also of exponential size and thus the existence of a solution can thus be checked in space exponential in the size of the input formula φ_0 since $\exists\mathbb{R} \subseteq \text{PSPACE}$. It remains to verify that every world type is realizable, in time double exponential in the size of φ_0 . The lower bound is inherited from satisfiability in GF. For part (b), we can argue analogously with the difference that realizability can be checked in exponential time. For part (c), observe that in this case $\mathcal{E}(Q, R)$ is a system of linear inequalities and can thus be solved in polynomial time. The lower bound follows from the fact that the NEXPTIME -hard modal logic S5_{ALC} [8] is contained in this fragment.

Other FO fragments such as FO_2 and MonaFO are not closed under unions of types. Consider for example the FO_2 sentence $\psi = \forall x(\forall y R(x, y) \vee \forall y \neg R(y, x))$ which states that R is either the full relation or the empty relation. It does this in a slightly unorthodox way to ensure that no sentence from $\text{sub}_x(\psi)$ can distinguish the two cases. But the cases are distinguished in types because if R is full, then every type contains the formula $\forall y R(x, y)$ and if R is empty, then every type contains its negation. It is thus easy to show that closure under unions of types fails.

7 Conclusion

We have analyzed the reasons for the bad computational behaviour of ProbFO and we have shown that, unlike other natural restrictions that fail to establish recursive enumerability and decidability, monodicity is able to tame ProbFO computationally. We thus believe that monodic ProbFO lays a promising foundation for identifying decidable and useful probabilistic logics for computer science.

An interesting direction for further research is to enrich monodic ProbFO with additional expressive power that enables more complex and succinct statements about independence and conditioning. Note that existing decidable probabilistic first-order logics used in statistical relational learning such as Markov logic [9, 21] are largely orthogonal to monodic ProbFO as they typically assume bounded domains and their main use is to encode a fixed distribution for a propositional theory (over ground instances).

Another important extension to be investigated is to combine statistical and subjective probabilities in a probabilistic FO logic. A basic version of ProbFO that combines both kinds of probability was considered by Halpern [12] under the name *type-3 ProbFO*, and later refined to include stronger forms of independence and conditioning [19]. Adapting the quasi-model machinery to type-3 ProbFO and the mentioned extensions is a challenging open research objective.

We have given a tight upper complexity bound for monodic ProbGF . An open problem that remains is to determine the exact computational complexity of other relevant fragments of monodic ProbFO such as monodic ProbMonaFO , ProbFO_2 , and negation-guarded monodic ProbFO .

Acknowledgments. This work was supported by the DFG project Probabilistic Description Logics (LU1417/1-1, SCHR1118/6-1).

References

1. M. Abadi and J. Halpern. Decidability and expressiveness for first-order logics of probability. *Inf. Comput.*, 112:1–36, 1994.
2. H. Andréka, J. van Benthem, and I. Németi. Back and forth between modal logic and classical logic. *Logic J. IGPL*, 3:685–720, 1995.
3. E. Artin and O. Schreier. Algebraische Konstruktion reeller Körper. *Abh. Math. Sem. Univ. Hamburg*, 5:85–99, 1927.
4. F. Bacchus. *Representing and reasoning with probabilistic knowledge - a logical approach to probabilities*. MIT Press, 1990.
5. F. Bacchus, A. Grove, D. Koller, and J. Halpern. From statistics to beliefs. In *Artificial Intelligence, AAAI 1992*, pp. 602–608. AAAI Press/The MIT Press, 1992.
6. V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. In *Automata, Languages and Programming, ICALP 2011*, vol. 6756 of LNCS, pp. 356–367. Springer, 2011.
7. H. B. Enderton. *A mathematical introduction to logic*. Academic Press, 1972.
8. D. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyashev. *Many-dimensional modal logics: theory and applications*, vol. 148 of *Studies in Logic*. Elsevier, 2003.
9. L. Getoor and B. Taskar. *Introduction to Statistical Relational Learning*. MIT Press, 2007.
10. E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bull. Symb. Log.*, 3:53–69, 1997.
11. V. Gutiérrez-Basulto, J. Jung, C. Lutz, and L. Schröder. A closer look at the probabilistic description logic prob- \mathcal{EL} . In *Artificial Intelligence, AAAI 2011*. AAAI Press, 2011.
12. J. Halpern. An analysis of first-order logics of probability. *Artif. Intell.*, 46:311–350, 1990.
13. D. Harel. Recurring dominoes: making the highly undecidable highly understandable. *Ann. Discrete Math.*, 24:51–72, 1985.
14. I. Hodkinson. Monodic packed fragment with equality is decidable. *Stud. Log.*, 72:185–197, 2002.
15. I. Hodkinson. Complexity of monodic guarded fragments over linear and real time. *Ann. Pure Appl. Logic*, 138:94–125, 2006.
16. I. Hodkinson, R. Kontchakov, A. Kurucz, F. Wolter, and M. Zakharyashev. On the computational complexity of decidable fragments of first-order linear temporal logics. In *Proc. TIME-ICTL 2003*, pp. 91–98. IEEE Computer Society, 2003.
17. I. Hodkinson, F. Wolter, and M. Zakharyashev. Decidable fragment of first-order temporal logics. *Ann. Pure Appl. Logic*, 106:85–134, 2000.
18. I. Hodkinson, F. Wolter, and M. Zakharyashev. Monodic fragments of first-order temporal logics: 2000-2001 A.D. In *Logic for Programming, Artificial Intelligence, and Reasoning, LPAR 2001*, vol. 2250 of LNCS, pp. 1–23, 2001.
19. D. Koller and J. Halpern. Irrelevance and conditioning in first-order probabilistic logic. In *Proc. AAAI/IAAI 1096*, pp. 569–576. AAAI Press / The MIT Press, 1996.
20. C. Lutz and L. Schröder. Probabilistic description logics for subjective uncertainty. In *Principles of Knowledge Representation and Reasoning, KR 2010*. AAAI Press, 2010.
21. M. Richardson and P. Domingos. Markov logic networks. *Machine learning*, 62:107–136, 2006.
22. M. Schaefer. Complexity of some geometric and topological problems. In *Graph Drawing, GD 2009*, vol. 5849 of LNCS, pp. 334–344. Springer, 2010.
23. F. Wolter and M. Zakharyashev. Axiomatizing the monodic fragment of first-order temporal logic. *Ann. Pure Appl. Logic*, 118:133–145, 2002.

