Asymptotic Conditional Probabilities: The Non-unary Case*

Adam J. Grove
NEC Research Institute
4 Independence Way
Princeton, NJ 08540
email: grove@research.nj.nec.com

Joseph Y. Halpern
IBM Almaden Research Center
650 Harry Rd.
San Jose, CA 95120
email: halpern@almaden.ibm.com

Daphne Koller
Department of Computer Science
Stanford University
Stanford, CA 94305
email: daphne@cs.stanford.edu

Abstract

Motivated by problems that arise in computing *degrees of belief*, we consider the problem of computing asymptotic conditional probabilities for first-order sentences. Given first-order sentences φ and θ, we consider the structures with domain $\{1, \ldots, N\}$ that satisfy θ, and compute the fraction of them in which φ is true. We then consider what happens to this fraction as N gets large. This extends the work on 0-1 laws that considers the limiting probability of first-order sentences, by considering asymptotic *conditional* probabilities. As shown by Liogon'kiĭ [Lio69], if there is a non-unary predicate symbol in the vocabulary, asymptotic conditional probabilities do not always exist. We extend this result to show that asymptotic conditional probabilities do not always exist for any reasonable notion of limit. Liogon'kiĭ also showed that the problem of deciding whether the limit exists is undecidable. We analyze the complexity of three problems with respect to this limit: deciding whether it is well defined, whether it exists, and whether it lies in some nontrivial interval. Matching upper and lower bounds are given for all three problems, showing them to be highly undecidable.

*Some of this research was done while all three authors were at IBM Almaden Research Center. During this work, the first author was at Stanford University, and was supported by an IBM Graduate Fellowship. This research was sponsored in part by the Air Force Office of Scientific Research (AFSC), under Contract F49620-91-C-0080. The United States Government is authorized to reproduce and distribute reprints for governmental purposes. Some of this research appeared in preliminary form in a paper entitled "Asymptotic conditional probabilities for first-order logic", which appears in *Proceedings 24th ACM Symp. on Theory of Computing*, 1992, pages 294–305. This paper appears in the *Journal of Symbolic Logic*.

1 Introduction

Suppose we have a sentence θ expressing facts that are known to be true, and another sentence φ whose truth is uncertain. Our knowledge θ is often insufficient to determine the truth of φ: both φ and its negation may be consistent with θ. In such cases, it can be useful to assign a *probability* to φ, based on θ. One important application of this idea—indeed, the one that has provided most of our motivation—is in the domain of decision theory and artificial intelligence. Consider an agent (or expert system) whose knowledge consists of some facts θ, who would like to assign a *degree of belief* to a particular statement φ. For example, a doctor may want to assign a degree of belief to the hypothesis that a patient has a particular illness, based on the symptoms exhibited by the patient together with general information about symptoms and diseases. Since the actions the agent takes may depend crucially on this value, we would like techniques for computing degrees of belief in a principled manner.

The difficulty of defining a principled technique for computing the probability of φ given θ, and then actually computing that probability, depends in part on the language and logic being considered. In decision theory, applications often demand the ability to express statistical knowledge (for instance, correlations between symptoms and diseases) as well as first-order knowledge. Work in the field of *0-1 laws* (which, as discussed below, is closely related to our own) has examined some higher-order logics as well as first-order logic. Nevertheless, the pure first-order case is still difficult, and is important because it provides a foundation for all extensions. In this paper and in [GHK93] we address the problem of computing conditional probabilities in the first-order case. In a companion paper [GHK94], we consider the case of statistical knowledge.

The general problem of assigning probabilities to first-order sentences has been well studied (cf. [Gai60, Gai64]). In this paper, we investigate two specific formalisms for computing probabilities, based on the same basic approach. The approach is based on an old idea, that goes back to Laplace [Lap20]. It is essentially an application of what has been called *the principle of insufficient reason* [Kri86] or *the principle of indifference* [Key21]. The idea is to assign equal degree of belief to all basic "situations" consistent with the known facts. The two formalisms we consider differ only in how they interpret "situation". We discuss this in more detail below.

In many applications, including the one of most interest to us, it makes sense to consider finite domains only. In the first-order case, our approach essentially generalizes the methods used in the work on 0-1 laws to the case of conditional probabilities. (See Compton's overview [Com88] for an introduction to this work.) Assume, without loss of generality, that the domain is $\{1,\ldots,N\}$ for some natural number N. As we said above, we consider two notions of "situation". In the *random-worlds method*, the possible situations are all the worlds, or first-order models, with domain $\{1,\ldots,N\}$ that satisfy the constraints θ. Based on the principle of indifference, we assume that all worlds are equally likely. In order to assign a probability to φ, we therefore simply compute the fraction of them in which the sentence φ is true. The random-worlds approach views each individual in $\{1,\ldots,N\}$ as having a distinct name (even though the name may not correspond to any constant in the vocabulary). Thus, two worlds that are isomorphic with respect to the symbols in the vocabulary are still treated as distinct situations. In some cases, however, we may believe that all relevant distinctions are captured by our vocabulary, and that isomorphic worlds are not truly distinct. The *random-structures* *method* attempts to capture this intuition by considering a situation to be a *structure*—an *isomorphism class* of worlds. This corresponds to assuming that individuals are distinguishable only if they differ with respect to properties definable by the language. As before, we assign a probability to φ by computing the fraction of the structures that satisfy φ among those structures that satisfy θ.[1]

Since we are computing probabilities over finite models, we have assumed that the domain is $\{1, \ldots, N\}$ for some N. However, we often do not know the precise domain size N. In many cases, we know only that N is large. We therefore estimate the probability of φ given θ by the asymptotic limit, as N grows to infinity, of this probability over models of size N.

Precisely the same definitions of asymptotic probability are used in the context of 0-1 laws for first-order logic, but without allowing arbitrary prior information θ. The original 0-1 law, proved independently by Glebskiĭ et al. [GKLT69] and Fagin [Fag76], states that the asymptotic probability of any first-order sentence φ with no constant or function symbols is either 0 or 1. Intuitively, such a sentence is true in almost all finite structures, or in almost none. Interestingly, this 0-1 law holds under both the random-worlds and the random-structures methods; in fact, both approaches lead to the same asymptotic probability in this case [Fag77].

Our work differs from the original work on 0-1 laws in two respects. The first is relatively minor: we need to allow the use of constant symbols in φ, as they are necessary when discussing individuals (such as patients). Although this is a minor change, it is worth observing that it has a significant impact: It is easy to see that once we allow constant symbols, the asymptotic probability of a sentence φ is no longer either 0 or 1; for example, the asymptotic probability of $P(c)$ is $\frac{1}{2}$. Moreover, once we allow constant symbols, the asymptotic probability under random worlds and under random structures need not be the same. The more significant difference, however, is that we are interested in the asymptotic *conditional* probability of φ, given some prior knowledge θ. That is, we want the probability of φ over the class of finite structures defined by θ.

Some work has already been done on aspects of this question. Fagin [Fag76] and Liogon'kiĭ [Lio69] independently showed that asymptotic conditional probabilities do not necessarily converge to any limit. Subsequently, 0-1 laws were proved for special classes of first-order structures (such as graphs, tournaments, partial orders, etc.; see the overview paper [Com88] for details and further references). In many cases, the classes considered could be defined in terms of first-order constraints. Thus, these results can be viewed as special cases of the problem that we are interested in: computing asymptotic *conditional* probabilities relative to structures satisfying the constraints of a knowledge base. Lynch [Lyn85] showed that, for the random-worlds method, asymptotic probabilities exist for first-order sentences involving unary functions, although there is no 0-1 law. (Recall that the original 0-1 result is specifically for first-order logic *without* function symbols.) This can also be viewed as a special case of an asymptotic conditional probability for first-order logic without functions, since we can replace the unary functions by binary predicates, and condition on the fact that they are functions.

[1]The random-worlds method considers what has been called in the literature *labeled* structures, while the random-structures method considers *unlabeled* structures [Com88]. We choose to use our own terminology for the random-worlds and random-structures methods, rather than the terminology of labeled and unlabeled. This is partly because we feel it is more descriptive, and partly because there are other variants of the approach, that are useful for our intended application, and do not fit into the standard labeled/unlabeled structures dichotomy (see [BGHK95]).

The most comprehensive work on this problem is the work of Liogon'kiĭ [Lio69].[2] In addition to pointing out that asymptotic conditional probabilities do not exist in general, he shows that it is undecidable whether such a probability exists. He then investigates the special case of conditioning on formulas involving unary predicates only (but no equality). In this case, he proves that the asymptotic conditional probability does exist and can be effectively computed, even if the left side of the conditional has predicates of arbitrary arity and equality.

We extend the results of [Lio69] in a number of ways. We first show, in Section 3, that under any standard weakening of the concept of limit, asymptotic conditional probabilities still do not exist. We define three independent questions related to the asymptotic conditional probability: deciding whether it is well defined (i.e., is there an infinite sequence of probabilities to take the limit over); deciding whether it exists, given that it is well defined; and computing or approximating it, given that it exists. We show in Section 4 that all three problems are undecidable, and precisely characterize the degree of their undecidability. These results continue to hold for many quite restrictive sublanguages of first-order logic. We then present one "positive" result: In perhaps the most restrictive sublanguage that is still of any interest, if there is a fixed, finite vocabulary, and the quantifier depths of φ and θ are bounded, there is a linear time algorithm that computes the asymptotic conditional probability of φ given θ. Moreover, for each fixed vocabulary and fixed bound on quantifier depth, we can construct a finite set of algorithms, one of which is guaranteed to be one that solves the problem. However, it follows from our undecidability results that we cannot tell which algorithm is the correct one. So even this result holds no real promise. In a companion paper [GHK93], we extend Liogon'kiĭ's results for the case of conditioning on unary formulas. This special case turns out to be quite important for our application; see [GHK94, BGHK95, BGHK94].

Our undecidability results are of more than purely technical interest. The random-worlds method is of considerable theoretical and practical importance. We have already mentioned its relevance to computing degrees of belief. There are well-known results from physics that show the close connection between the random-worlds method and *maximum entropy* [Jay78]; some formalization of similar results, but in a framework that is closer to that of the current paper, can be found in [PV89, GHK94]. Essentially, the results say that in certain cases the asymptotic probability can be computed using maximum entropy methods.[3]

Given the wide use of maximum entropy, and its justification in terms of the random-worlds method, our results showing that it is not as widely applicable as one might hope come as somewhat of a surprise. Indeed, the difficulties of using the method once we move to non-unary predicates seem not to have been fully appreciated. In retrospect, this is not that hard to explain; in almost all applications where maximum entropy has been used (and where its application can be best justified in terms of the random-worlds method) the database is described in terms of unary predicates (or, equivalently, unary functions with a finite range). For example, in physics applications we are interested in such predicates as quantum state (see [DD85]). Similarly, AI applications and expert systems [Che83] typically use only unary predicates such as symptoms and diseases.