Appendix

Proofs for Section 3

Theorem 1. *Validity in ProbFO is Π_1^1 -hard even if quantification over field variables is disallowed and only two object variables are admitted.*

Proof. The following version of the tiling problem is known to be Σ_1^1 -hard [13]. A *tiling problem* is a quadruple $P = (T, H, V, t_r)$, where T is a finite set of *tile types*, $H, V \subseteq T \times T$ are the *horizontal and vertical matching conditions*, and $t_r \in T$ is the *recurrent tile*. A *solution* to P is a mapping $\tau : \mathbb{N} \times \mathbb{N} \rightarrow T$ such that

- $(\tau(i, j), \tau(i + 1, j)) \in H$ for all $i, j \geq 0$;
- $(\tau(i, j), \tau(i, j + 1)) \in V$ for all $i, j \geq 0$;
- there are infinitely many $j \geq 0$ such that $\tau(0, j) = t_r$.

Let P be a tiling problem as above. We use the FO dimension to represent the horizontal dimension of the grid, and possible worlds for the vertical dimension. To represent the successor relation in the horizontal direction, we introduce a binary relation R . Of course, R should be unbounded and rigid, there should be exactly one domino tile at every grid element, and the tiling should be compatible with horizontal successors:

$$\mathbf{w}(\forall x \exists y R(x, y)) = 1 \quad (1)$$

$$\forall x \forall y R(x, y) \Rightarrow \mathbf{w}(R(x, y)) = 1 \quad (2)$$

$$\mathbf{w}(\forall x (\bigvee_{t \in T} X_t(x) \wedge \bigwedge_{t, t' \in T, t \neq t'} \neg(X_t(x) \wedge X_{t'}(x)))) = 1 \quad (3)$$

$$\mathbf{w}(\forall x \forall y (R(x, y) \Rightarrow \bigvee_{(t, t') \in H} X_t(x) \wedge X_{t'}(y))) = 1. \quad (4)$$

To represent the successor relation in the vertical direction, we use probabilities. More precisely, a grid node in row i is represented by a domain element that satisfies A with probability $1/2^i$. To make this work, we first enforce that the probability of any element to satisfy A is $1/2^i$ for some $i \in \mathbb{N}$, and that all probabilities of this form are present:

$$\forall x (\mathbf{w}(A(x)) = 1 \vee \exists y (\mathbf{w}(A(y)) = 2\mathbf{w}(A(x)))) \quad (5)$$

$$\forall x \exists y 2\mathbf{w}(A(y)) = \mathbf{w}(A(x)) \quad (6)$$

These probabilities, though, are still associated with the FO dimension. To transfer our probability scheme to the dimension of possible worlds, we force that in every world, there is at least one element that satisfies the unary predicate M and all elements that satisfy M agree on the probability of satisfying A ; moreover, every element is marked in at least one world:

$$\mathbf{w}(\exists x M(x)) = 1 \wedge \quad (7)$$

$$\mathbf{w}(\forall x \forall y ((M(x) \wedge M(y)) \Rightarrow \mathbf{w}(A(x)) = \mathbf{w}(A(y)))) = 1 \quad (8)$$

$$\forall x \mathbf{w}(M(x)) > 0. \quad (9)$$

In this way, every *world* is associated with a unique probability: the probability of the M -marked elements to satisfy A ; conversely, for each probability $p = 1/2^i$, there is a world associated with p .⁵ We can now enforce as follows that the tiling is compatible also with vertical successors:

$$\mathbf{w}(\forall x \exists y \bigvee_{(t,t') \in V} (X_t(x) \wedge M(y) \wedge \mathbf{w}(\psi_{t'}(x,y)) = 1)) = 1 \quad (10)$$

where

$$\psi_{t'}(x,y) = \exists x (M(x) \wedge \mathbf{w}(A(y)) = 2\mathbf{w}(A(x))) \Rightarrow X_{t'}(x).$$

It remains to enforce that the recurring tile t_r occurs infinitely often. We first introduce a new unary concept name C_0 that marks the first column and ensure that, in this column, the recurring tile t_r occurs at least once:

$$\exists x (\mathbf{w}(C_0(x)) = 1 \wedge \mathbf{w}(X_{t_r}(x)) > 0). \quad (11)$$

Now infinite occurrence of t_r can be expressed as follows:

$$\begin{aligned} \mathbf{w}(\forall x ((C_0(x) \wedge X_{t_r}(x)) \Rightarrow \\ \exists y (\exists x (M(x) \wedge \mathbf{w}(A(y)) < \mathbf{w}(A(x))) \wedge \\ \mathbf{w}(\vartheta(x,y)) = 1))) = 1 \end{aligned} \quad (12)$$

where

$$\vartheta(x,y) = \exists x (M(x) \wedge \mathbf{w}(A(y)) = \mathbf{w}(A(x))) \Rightarrow X_{t_r}(x).$$

Let φ_P be the conjunction of all ProbFO sentences above. It remains to show that P has a solution iff φ_P is satisfiable.

“if”: Let $\mathfrak{M} = (D, W, \mu, \pi)$ be a probabilistic structure that satisfies φ_P . By (5), there is some $e \in D$ satisfying A with probability 1. By (9), there is some world w_0 associated to this e , thus the probability associated to w_0 is 1. Starting at w_0 , fix an infinite sequence of worlds w_0, w_1, w_2, \dots such that the associated probability of w_i is $1/2^i$ (this is possible as, by (6), every such probability exists).

By (11), there is some $d_0 \in D$ that satisfies C_0 . Observe that due to (12) d_0 satisfies X_{t_r} in infinitely many worlds among w_0, w_1, \dots . Starting at d_0 , fix an infinite sequence of elements d_0, d_1, d_2, \dots such that d_i is related to d_{i+1} by R , for all $i \geq 0$ and in all worlds $w \in W$ (possible because of (1) and (2)). We can read off a mapping $\tau : \mathbb{N} \times \mathbb{N} \rightarrow T$ as follows:

$$\tau(i,j) = t \quad \Leftrightarrow \quad d_i \in \pi(X_t, w_j).$$

The mapping τ is well-defined because of (3) and the horizontal matching condition is satisfied due to (4). Finally observe that the vertical matching condition is satisfied because of (10).

⁵ Note that the probabilities associated with worlds in the described way are not the probability that is assigned by the probabilistic structure to that world via its distribution over worlds; in particular, the probabilities associated with worlds need not sum up to one.