[2]In an earlier version of this paper [GHK92], we stated that, to our knowledge, no work had been done on the general problem of asymptotic conditional probabilities. We thank Moshe Vardi for pointing out Liogon'kiĭ's work to us.

[3]These results are of far more interest when there are statistical assertions in the language, so we do not discuss them here (see [PV89, GHK94] for more details).

It is interesting to note that in [Car52], where Carnap considers a continuum of methods for inductive reasoning (which includes the random-worlds method and a variant of the random-structures method), he considers only the unary case for all of them, without any comment or justification. He does provide some justification in [Car50], as well as expressing concern that the case of non-unary predicates may cause difficulties (although he presents no technical justification for this claim):

> ... the bulk of our inductive logic will deal only with properties of individuals [i.e., unary predicates], not with relations between individuals, except for those relations which are defined on the basis of properties. At the present time, this restriction seems natural and well justified, in view of the fact that deductive logic took more than two thousand years from its start with Aristotle to the first logic of relations (De Morgan, 1860). Inductive logic ... is only a few hundred years old. Therefore, it is not surprising to see that so far nobody has made an attempt to apply it to relations. ... The inclusion of relations in deductive logic causes obviously a certain increase in complexity. The corresponding increase in complexity for inductive logic is very much greater.

Carnap's allusion to the difficulty of adding relations to deductive logic is perhaps the observation—known at the time—that while first-order logic over a vocabulary with only unary predicate symbols is decidable, it becomes undecidable when we add non-unary predicates [DG79, Lew79]. The fact that there is an increase in complexity in inductive logic when we add non-unary predicates is not substantiated by Carnap, other than by the observation that very difficult combinatorial questions arise. As our results show, Carnap's concern about the difficulty of doing inductive reasoning with relations (non-unary predicates) is well founded.

2 Asymptotic conditional probabilities

Let Φ be a set of of predicate and function symbols, and let $\mathcal{L}(\Phi)$ (resp., $\mathcal{L}^-(\Phi)$) denote the set of first-order sentences over Φ with equality (resp., without equality). For the purpose of this paper, we assume that Φ is finite. We discuss the possibility of an infinite vocabulary in [GHK93]; it turns out that there are various ways to extend the relevant concepts to the infinite case, but that, for generally trivial reasons, all the results in this paper hold under any of these definitions. Therefore it is sufficient to work with the simplifying assumption of finiteness.

2.1 The random-worlds method

We begin by defining the random-worlds, or labeled, method. Given a sentence $\xi \in \mathcal{L}(\Phi)$, let $\#world_N^\Phi(\xi)$ be the number of worlds (first-order models) over Φ with domain $\{1, \ldots, N\}$ in which ξ is true. Since Φ is finite, so is $\#world_N^\Phi(\xi)$. We define

$$\Pr_N^{w,\Phi}(\varphi|\theta) = \frac{\#world_N^\Phi(\varphi \wedge \theta)}{\#world_N^\Phi(\theta)} .$$

At first glance, it seems that the value of $\Pr_N^{w,\Phi}(\varphi|\theta)$ depends on the choice of Φ. The following proposition shows that this is not the case.

Proposition 2.1: *Let Φ, Φ' be finite vocabularies, and let φ, θ be sentences in both $\mathcal{L}(\Phi)$ and $\mathcal{L}(\Phi')$. Then $\Pr_N^{w,\Phi}(\varphi|\theta) = \Pr_N^{w,\Phi'}(\varphi|\theta)$.*

Proof: We first prove the claim for the case $\Phi' = \Phi \cup \{R\}$ for some symbol $R \notin \Phi$. Let $\xi \in \mathcal{L}(\Phi)$ be an arbitrary sentence. A world over Φ' determines the denotations of the symbols in Φ, and the denotation of R. Let r be the number of possible denotations of R over a domain of size N. Since ξ does not mention R, it is easy to see that each model of ξ over Φ corresponds to r models of ξ over Φ', one for each possible denotation of R. Therefore, $\#\mathit{world}_N^{\Phi'}(\xi) = r \cdot \#\mathit{world}_N^{\Phi}(\xi)$. From this, we can deduce that

$$\Pr_N^{w,\Phi'}(\varphi|\theta) = \frac{\#\mathit{world}_N^{\Phi'}(\varphi \wedge \theta)}{\#\mathit{world}_N^{\Phi'}(\theta)} = \frac{r \cdot \#\mathit{world}_N^{\Phi}(\varphi \wedge \theta)}{r \cdot \#\mathit{world}_N^{\Phi}(\theta)} = \frac{\#\mathit{world}_N^{\Phi}(\varphi \wedge \theta)}{\#\mathit{world}_N^{\Phi}(\theta)} = \Pr_N^{w,\Phi}(\varphi|\theta) \ ,$$

as required.

Now, given arbitrary Φ and Φ', a straightforward induction on the cardinality of $\Phi' - \Phi$ shows that $\Pr_N^{w,\Phi\cup\Phi'}(\varphi|\theta) = \Pr_N^{w,\Phi}(\varphi|\theta)$. Similarly, we can show that $\Pr_N^{w,\Phi\cup\Phi'}(\varphi|\theta) = \Pr_N^{w,\Phi'}(\varphi|\theta)$. The result now follows. ∎

Based on this proposition, we omit reference to Φ in $\Pr_N^{w,\Phi}(\varphi|\theta)$, writing $\Pr_N^{w}(\varphi|\theta)$ instead.

We would like to define $\Pr_\infty^{w}(\varphi|\theta)$ as the limit $\lim_{N\to\infty} \Pr_N^{w}(\varphi|\theta)$. However, we must first deal with a technical problem in this definition: we must decide what to do if $\#\mathit{world}_N^{\Phi}(\theta) = 0$, so that $\Pr_N^{w}(\varphi|\theta)$ is not well defined. Liogon'kiĭ simply takes $\Pr_N^{w}(\varphi|\theta) = 1/2$ in this case; we take a somewhat more refined approach here.

It might seem reasonable to say that the asymptotic probability is not well defined if $\#\mathit{world}_N^{\Phi}(\theta) = 0$ for infinitely many N. However, suppose that θ is a sentence that is satisfiable only when N is even and, for even N, $\varphi \wedge \theta$ holds in one third of the models of θ. In this case, we might want to say that there is an asymptotic conditional probability of 1/3, even though $\#\mathit{world}_N^{\Phi}(\theta) = 0$ for infinitely many N. Thus, we actually consider two notions: the persistent limit, denoted $\Diamond\Box\Pr_\infty^{w}(\varphi|\theta)$, and the intermittent limit, denoted $\Box\Diamond\Pr_\infty^{w}(\varphi|\theta)$ (the prefixes stand for the temporal logic representation of the persistence and intermittence properties [MP92]). In either case, we say that the limiting probability is either not well defined, does not exist, or is some number between 0 or 1. The only difference between the two notions lies in when the limiting probability is taken to be well defined. This difference is made precise in the following definition.

Definition 2.2: Let $\mathcal{N}(\theta)$ denote the set $\{N : \#\mathit{world}_N^{\Phi}(\theta) \neq 0\}$. The asymptotic conditional probability $\Diamond\Box\Pr_\infty^{w}(\varphi|\theta)$ is *well defined* if $\mathcal{N}(\theta)$ contains all but finitely many N's; $\Box\Diamond\Pr_\infty^{w}(\varphi|\theta)$ is *well defined* if $\mathcal{N}(\theta)$ is infinite. If the asymptotic probability $\Diamond\Box\Pr_\infty^{w}(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^{w}(\varphi|\theta)$) is well defined, then we take $\Diamond\Box\Pr_\infty^{w}(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^{w}(\varphi|\theta)$) to denote $\lim_{N\to\infty, N\in\mathcal{N}(\theta)} \Pr_N^{w}(\varphi|\theta)$. ∎

Remark 2.3:

(a) If $\Diamond\Box\Pr_\infty^{w}(\varphi|\theta)$ is well defined, then so is $\Box\Diamond\Pr_\infty^{w}(\varphi|\theta)$. The converse is not necessarily true.

(b) If both $\Diamond\Box\Pr^w_\infty(\varphi|\theta)$ and $\Box\Diamond\Pr^w_\infty(\varphi|\theta)$ are well defined, then they are equal.[4] ∎

It follows from our results in [GHK93] that the two notions of limiting probability coincide if we restrict to unary predicates or to languages without equality.

2.2 The random-structures method

One way of thinking about the random-worlds method is that it treats each individual in $\{1, \ldots, N\}$ as having a distinct name or label (even though the name may not actually correspond to any constant in the vocabulary). This intuition explains why two worlds that are completely isomorphic as first-order structures (i.e., with respect to the symbols in the vocabulary) are nevertheless regarded as distinct worlds and are counted separately. The *random-structures* method, on the other hand, only counts the number of *(unlabeled) structures*, which we identify with isomorphism classes of worlds. Formally, we say that two worlds $\mathcal{W}_1$ and $\mathcal{W}_2$ are *isomorphic* (with respect to the language $\mathcal{L}$) if there is a bijective mapping f from the domain of $\mathcal{W}_1$ to the domain of $\mathcal{W}_2$, such that for every predicate symbol $P \in \mathcal{L}$, we have $P^{\mathcal{W}_1}(d_1, \ldots, d_k)$ iff $P^{\mathcal{W}_2}(f(d_1), \ldots, f(d_k))$, and similarly for the function and constant symbols in $\mathcal{L}$. Intuitively, this says that worlds that treat the symbols in the language in the same way are "really" the same, and so should only be counted once.

Given a structure $\mathcal{S}$ and a sentence $\xi \in \mathcal{L}(\Phi)$, all the worlds in $\mathcal{S}$ agree on the truth value they assign to ξ. Therefore, we can say that $\mathcal{S}$ satisfies (or does not satisfy) ξ. Let $\#struct^\Phi_N(\xi)$ be the number of structures with domain $\{1, \ldots, N\}$ over Φ that satisfy ξ. We can proceed, as before, to define

$$\Pr^{s,\Phi}_N(\varphi|\theta) = \frac{\#struct^\Phi_N(\varphi \wedge \theta)}{\#struct^\Phi_N(\theta)} .$$

We define the persistent limit, denoted $\Diamond\Box\Pr^{s,\Phi}_\infty(\varphi|\theta)$, and the intermittent limit, denoted $\Box\Diamond\Pr^{s,\Phi}_\infty(\varphi|\theta)$, in terms of $\Pr^{s,\Phi}_N(\varphi|\theta)$, in analogy to the earlier definitions for random-worlds.