“only if”. Given a solution τ to P we define a probabilistic structure $\mathfrak{M} = (D, W, \mu, \pi)$ by taking $D = W = \mathbb{N}$ and

$$i \in \pi(X_t, j) \Leftrightarrow \tau(i, j) = t \quad \text{for all } i, j \in \mathbb{N}, t \in T.$$

That is, domain elements model the horizontal dimension and worlds the vertical dimension of the tiling. It remains to give the interpretation of the relations C_0, R, A , and M and ensure that world i is associated with probability $1/2^i$ (in the sense explained above).

- $\mu(i) = 1/2^{i+1}$;
- $\pi(R, j) = \{(i, i+1) \mid i \in \mathbb{N}\}$ for all $j \in \mathbb{N}$;
- $\pi(A, j) = \{i \mid i \leq j\}$;
- $\pi(M, i) = \{i\}$ for all $i \in \mathbb{N}$;
- $\pi(C_0, i) = \{0\}$ for all $i \in \mathbb{N}$.

It is now not hard to verify that $\mathfrak{M} \models \varphi_P$. □

Theorem 2. *Validity in ProbFO is Π_1^0 -hard (i.e., not recursively enumerable) even if only one free object variable is allowed to occur in weight formulas.*

Proof. We reduce from finite validity in FO. Let φ be an FO sentence, take a fresh unary predicate P and start with enforcing that for every domain element, the probability to satisfy P is $1/2^i$ for some $i \in \mathbb{N}$:

$$\forall x \forall r (w(P(x)) = r \Rightarrow \exists y (r = 1 \vee w(P(y)) = 2r)).$$

Next, guarantee that there are no infinite decreasing chains and thus only finitely many probabilities of satisfying P actually occur:

$$\exists r (r > 0 \wedge \forall y w(P(y)) \geq r).$$

Note that there can still be infinitely many elements with identical probabilities of satisfying P . We cannot prevent this, but we can force that ‘having the same probability of satisfying P ’ is a congruence regarding all relations that occur in φ . We only illustrate for a binary predicate R :

$$\begin{aligned} & \forall r_1 \forall r_2 \forall x_1 \forall x_2 \forall y_1 \forall y_2 \\ & \left((R(x_1, y_1) \wedge w(P(x_1)) = r_1 \wedge w(P(x_2)) = r_1 \wedge \right. \\ & \quad \left. w(P(y_1)) = r_2 \wedge w(P(y_2)) = r_2 \right) \Rightarrow R(x_2, y_2) \end{aligned}$$

Let ψ be the conjunction of the above ProbFO formulas. It is now easy to see that φ is finitely valid iff $(\psi \Rightarrow \varphi)$ is a ProbFO validity. □

Proposition 4. *The following formulas are not expressible in monodic ProbFO:*

1. $w(P(x, y)) \sim p$ for binary predicates P , $0 \leq p \leq 1$, and $\sim \in \{<, \leq, =, \geq, >\}$;
2. $w(A(x)) > w(A(y))$ for unary predicates A .

Proof. For point 1, assume that $\psi(x, y)$ is a formula expressing $w(P(x, y)) \sim p$. Now, consider the sentence $\varphi = \forall x (\varphi_1(x) \wedge \varphi_2(x) \wedge \varphi_3(x))$ with

$$\begin{aligned}\varphi_1(x) &= \neg A(x) \Rightarrow (w(A(x)) > 0) \\ \varphi_2(x) &= A(x) \Rightarrow \exists y (\psi(x, y) \wedge \neg A(y)) \\ \varphi_3(x) &= \neg A(x) \Rightarrow \forall y (\psi(x, y) \Rightarrow \neg A(y))\end{aligned}$$

which was used before for showing the lack of a finite-world property [20]. It is not hard to show that φ is only satisfiable in models comprising infinitely many worlds.

For point 2, let $\psi(x, y)$ be a formula expressing $w(A(x)) > w(A(y))$ and assume a model $\mathfrak{M} = (D, W, \mu, \pi)$ of $\varphi = \forall x \exists y \psi(x, y)$. By Theorem 7, we can assume that W is finite. Fix some $d_0 \in D$ and an arbitrary $w \in W$. As \mathfrak{M} is a model of φ , there is an infinite sequence d_0, d_1, \dots such that $w(A(d_i)) > w(A(d_{i+1}))$ for all $i \geq 0$. As W is finite, each $w(A(d_i))$ is a finite sum $\sum_{w \in W_i} \mu(w)$ for some $W_i \subseteq W$. However, this is a contradiction since there are only finitely many such subsets W_i . \square

Theorem 5. *The valid sentences of monodic ProbFO⁼ are Π_1^0 -hard.*

Proof. We again reduce finite validity in FO. Let φ be an FO sentence and take a fresh unary predicate P and a fresh constant symbol c . We enforce that for every domain element, the probability of satisfying P is at least as high as the probability that c satisfies P . Moreover, in every world, at most one element satisfies P :

$$\begin{aligned}w(P(c)) &> 0 \\ \forall x w(P(x)) &\geq w(P(c)) \\ w(\forall x \forall y ((P(x) \wedge P(y)) \rightarrow x = y)) &= 1\end{aligned}$$

It is easy to see that the above formulas guarantee that there are only finitely many domain elements: by the last formula, the probabilities of all elements to satisfy P sum up to at most one. By the first two lines, every probability in this sum is larger than some fixed positive real number r . \square

Proofs for Section 4

Theorem 7. *A monodic ProbFO sentence φ_0 is satisfiable iff it is satisfied in some quasi-model. Moreover, any satisfiable monodic ProbFO sentence is satisfied in a probabilistic structure with finitely many worlds.*

In the direction “ \Leftarrow ” of the proof for Theorem 7 we need the following lemma which can be proved based on the fact that we consider FO *without* equality.