It is clear that $\#world^\Phi_N(\theta) = 0$ iff $\#struct^\Phi_N(\theta) = 0$, so that well definedness (both persistent and intermittent) is equivalent for the two methods, for any φ, θ.

Proposition 2.4: *For any $\varphi, \theta \in \mathcal{L}(\Phi)$, $\Diamond\Box\Pr^w_\infty(\varphi|\theta)$ (resp., $\Box\Diamond\Pr^w_\infty(\varphi|\theta)$) is well defined iff $\Diamond\Box\Pr^{s,\Phi}_\infty(\varphi|\theta)$ (resp., $\Box\Diamond\Pr^{s,\Phi}_\infty(\varphi|\theta)$) is well defined.*

As the following example shows, for the random-structures method the analogue to Proposition 2.1 does not hold; the value of $\Pr^{s,\Phi}_N(\varphi|\theta)$, and even the value of the limit, depends on the choice of Φ. This example, together with Proposition 2.1, also demonstrates that the values of conditional probabilities generally differ between the random-worlds method and the random-structures method. By way of contrast, Fagin [Fag76] showed that the random-worlds and random-structures methods give the same answers for unconditional probabilities, if we do not have constant or function symbols in the language.

[4] When we say that two limits are equal, we mean that one is well defined iff the other is, one exists iff the other does, and if they both exist then they are equal.

Example 2.5: Consider $\Phi = \{P\}$ for a unary predicate P. Let θ be $\exists! x\, P(x) \vee \neg\exists x\, P(x)$ (where, as usual, "$\exists!$" means "exists a unique"), and let φ be $\exists x\, P(x)$. For any domain size N, $\#struct_N^\Phi(\theta) = 2$. In one structure, there is exactly one element satisfying P and $N-1$ satisfying $\neg P$; in the other, all elements satisfy $\neg P$. Therefore, $\Diamond\Box\Pr_\infty^{s,\Phi}(\varphi|\theta) = \frac{1}{2}$.

Now, consider $\Phi' = \{P, Q\}$ for a new unary predicate Q. There are $2N$ structures where there exists an element satisfying P: the element satisfying P may or may not satisfy Q, and of the $N-1$ elements satisfying $\neg P$, any number between 0 and $N-1$ may also satisfy Q. On the other hand, there are $N+1$ structures where all elements satisfy $\neg P$: any number of elements between 0 and N may satisfy Q. Therefore, $\Pr_N^{s,\Phi'}(\varphi|\theta) = \frac{2N}{3N+1}$, and $\Diamond\Box\Pr_\infty^{s,\Phi'} = \frac{2}{3}$.

We know that the asymptotic limit for the random-worlds method will be the same, whether we use Φ or Φ'. Using Φ, notice that the single structure where $\exists! x\, P(x)$ is true contains N worlds (corresponding to the choice of element satisfying P), whereas the other possible structure contains only one world. Therefore, $\Pr_N^w(\varphi|\theta) = \frac{N}{N+1}$, and $\Diamond\Box\Pr_\infty^w(\varphi|\theta) = 1$. ∎

Although the two methods give different answers in general, there are important circumstances under which they agree. One particular case which is of interest to us in this paper is the following:

Proposition 2.6: *If Φ contains at least one non-unary predicate symbol that does not appear in θ, then*

$$\Box\Diamond\Pr_\infty^w(\varphi|\theta) = \Box\Diamond\Pr_\infty^{s,\Phi}(\varphi|\theta)\ .$$

This proposition is a special case of Corollary 2.7 in [GHK93], so we do not prove it here.

3 Nonexistence results

In this section, we show that the limiting probability $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ (and hence $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$) does not always exist. In fact, for most reasonable concepts of limit (including, for example, the Cesàro limit), there are sentences for which the sequence $\Pr_N^w(\varphi|\theta)$ does not converge.

3.1 Nonexistence for conventional limits

As we mentioned above, the fact that asymptotic conditional probabilities do not always exist is well known.

Theorem 3.1: [Lio69, Fag76] *Let Φ be a vocabulary containing at least one non-unary predicate symbol. There exist sentences $\varphi, \theta \in \mathcal{L}(\Phi)$ such that neither $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ nor $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., neither $\Box\Diamond\Pr_\infty^{s,\Phi}(\varphi|\theta)$ nor $\Diamond\Box\Pr_\infty^{s,\Phi}(\varphi|\theta)$) exists, although both are well defined.*

Proof: Fagin's proof of this theorem is quite straightforward. Let R be a binary predicate in Φ (although, clearly, any non-unary predicate will do). Using R and equality, it is not hard to construct sentences φ_{even} and φ_{odd} such that:

- φ_{even} and φ_{odd} both force R to be a symmetric antireflexive binary relation that divides the domain elements into pairs, where i, j is a pair precisely when $R(i,j)$. Both φ_{even} and φ_{odd} force each element to be paired up with at most one other element.

- φ_{even} forces the pairing to be complete; that is, each element is paired up with precisely one domain element. It is clear that φ_{even} is satisfiable if and only if the domain size is even.

- φ_{odd} forces the pairing to be almost-complete; that is, all elements but one are perfectly paired. It is clear that φ_{odd} is satisfiable if and only if the domain size is odd.

We then take φ to be φ_{odd} and θ to be $\varphi_{even} \vee \varphi_{odd}$. Clearly, $\Pr_N^w(\varphi|\theta)$ alternates between 0 and 1 as N increases, and does not approach an asymptotic limit. ∎

Although this shows that the asymptotic limit does not exist in general, a good argument can be made that in this case there is a reasonable degree of belief that one can hold. In the absence of any information about domain size, 1/2 seems the natural answer. Perhaps if we modified our definition of asymptotic probability slightly, we could increase the applicability of our techniques.

There is indeed a reasonable modification that will let us assign a degree of belief of 1/2 in this case: we can use the Cesàro limit instead of the conventional limit.[5] The Cesàro limit of a sequence $s_1, s_2, \ldots$ is the conventional limit of the sequence $s_1, (s_1 + s_2)/2, (s_1 + s_2 + s_3)/3, \ldots$, whose kth element is the average of the first k elements of the original sequence. It is well known that if the conventional limit exists, then so does the Cesàro limit, and they are equal. However, there are times when the Cesàro limit exists and the conventional limit does not. For example, for a sequence of the form $1, 0, 1, 0, \ldots$ (which, of course, is precisely the sequence that arises in the proof of Theorem 3.1), the conventional limit does not exist, but the Cesàro limit does, and is 1/2.

Unfortunately, we show in Section 3.3 that for any definition of limit satisfying some very basic restrictions, the limit of the conditional probabilities may not exist. In particular, the Cesàro limit satisfies these restrictions; therefore, even for Cesàro limits, the non-existence problem still arises.

3.2 Simulating Turing machines

Before we prove the nonexistence theorem, we present the construction on which it is based. All of our lower bounds are also based on this construction. The main idea is the well-known fact that we can use first-order sentences, interpreted over finite domains, to encode (arbitrarily long) prefixes of the computation of a deterministic Turing machine (see [Tra50]). That is, given a Turing machine $\mathbf{M}$, we can define a sentence $\theta_{\mathbf{M}}$ such that any finite model satisfying $\theta_{\mathbf{M}}$ encodes a finite prefix of the computation of $\mathbf{M}$ on empty input. The exact construction is fairly standard, but requires many details; we present only an outline here.

The following definition will turn out to be useful.

Definition 3.2: Let ξ be a formula, and let $\omega(x)$ be a formula with a single free variable x. We define ξ *restricted to* ω to be the formula $\xi' \wedge \xi_\omega$, where ξ' is a conjunction of formulas $\omega(z)$ for any constant or free variable z appearing in ξ, and ξ_ω is defined by induction on the structure of formulas as follows:

[5]We remark that Cesàro limits have been used before in the context of 0-1 laws; see Compton's overview [Com88] for details and further references.

- $\xi_\omega = \xi$ for any atomic formula ξ,
- $(\neg\xi)_\omega = \neg\xi_\omega$,
- $(\xi \wedge \xi')_\omega = \xi_\omega \wedge \xi'_\omega$,
- $(\forall y\, \xi(y))_\omega = \forall y(\omega(y) \Rightarrow \xi_\omega(y))$. ∎

Intuitively, ξ restricted to ω holds if ξ holds on the submodel consisting of the set of elements which satisfy ω.

Given a deterministic Turing machine $\mathbf{M}$, we construct $\theta_{\mathbf{M}}$ as follows. Think of the computation of $\mathbf{M}$ as consisting of a sequence of instantaneous descriptions (IDs), which specify the head position, state, and the contents of (at least) that part of the tape which has been read or written so far. Without loss of generality, we can assume that the jth ID contains exactly the first j symbols on the tape (padding it with blanks if necessary). The construction uses two binary predicate symbols, H and V, to impose a matching "layered" structure on the elements of a finite domain.

More specifically, we force the domain to look like a sequence of n layers for some n, where there are exactly j elements in the jth layer for $1 \leq j < n$, but the last layer may be "incomplete", and have less than n elements. (This ensures that such a partition of domain elements into layers is possible for any domain size.) We construct each layer separately, by assigning each element a *horizontal successor*. The horizontal successor of the ith element in the jth layer is the $(i+1)$st element in the jth layer. This successor must exist except when i is the last element in the layer ($i = j$), or j is the last (and possibly incomplete) layer ($j = n$). We connect one layer to the next by assigning each element a *vertical successor*. The vertical successor of the ith element in the jth layer is the ith element in the $(j+1)$st layer. This successor must exist except if j is the last layer ($j = n$), and possibly if j is the next-to-last layer ($j = n-1$). These two types of successor relationship are captured using H and V: $H(x, y)$ holds iff y is the horizontal successor of x, and $V(x, y)$ holds iff y is the vertical successor of x. Straightforward assertions in first-order logic can be used to constrain H and V to have the right properties.

We use the jth layer to encode the jth ID, using unary predicates to encode the contents of each cell in the ID and the state of the machine $\mathbf{M}$. It is straightforward to write a sentence $\theta_{\mathbf{M}}$ that ensures that this simulation of the Turing machine starts correctly, and continues according to the rules of $\mathbf{M}$. It follows that there is an exact one-to-one correspondence between finite models of $\theta_{\mathbf{M}}$ and finite prefixes of computations of $\mathbf{M}$, as required.