Lemma 16. *If a world type T is realizable in a structure $\mathfrak{A} = (A, \pi)$ such that every type $t \in T$ is realized by κ_t elements in A , then for any family of cardinals $\lambda_t, t \in T$, with $\lambda_t \geq \kappa_t$ for all $t \in T$, there is a structure $\mathfrak{A}' = (A', \pi')$ such that every type $t \in T$ is realized by precisely λ_t elements in A' .*

Proof (of Theorem 7). “ \Rightarrow ”: Let $\mathfrak{M} = (D, W, \mu, \pi)$ be a probabilistic structure that satisfies φ_0 , that is $(\mathfrak{M}, w_0) \models \varphi_0$. We can w.l.o.g. assume that w_0 is the unique world with $\mu(w_0) = 0$: observe first, that all worlds with probability 0 (except for w_0) can be dropped without changing \mathfrak{M} being a model of φ_0 . If after this transformation $\mu(w_0) > 0$ we can add a world w'_0 which is essentially a copy of w_0 with probability $\mu(w'_0) = 0$.

We define a quasi-model (T_0, Q, R) satisfying φ_0 . For this purpose, we lift the definition of the functions $\text{tp}(\cdot)$ to probabilistic structures:

$$\begin{aligned} \text{tp}(\mathfrak{M}, d, w) &= \{\psi \in \text{sub}_x(\varphi_0) \mid \mathfrak{M}, w \models \psi(d)\} \\ \text{tp}(\mathfrak{M}, w) &= \{\text{tp}(\mathfrak{M}, d, w) \mid d \in D\}. \end{aligned}$$

Set $T_0 = \text{tp}(\mathfrak{M}, w_0)$ and $Q = \{\text{tp}(\mathfrak{M}, w) \mid w \in W, \mu(w) > 0\}$. Obviously, every $T \in Q \cup \{T_0\}$ is realizable. Next, define a set $R = \{r_d \mid d \in D\}$ where each function r_d is defined as

$$r_d(T) := \{\text{tp}(\mathfrak{M}, d, w) \mid w \in W, \text{tp}(\mathfrak{M}, w) = T\}.$$

Obviously, each function r_d is a run through Q . To show that $\mathcal{E}(Q, R)$ is positively satisfiable, choose for each $r \in R$ a domain element $d(r) \in D$ such that $r_{d(r)} = r$. Then define values x_T^* for every $T \in Q$ and $x_{r,t,T}^*$ for every $r \in R, T \in Q$, and $t \in r(T)$ by taking:

$$\begin{aligned} x_T^* &:= \mu(\chi(T)) \\ x_{r,t,T}^* &:= \mu(\chi(T) \wedge t(d(r))) \end{aligned}$$

It remains to check that the values $x_T^*, x_{r,t,T}^*$ present a positive solution to $\mathcal{E}(Q, R)$ from Definition 6. It should be clear that all values are positive. An important consequence of maximality of the types is that they are pairwise contradictory, i.e., for any two distinct types t_1, t_2 the formula $t_1(x) \wedge t_2(x)$ is unsatisfiable.

- The equation under Item 1 is clearly satisfied as

$$\sum_{T \in Q} x_T^* = \sum_{T \in Q} \mu(\chi(T)) = \mu\left(\bigvee_{T \in Q} \chi(T)\right) = 1$$

where the second equality holds as $\chi(T)$ and $\chi(T')$ are contradictory for any distinct $T, T' \in Q$ and the last equality holds as any world $w \in W$ satisfies $\chi(T)$ for some $T \in Q$.

- The equations under Item 2 are satisfied for every $r \in R, T \in Q$ since

$$\begin{aligned} \sum_{t \in r(T)} x_{r,t,T}^* &= \sum_{t \in r(T)} \mu(\chi(T) \wedge t(d(r))) \\ &= \sum_{t \in T} \mu(\chi(T) \wedge t(d(r))) \\ &= \mu(\chi(T) \wedge \bigvee_{t \in T} t(d(r))) \\ &= \mu(\chi(T)) = x_T^*. \end{aligned}$$

The second equality holds since $\mu(\chi(T) \wedge t(d(r))) = 0$ in case $t \notin r(T)$. The third equality holds as types are pairwise contradictory. The fourth equality holds as $\mathfrak{M} \models \chi(T) \Rightarrow \bigvee_{t \in T} t(d(r))$.

- For checking the equations under Item 3, let $f_1 \sim f_2 \in r$ for some $r \in R$. It suffices to show that $[f_i]_r = [f_i(d(r))]_{\mathfrak{M}, w}$ for all $w \in W$ and $i \in \{1, 2\}$. (Note that each f_i has at most one free variable, whose valuation we indicate by substituting $d(r)$ into the term f_i , then dispensing with further mention of a valuation). This is shown by induction on the structure of field terms f . The cases when f equals 0 , 1 , $f' + f''$, or $f' \times f''$ are clear. So it remains to consider the case $f = w(\psi(x))$. We have that $[w(\psi(d(r)))]_{\mathfrak{M}, w}$ equals

$$\begin{aligned}
& \mu(\{v \mid (\mathfrak{M}, v) \models \psi(d(r))\}) \\
&= \mu(\{v \mid (\mathfrak{M}, v) \models \bigvee_{T \in Q} \bigvee_{t \in T, \psi(x) \in t} (\chi(T) \wedge t(d(r)))\}) \\
&= \sum_{T \in Q} \sum_{t \in T, \psi(x) \in t} \mu(\chi(T) \wedge t(d(r))) \\
&= \sum_{T \in Q} \sum_{t \in r(T), \psi(x) \in t} \mu(\chi(T) \wedge t(d(r))) \\
&= \sum_{T \in Q} \sum_{t \in r(T), \psi(x) \in t} x_{r,t,T}^*
\end{aligned}$$

Note that the first equality holds as the big disjunction covers all possibilities. The second equality holds as all the $\chi(T)$ are pairwise contradictory and the types are pairwise contradictory. The third holds as $\mu(\chi(T) \wedge t(d(r))) = 0$ in case $t \notin r(T)$.

It remains to remark that by assumption $\varphi_0 \in t$ for all $t \in T_0$ and thus (T_0, Q, R) is a quasi-model satisfying φ_0 .

“ \Leftarrow ”: Let (T_0, Q, R) be a quasi-model satisfying φ_0 . Hence, $\varphi_0 \in t$ for some $t \in T_0$, every $T \in Q \cup \{T_0\}$ is realizable and $\mathcal{E}(Q, R)$ has a positive solution $x_T^*, T \in Q, x_{r,t,T}^*, r \in R, T \in Q, t \in r(T)$. Note that we can assume w.l.o.g. that $r(T_0)$ is a singleton set for each $r \in R$: As the types in T_0 do not contribute to the equation system $\mathcal{E}(Q, R)$, we can replace a run r with $r(T_0) = \{t_1, \dots, t_k\}$ by k runs r_1, \dots, r_k defined as $r_i(T_0) = \{t_i\}$ and $r_i(T) = r(T)$ for all $T \in Q$ and $i \in \{1, \dots, k\}$.

For constructing a probabilistic structure \mathfrak{M} that satisfies φ_0 , it would be convenient to use the world types in Q as worlds. Since runs can associate more than one type with a world type, though, this is not sufficient. We thus need to subdivide each $T \in Q$ into several worlds, each accommodating a type that a given run assigns to T . This has to be done in a careful way since we have to do this simultaneously for *all runs* while also ensuring that all types in T are realized in each of the worlds T is subdivided into.

Let $r \in R$ and $T \in Q$. A *subdivision of T for r* is a tuple $s = (b_1, \dots, b_n, \zeta)$ such that $b_1 < b_2 < \dots < b_n = x_T^*$, $n = |r(T)| + 1$, and ζ is a surjective function that assigns to every b_i a type $\zeta(b_i) \in r(T)$ such that for all $t \in r(T)$ we have $\sum_{i \in [1, n], \zeta(b_i) = t} (b_i - b_{i-1}) = x_{r,t,T}^*$ where, here and in what follows, $b_0 := 0$. Intuitively, the interval $[0, x_T^*]$ represents the probability covered by all worlds of type T

and we subdivide this range into the intervals $(b_i, b_{i+1}]$, with $i < n$. Elements described by the run r then have type $\zeta(b_{i+1})$ in the interval $(b_i, b_{i+1}]$. For easier reference, we say for $p \in (0, x_T^*]$ that s has type t at p if $\zeta(b_i) = t$ and $p \in (b_{i-1}, b_i]$.

A *subdivided run* is a pair (r, S) with $r \in R$ a run and S a function that assigns with every $T \in Q$ a subdivision $S(T)$ of T for r . We could use the subintervals identified by a subdivided run (r, S) as worlds if we had only the single run r . Since this is not the case, though, worlds are defined by combining a sufficiently rich set of subdivisions in an appropriate way. This set is identified by the next claim.

Claim 1. There is a finite set Γ of subdivided runs such that

- (*) for all $T \in Q$, $t \in T$, and $p \in (0, x_T^*]$, there is some $(r, S) \in \Gamma$ such that $S(T)$ has type t at p .

Proof of Claim 1. It suffices to show the statement for each $T \in Q$, $t \in T$. So fix some $T \in Q$ and $t \in T$. By definition of (Q, R) , there is some $r \in R$ with $t \in r(T)$ and, by assumption, $x_{r,t,T}^* > 0$. Obviously, there is a finite set of intervals $(y_1, z_1], \dots, (y_k, z_k]$ each of length $x_{r,t,T}^*$ which cover $(0, x_T^*]$. For each such interval $(y_i, z_i]$ we can find a subdivision $s_i = (b_1, \dots, b_n, \zeta)$ of T for r such that s_i has type t at p for all $p \in (y_i, z_i]$: for $n > r(T)$, we can always arrange the n intervals of a subdivision such that one interval $(b_{j-1}, b_j]$ is precisely $(y_i, z_i]$ and $\zeta(b_j) = t$.

Finally, note that for $T' \in Q$, $T' \neq T$ we can trivially fix a subdivision $s_{T'}$ of T' for r . Combining these with the subdivisions identified above yields a finite set of subdivided runs satisfying (*) for T, t . This finishes the proof of Claim 1.

Let Γ be the set of subdivided runs from Claim 1. For every $T \in Q$, let $\mathfrak{Z}(T)$ denote the set of all values b_i that occur in a subdivision $s = (b_1, \dots, b_n, \zeta)$ for T and r , for some $(r, S) \in \Gamma$ with $S(T) = s$. Further assume that $\mathfrak{Z}(T) = \{z_1^T, \dots, z_{m_T}^T\}$ with $z_1^T < \dots < z_{m_T}^T$. We now define a probabilistic structure $\mathfrak{M} = (D, W, \pi, \mu)$. First, define

$$\begin{aligned} W &= \{(T, z) \mid T \in Q, z \in \mathfrak{Z}(T)\} \cup \{(T_0, 0)\} \\ \mu(T, z_i^T) &= z_i^T - z_{i-1}^T \text{ for all } T \in Q \text{ and } 1 \leq i \leq m_T \\ &\text{(where } z_0^T := 0). \end{aligned}$$

Notice that every world (T, z) with $T \in Q$ has positive probability and the additional world $(T_0, 0)$ has probability 0. Next, we define the domain D of \mathfrak{M} . By assumption, every $T \in Q \cup \{T_0\}$ is realizable. Hence, for every world type $T \in Q \cup \{T_0\}$ we can fix an FO structure \mathfrak{A}_T that realizes it. Let $\mathfrak{A}_{T,z} = (A_{T,z}, \pi_{T,z})$, for $(T, z) \in W$, be pairwise disjoint copies of \mathfrak{A}_T . Then D is the disjoint union of all $A_{T,z}$, $(T, z) \in W$.

It remains to give the interpretation function π of \mathfrak{M} . For this purpose, we associate with every $d \in D$ a function σ_d that associates every world with a type as follows.

- If $d \in A_{T,z}$ for some $T \in Q$, then choose a $(r, S) \in \Gamma$ that has type $\text{tp}(\mathfrak{A}_{T,z}, d)$ at $S(T)$ (possible due to (*));
- if $d \in A_{T_0,0}$, choose some $(r, S) \in \Gamma$ with $r(T_0) = \{\text{tp}(\mathfrak{A}_{T_0,0}, d)\}$.

Then, define $\sigma_d(T', z')$ as the type of $S(T')$ at z' for all $(T', z') \in W$ and $\sigma_d(T_0, 0) = t_0$ where $\{t_0\} = r(T_0)$. Obviously, σ_d satisfies the following properties:

- If $d \in A_{T,z}$, then $\sigma_d(T, z) = \text{tp}(\mathfrak{A}_{T,z}, d)$;
- For all $(T, z) \in W$, we have $\sigma_d(T, z) \in T$.