We have assumed that two binary and several unary predicate symbols are available. In fact, it is possible to do all the necessary encoding using only a single binary (or any non-unary) predicate symbol. Because this observation will be important later, we sketch how the extra predicate and constant symbols can be eliminated. First, note that the predicates H and V can be encoded using a single predicate R. Since H holds only between elements in the same layer, and V only between elements in two consecutive layers, we can define $R(x, y)$ to mean $H(x, y)$ in the first case, and $V(x, y)$ in the second (we can construct the sentences so that it is easy to tell whether two elements are in the same layer). Any unary predicate P used in the construction can be eliminated by replacing $P(x)$ with $R(c, x)$ for some special constant symbol c. We then replace $\theta_{\mathbf{M}}$ with $\theta_{\mathbf{M}}$ restricted to $x \neq c$, as in Definition 3.2, thus making the denotation of c a distinguished element which does not participate in the construction of the Turing machine. Finally, it is possible to eliminate the use of constant symbols by using additional variables quantified with "exists unique"; we omit details. However, note for future reference that for every constant we eliminate, we increase the quantifier depth of the formula by one.

This construction has another very useful property. First, note that the layered structure imposed by H and V ensures that every domain element plays a unique role (i.e., for each element we can find a first-order formula with one free variable which holds of that element and no other). So if we (nontrivially) permute the domain elements in one model, we obtain a different (although isomorphic) model. This property has been called *rigidity*. Rigidity implies that, if the domain size is N, every isomorphism class of worlds satisfying $\theta_{\mathbf{M}}$ contains exactly $N!$ worlds. The first important corollary of this is that, for any φ and θ such that $\theta \Rightarrow \theta_{\mathbf{M}}$ is valid, $\Pr_N^{s,\Phi}(\varphi|\theta) = \Pr_N^w(\varphi|\theta)$. Second, note that any two size N models of $\theta_{\mathbf{M}}$ are isomorphic (because the machine $\mathbf{M}$ is assumed to be deterministic and thus has a unique computation path when started on the empty input). From this observation and rigidity, we conclude that the number of size N models of $\theta_{\mathbf{M}}$ is exactly $N!$; this fact will also be useful later.

3.3 Weaker limits

Fagin's non-existence example in Theorem 3.1 was based on a sequence $\Pr_N^w(\varphi|\theta)$ that consistently alternated between 0 and 1. We mentioned at the end of Section 3.1 that using the Cesàro limit in place of the conventional limit when computing the limit of this sequence gives us the plausible answer of $\frac{1}{2}$. This may lead us to hope that by replacing the conventional limit in our definition of asymptotic conditional probability, we can circumvent the nonexistence problem. Unfortunately, this is not the case. It is relatively easy to construct examples that show that even Cesàro limits of the conditional probabilities $\Pr_N^w(\varphi|\theta)$ do not necessarily converge. In this section, we will prove a far more general theorem. Essentially, the theorem shows that no reasonable notion of limit will ensure convergence in all cases. We begin by describing the general framework that allows us to formalize the notion of "reasonable notion of limit".

The Cesàro limit is only one of many well-studied *summability* techniques that weaken the conventional definition of convergence for infinite sequences. These are techniques which try to assign "limits" to sequences that do not converge in the conventional sense. There is a general framework for summability techniques, which we now explain.[6] (See, for example, [PS72] for further details.)

Let $A = (a_{ij})$ be an infinite square matrix; that is, a_{ij} is a (possibly complex) number for each pair of natural numbers i, j. Let $(s_i) = s_1, s_2, s_3, \ldots$ be an infinite sequence. Suppose that, for all i, the series $\sum_{j=1}^{\infty} a_{ij} s_j$ converges, say to sum S_i. Then the new sequence (S_i) is called the *A-transform* of (s_i). The idea is that (S_i) may converge to a limit, even if (s_i) does not. The standard notion of limit can be obtained by taking $a_{ii} = 1$ and $a_{ij} = 0$ if $i \neq j$. The Cesàro limit can be obtained by taking $a_{ij} = 1/i$ if $j \leq i$, and $a_{ij} = 0$ otherwise.

[6]Summability theory is so named because one application is to find a way of assigning a "sum" to series that are divergent according to the conventional notion of limit. However, the theory addresses the problem of convergence for any sequence, whether or not it arises naturally as a sequence of partial sums.

Not every transform makes intuitive sense as a weakened notion of convergence. It would seem reasonable to require, at the very least, the following conditions of a matrix transform A.

- *Computability.* There should be a recursive function f such that $f(i,j)$ is the entry a_{ij} of the matrix A. It is difficult to see how we could actually use a transform whose elements could not be effectively computed.
- *Regularity.* If a sequence converges (in the conventional sense), say to limit ℓ, then the A-transform should exist and converge to ℓ. This ensures that we really do obtain a more general notion of convergence.

The regularity condition has been well studied. The following three conditions are known to be necessary and sufficient for A to be regular. (This result is known as the Silverman-Toeplitz theorem; see [PS72].)

R1. $\lim_{i\to\infty} a_{ij} = 0$, for all j,

R2. $\lim_{i\to\infty} \sum_{j=1}^{\infty} a_{ij} = 1$, and

R3. there exists M such that $\sum_{j=1}^{\infty} |a_{ij}| < M$, for all i.