These properties imply that in every world $(T, z) \in W$ the types t that occur (that is, for some d we have $\sigma_d(T, z) = t$) are precisely the types from T . More formally, for every $(T, z) \in W$ and $t \in T$ define κ_t as the cardinality of $\{d \in D \mid \sigma_d(T, z) = t\}$. By the first property above, we have for all $t \in T$:

$$\kappa_t \geq |\{d \in A_{T,z} \mid \text{tp}(\mathfrak{A}_{T,z}, d) = t\}|.$$

Thus, Lemma 16 implies the existence of models $\mathfrak{B}_{T,z} = (D, \pi'_{T,z})$ that realize T and such that for all $d \in D$ we have

$$\sigma_d(T, z) = \text{tp}(\mathfrak{B}_{T,z}, d).$$

To finish the definition of \mathfrak{M} , set for all predicate names P and all worlds $(T, z) \in W$:

$$\pi(P, (T, z)) = \pi'_{T,z}(P).$$

In order to prove correctness of the construction we show the following claim. There, it is convenient to write $\psi(x)$ even if ψ is a sentence, i.e., x does not occur freely in ψ .

Claim 2. For all valuations ν , all $\psi(x) \in \text{sub}_x(\varphi_0)$, and all $(T, z) \in W$ we have

$$(\mathfrak{B}_{T,z}, \nu) \models \overline{\psi(x)} \iff (\mathfrak{M}, (T, z), \nu) \models \psi(x).$$

Proof of Claim 2. The induction base, i.e., $\psi(x) = P(x_1, \dots, x_k)$, is clear since then $\psi(x) = \overline{\psi(x)}$ and, by construction, $\pi'_{T,z}(P) = \pi(P, (T, z))$. The induction steps for \wedge , \neg , $\forall x$ follow immediately from induction hypothesis.

Now let $\psi(x) = f_1 \leq f_2$ and thus $\overline{\psi(x)} = P_\psi(x)$. Assume $\nu(x) = d$ and $d \in A_{T^*, z^*}$ for some $(T^*, z^*) \in W$ and fix the subdivided run (r, S) used in the construction of σ_d . By definition of π , $(\mathfrak{B}_{T,z}, \nu) \models P_\psi(x)$ is equivalent to $\psi(x) \in \text{tp}(\mathfrak{B}_{T,z}, d)$. As $\psi(x)$ is a weight formula and r satisfies the coherence condition, this is equivalent to the fact that $\psi(x) \in r$. By Item 3 of Definition 6 and maximality of types, this is the case if, and only if, it holds $[f_1]_r \leq [f_2]_r$ (with the values $x_{r,t,T}^*$ from the fixed solution). By definition of subdivided runs and the construction of the set of worlds, this is equivalent to $[f_1]_{(\mathfrak{M}, w, \nu)} \leq [f_2]_{(\mathfrak{M}, w, \nu)}$ for any $w \in W$ (in particular, a subdivision for r – and thus a subdivided run – does not change the probability of some type t , it only gives some arrangement). Finally, this is equivalent to $(\mathfrak{M}, (T, z), \nu) \models f_1 \leq f_2$. This finishes the proof of Claim 2.

It remains to verify that $(\mathfrak{M}, (T_0, 0)) \models \varphi_0$. By definition of a quasi-model, we have $\varphi_0 \in t$ for some $t \in T_0$. By definition of \mathfrak{M} , there is some $d \in D$ with $\mathfrak{B}_{T_0,0} \models \overline{t(d)}$. Hence, $(\mathfrak{B}_{T_0,0}, \nu) \models \overline{\varphi_0}$ for any valuation ν . By the above claim, $(\mathfrak{M}, (T_0, 0), \nu) \models \varphi_0$, and thus $(\mathfrak{M}, (T_0, 0)) \models \varphi_0$. \square

Proofs for Section 5

In order to show completeness we first prove some auxiliary statements. We say that $\varphi_1, \dots, \varphi_m$ are pairwise mutually exclusive if $PC \vdash \overline{\varphi_i} \Rightarrow \neg \overline{\varphi_j}$ for $i \neq j$.

- Lemma 17.** (1) If $\varphi_1, \dots, \varphi_k$ are pairwise mutually exclusive, then $AX_2 \vdash w(\varphi_1 \vee \dots \vee \varphi_k) = w(\varphi_1) + \dots + w(\varphi_k)$.
 (2) If $AX_2 \vdash \varphi$, then $AX_2 \vdash w(\varphi) = 1$.
 (3) $AX_2 \vdash w(\varphi) + w(\neg\varphi) = 1$.
 (4) $AX_2 \vdash w(\varphi \wedge \psi) \leq w(\varphi)$.
 (5) $AX_2 \vdash w(\varphi \wedge \theta) > 0 \Rightarrow \theta$ provided that every predicate symbol in θ appears inside a weight term $w(\psi)$.

Proof. Claims (1)–(4) are shown in [12]. For claim (5), assume the contrary, i.e., $w(\varphi \wedge \theta) > 0 \wedge \neg\theta$ is consistent. By Axiom PW_1 , also $w(\varphi \wedge \theta) > 0 \wedge w(\neg\theta) = 1$ is consistent. By claim (3) of this Lemma, also $w(\varphi \wedge \theta) > 0 \wedge w(\theta) = 0$ is consistent. By claim (4), we get that $w(\theta) > 0 \wedge w(\theta) = 0$ is consistent, contradiction. \square

Lemma 18. If a monodic ProbFO sentence φ_0 is consistent, then there is a quasi-model satisfying φ_0 .

Proof. For a type t , let us denote with $wf(t)$ the set of all weight formulas or their negations contained in t . Moreover, we say that two types t, t' agree on weight formulas if $wf(t) = wf(t')$. Denote with \widehat{W} (resp., W) be the set of all (resp., all realizable) world types for φ_0 . We begin with showing the following claim.

Claim 1. There is a realizable world type T_0 and a type $t \in T_0$ with $\varphi_0 \in t$ and a set Q of realizable world types such that

- (i) the formula

$$\vartheta = \chi(T_0) \wedge \left(\sum_{T \in Q} w(\chi(T)) = 1 \right) \wedge \left(\bigwedge_{T \in Q} w(\chi(T)) > 0 \right)$$

is consistent;

- (ii) for each $T \in Q, t \in T$ there is a $t_0 \in T$ agreeing on the weight formulas with t .

Proof of Claim 1. We begin by noting that, if φ_0 is consistent, then also $\chi(T_0)$ is consistent for some realizable world type T_0 containing some $t \in T_0$ with $\varphi_0 \in t$. Next, observe that $PC \vdash \bigvee_{T \in \widehat{W}} \chi(T)$ and for any not realizable world type $T' \in \widehat{W}$ we have $PC \vdash \neg\chi(T')$. Thus, we obtain $PC \vdash \bigvee_{T \in W} \chi(T)$. Lemma 17(2) now yields $AX_2 \vdash w(\bigvee_{T \in W} \chi(T)) = 1$. As for distinct world types T, T' we know that $\chi(T)$ and $\chi(T')$ are mutually exclusive, Lemma 17(1) implies $AX_2 \vdash \sum_{T \in W} w(\chi(T)) = 1$. Hence, the formula $\chi(T_0) \wedge \sum_{T \in W} w(\chi(T)) = 1$ is consistent. Reasoning by cases (in *OF*) yields that we can identify a subset $Q \subseteq W$ such that ϑ is consistent; thus, we are finished with Item (i).

For Item (ii) assume some $T^* \in Q, t^* \in T$ that does not agree on the weight formulas with any type from T_0 . Observe that $AX_2 \vdash \chi(T_0) \Rightarrow \forall x \bigvee_{t \in T_0} wf(t(x))$.

On the other hand, we have that $AX_2 \vdash \vartheta \Rightarrow w(\chi(T^*)) > 0$ and thus $AX_2 \vdash \vartheta \Rightarrow w(\exists x t^*(x)) > 0$ by Lemma 17(4). By PW_4 , we obtain $AX_2 \vdash \vartheta \Rightarrow \exists x w(t^*(x)) > 0$. Applying Lemma 17(5) yields $AX_2 \vdash \vartheta \Rightarrow \exists x wf(t^*(x))$. Overall, consistency of ϑ yields consistency of $\forall x \bigvee_{t \in T_0} wf(t(x)) \wedge \exists x wf(t^*(x))$ which is a contradiction by PC , the assumption that t^* does not agree on the weight formulas with any $t \in T_0$, and maximality of types. This finishes the proof of Claim 1.

Subsequently, we need some auxiliary formulas that can be derived in AX_2 .

Claim 2. The following formulas can be derived:

$$AX_2 \vdash \vartheta \Rightarrow w(\psi(x)) = \sum_{S \in Q} \sum_{s \in S, \psi(x) \in s} w(\chi(S) \wedge s(x)); \quad (13)$$

$$AX_2 \vdash w(\chi(S)) = \sum_{s \in S} w(\chi(S) \wedge s(x)) \quad \text{for all } S \in W. \quad (14)$$

Proof of Claim 2. For deriving (13), note that we have $PC \vdash \psi(x) \equiv \bigvee_{S \in \widehat{W}} \bigvee_{s \in S, \psi(x) \in s} (\chi(S) \wedge s(x))$; thus RPW and Lemma 17(1) lead to

$$AX_2 \vdash w(\psi(x)) = \sum_{S \in \widehat{W}} \sum_{s \in S, \psi(x) \in s} w(\chi(S) \wedge s(x)).$$

It remains to note that that $AX_2 \vdash \vartheta \Rightarrow w(\chi(S)) = 0$ for all $S \notin Q$; by Lemma 17(4), also $AX_2 \vdash \vartheta \Rightarrow w(\chi(S) \wedge s(x)) = 0$ for all $s \in S$.

For (14), observe that we have for all world types $S \in W$:

$$PC \vdash \chi(S) \equiv \chi(S) \wedge \bigvee_{s \in S} s(x).$$

Hence, $AX_2 \vdash w(\chi(S)) = w(\chi(S) \wedge \bigvee_{s \in S} s(x))$, and Lemma 17(1) leads to the result. This finishes the proof of Claim 2.

The next step is to show how consistency of ϑ induces a run through every type $t \in T$, $T \in Q$. For this purpose, fix some $T \in Q$, a type $t \in T$, and the set $wf(t) = \{\theta_1, \dots, \theta_m\}$ of weight formulas (or their negations) in t . Obviously, we have $AX_2 \vdash \vartheta \Rightarrow w(\chi(T)) > 0$. As $\chi(T)$ contains the conjunct $\exists x t(x)$, we can derive $\vartheta \Rightarrow \exists x w(\chi(T) \wedge t(x)) > 0$ using PW_4 . By applying Lemma 17(5) to $\theta_1, \dots, \theta_m$ we can also derive

$$AX_2 \vdash \vartheta \Rightarrow \exists x w(\chi(T) \wedge t(x)) > 0 \wedge \theta_1(x) \wedge \dots \wedge \theta_m(x) \quad (15)$$

By formula (13), we can replace each $\theta_j(x)$ in (15) by the formula $\theta'_j(x)$ that is obtained by substituting every $w(\psi(x))$ with $\sum_{S \in Q} \sum_{s \in S, \psi(x) \in s} w(\chi(S) \wedge s(x))$. Moreover, we can add formula (14) for every $S \in Q$ as conjunct. We thus obtain the following:

$$AX_2 \vdash \vartheta \Rightarrow \exists x w(\chi(T) \wedge t(x)) > 0 \wedge \theta'_1(x) \wedge \dots \wedge \theta'_m(x) \wedge \bigwedge_{S \in Q} w(\chi(S)) = \sum_{s \in S} w(\chi(S) \wedge s(x))$$

Recall that $AX_2 \vdash \vartheta \Rightarrow w(\chi(S)) > 0$ for each $S \in Q$. Thus, in the sum in the last line, $w(\chi(S) \wedge s(x))$ has to be positive for at least one $s \in S$. Using reasoning by cases, there exists a choice function selecting the set of all such s for each S . More precisely, there is some function r such that $t \in r(T)$ (due to the conjunct $w(\chi(T) \wedge t(x)) > 0$), $r(S) \neq \emptyset$ for all $S \in Q$, and the following is consistent:

$$\vartheta \wedge \exists x \left(\bigwedge_{j=1}^m \theta'_j(x) \wedge \bigwedge_{S \in Q, s \in r(S)} w(\chi(S) \wedge s(x)) > 0 \right. \\ \left. \bigwedge_{S \in Q} w(\chi(S)) = \sum_{s \in r(S)} w(\chi(S) \wedge s(x)) \right)$$

Observe now that all types $s \in r(S)$, $S \in Q$ actually agree on the weight formulas; in particular, for all such s we have that $wf(s) = wf(t)$. Assume to the contrary, $\neg\theta_j(x) \in s$ for some $s \in r(S)$. By Lemma 17(5), we have $AX_2 \vdash w(\psi(S) \wedge s(x)) > 0 \Rightarrow \neg\theta_j(x)$ and thus $AX_2 \vdash w(\psi(S) \wedge s(x)) > 0 \Rightarrow \neg\theta'_j(x)$, which contradicts consistency of the above. By Item (ii) of Claim 1, r can be extended to be a run through $Q \cup \{T_0\}$ by setting $r(T_0) = \{t_0\}$ for some $t_0 \in T_0$ that agrees with t on weight formulas.

Repeating the above steps for each $T \in Q$, $t \in T$ leaves us with a run $r_{t,T}$ for each such pair t, T . Take $R = \{r_{t,T} \mid t \in T, T \in Q\}$. We also obtain a consistent formula of the form

$$\bigwedge_{T \in Q} w(\chi(T)) > 0 \wedge \sum_{T \in Q} w(\chi(T)) = 1 \wedge \\ \bigwedge_{r \in R} \left(\bigwedge_{\theta \in r} \theta'(x_r) \wedge \bigwedge_{S \in Q, s \in r(S)} w(\chi(S) \wedge s(x_r)) > 0 \wedge \right. \\ \left. \bigwedge_{S \in Q} w(\chi(S)) = \sum_{s \in r(S)} w(\chi(S) \wedge s(x_r)) \right) \quad (16)$$

where $\chi(T_0)$ and existential quantification of the variables $x_r, r \in R$ is omitted (can be omitted as we study consistency). Let ϑ' be obtained from (16) by replacing each $w(\chi(T))$ with x_T and each $w(\chi(T) \wedge t(x))$ in some conjunct for $r \in R$ with $x_{r,t,T}$. Note that ϑ' is a formula in the language of ordered fields; in fact, ϑ' is equivalent to $\mathcal{E}(Q, R)$ when we explicitly add positivity of the solution to $\mathcal{E}(Q, R)$. Consistency of (16) implies that ϑ' is consistent in the theory of ordered fields (otherwise, being quantifier free, its negation could, by Herbrand's theorem, be proved using propositional reasoning with instances of the ordered field axioms, and this proof would translate into one of the negation of (16) using only the instances included in our axioms OF). Therefore, ϑ' is satisfiable in some ordered field F , and hence by the Artin-Schreier theorem in some real-closed field (the real-closure of F), thus also in the real numbers. This implies that $\mathcal{E}(Q, R)$ has a positive solution over the reals, i.e. (T_0, Q, R) is a quasi-model satisfying φ_0 . \square

Proofs for Section 6

Theorem 14. *If \mathcal{L} is closed under union of types, then for every satisfiable monodic Prob \mathcal{L} sentence φ_0 , there is a quasi-model (T_0, Q, R) satisfying φ_0 such that any two distinct world types do not agree on sentences.*

Proof. Let φ_0 be a satisfiable monodic Prob \mathcal{L} sentence. By Theorem 7, there is a quasi-model (T_0, Q, R) satisfying φ_0 . Let us write $T \equiv_0 T'$ if two sets of types T, T' agree on sentences, that is, any for any $t \in T, t' \in T'$ we have that $t \equiv_0 t'$. Define a new quasi-model (T_0, Q', R') as follows:

- $Q' = \{\bigcup_{T' \in Q, T \equiv_0 T'} T' \mid T \in Q\}$;
- for each $r \in R$ define $r' \in R'$ by taking $r'(T_0) = r(T_0)$ and for each $T' \in Q'$:

$$r'(T') = \bigcup_{T \in Q, T \equiv_0 T'} r(T).$$

It should be clear that R' is a set of runs through $Q' \cup \{T_0\}$ and that (T_0, Q', R') is a quasi-model candidate.

Observe first that each $T \in Q'$ is realizable as it is a finite union of realizable $T' \in Q$ and, by assumption, \mathcal{L} is closed under union of types. Next, we show that $\mathcal{E}(Q', R')$ has a positive solution. By assumption, we can fix a solution x_T^* for every $T \in Q$ and $x_{r,t,T}^*$ for every $r \in R, T \in Q$, and $t \in r(T)$ of $\mathcal{E}(Q, R)$. For each $r' \in R'$, fix some (arbitrary) run $\bar{r} \in R$ such that $\bar{r}' = r'$ (in general, there might be different runs $r_1, r_2 \in R$ such that $r_1' = r_2'$). Now, define values $y_{T'}^*$ and $y_{r',t,T'}^*$:

$$\begin{aligned} y_{T'}^* &:= \sum_{T \in Q, T \equiv_0 T'} x_T^* \\ y_{r',t,T'}^* &:= \sum_{T \in Q, T \equiv_0 T'} \sum_{t \in \bar{r}(T)} x_{\bar{r},t,T}^*. \end{aligned}$$

We show that these values are a positive solution of $\mathcal{E}(Q', R')$.

- The equation under Item 1 of Definition 6 is satisfied as \equiv_0 partitions Q and it was satisfied in $\mathcal{E}(Q, R)$.
- For the equations under Item 2 we have:

$$\begin{aligned} \sum_{t \in r'(T')} y_{r',t,T'}^* &= \sum_{t \in r'(T')} \sum_{T \in Q, T \equiv_0 T'} \sum_{t \in \bar{r}(T)} x_{\bar{r},t,T}^* \\ &= \sum_{T \in Q, T \equiv_0 T'} x_T^* = y_{T'}^*. \end{aligned}$$

– For seeing that the equations under Item 3 are satisfied it suffices to note that:

$$\begin{aligned}
[\mathbf{w}(\chi(x))]_{r'} &= \sum_{T' \in Q'} \sum_{t \in r'(T'), \psi(x) \in t} y_{r', t, T'}^* \\
&= \sum_{T' \in Q'} \sum_{t \in r'(T'), \psi(x) \in t} \sum_{T \in Q, T \equiv_0 T'} \sum_{t \in \bar{r}(T)} x_{\bar{r}, t, T}^* \\
&= \sum_{T \in Q} \sum_{t \in \bar{r}(T), \psi(x) \in t} x_{\bar{r}, t, T}^* \\
&= [\mathbf{w}(\psi(x))]_{\bar{r}}
\end{aligned}$$

The third equality is the most subtle one: observe that the sums in the second line range over all $T \in Q$ and do not count any $T \in Q$ twice. Moreover, only types t with $\psi(x) \in t$ are considered. \square