In our setting—where the motivation is assigning degrees of belief—we can give an fairly intuitive interpretation to many regular summability methods. Fix a value for i and suppose that (1) for all j, a_{ij} is real and nonnegative, and (2) $\sum_{j=1}^{\infty} a_{ij} = 1$. Then the sequence $a_{i1}, a_{i2}, \ldots$ can also be viewed as a probability distribution over possible domain sizes. Given that one accepts the basic random-worlds framework for assigning degrees of belief relative to a particular domain size, it seems plausible that $\sum_{N=1}^{\infty} a_{iN} \Pr_N^w(\varphi|\theta)$ should be one's degree of belief in φ given θ, if the uncertainty about the correct domain size is captured by $a_{i1}, a_{i2}, \ldots$ (and if the probability is defined for all finite N; we discuss how to relax this below). For example, row i of the Cesàro matrix would be appropriate for someone who knows for certain that there are i or less individuals, but subject to this assigns equal degree of belief to each of the i possibilities. However, no single distribution over the natural numbers seems to accurately model the situation where *all* we know is that "the domain size is large." For one thing, any distribution gives nonzero probability to particular domain sizes, which seems to involve some commitment to scale. Instead, we can consider a sequence of distributions, such that the degree of belief in any particular domain size tends to zero. These assumptions imply conditions R1–R3, and therefore suffice to guarantee regularity. Furthermore, they are satisfied by almost all summability transforms considered in the literature.

One subtle problem concerning our application of summability transforms is that some terms in the sequence $\Pr_N^w(\varphi|\theta)$ (or $\Pr_N^{s,\Phi}(\varphi|\theta)$) may not exist. Throughout the following, we adopt perhaps the simplest solution to this difficulty, which is to apply the transform to the subsequence generated by just those domain sizes for which the probability exists (i.e., for which θ is satisfiable).

We are now in a position to state the main result of this section: No summability technique covered by this framework can guarantee convergence for asymptotic conditional probabilities. This is so even if the vocabulary consists of a single binary predicate symbol.

Theorem 3.3: *Let A be any computable regular matrix transform, and let Φ be a vocabulary containing at least one non-unary predicate symbol. There exist $\varphi, \theta \in \mathcal{L}(\Phi)$ such that the A-transform of the sequence $\Pr_N^w(\varphi|\theta)$ (resp., $\Pr_N^{s,\Phi}(\varphi|\theta)$) exists, but does not converge.*

Proof: In the following, let U be a rational number within 0.01 of $\limsup_{i\to\infty} \sum_{j=1}^{\infty} |a_{ij}|$, i.e., $|U - \limsup_{i\to\infty} \sum_{j=1}^{\infty} |a_{ij}|\,| < 0.01$. We will use U as a parameter to the algorithm we are about to construct. Notice that although the existence of an appropriate U is guaranteed by R3, we may not be able to compute its value. Thus, the proof we are about to give is not necessarily constructive. On the other hand, this is the only nonconstructive aspect of our algorithm. A value for U is computable in many cases of interest (for example, if a_{ij} is nonnegative for all i and j, then we can take $U = 1$); in these cases, our proof becomes constructive. Let $i_{\min}$ be such that whenever $i \geq i_{\min}$, we have $\sum_{j=1}^{\infty} |a_{ij}| < U + 0.01$. Such an $i_{\min}$ must exist (because of the way U is defined); it is not necessarily computable either, but the following does not actually depend on its value (i.e., we refer to $i_{\min}$ only when proving that the constructed machine works as required).

We use the value of U in the construction of a three-tape four-head Turing machine $\mathbf{M}$. Tape 2 of $\mathbf{M}$ will always (after the first step) contain an alternating sequence of 0's and 1's. The sentence $\theta_{\mathbf{M}}$ is constructed so that finite models of θ encode partial computations of $\mathbf{M}$, exactly as outlined in Section 3.2. The sentence φ is chosen to be true only in models of θ where the last element written on tape 2 is 1. Note that, as usual, we can assume that $\varphi, \theta_{\mathbf{M}} \in \mathcal{L}(\{R\})$ for a binary predicate symbol R.

The idea of the proof is as follows. Suppose b_j is the truth value of φ (either 0 or 1) in a domain of size j, and let $c_i = \sum_{j=1}^{\infty} a_{ij}\, b_j$. Obviously, the sequence (b_j) is determined by the times at which $\mathbf{M}$ writes a new symbol to tape 2. We construct $\mathbf{M}$ to guarantee that the sequence (b_j) has appropriately spaced runs of zeros and ones, so that there are infinitely many i where c_i is greater than 0.9 and infinitely many i where c_i is less than 0.1. This ensures that the sequence (c_i) does not converge.

As we have said, $\mathbf{M}$ is a three-tape four-head Turing machine. Heads 1a and 1b read tape 1, head 2 reads tape 2, and head 3 reads tape 3. We assume that any subset of heads can move in the same step. Tape 1 is used for keeping track, in unary, of the number of steps that $\mathbf{M}$ has taken so far. Tape 2 contains an alternating sequence of 0's and 1's. As we have indicated, the goal of the rest of the construction will be to ensure that tape 2 is updated at appropriate intervals. Finally, tape 3 is a work tape, used for all necessary calculations.

Every fourth step, head 1a writes a 1 at the right end of tape 1, and then moves one step to the right. This is done independently of the operation of the rest of the machine. Thus, if we represent the number written on tape 1 at a certain point as m, the actual number of steps taken by $\mathbf{M}$ up to that point is between $4m$ and $4m + 3$. Moreover, if we assume (as we do without loss of generality) that the size of the ith ID of the computation of $\mathbf{M}$ is i, then to encode the first i steps of the computation we need a domain of size $i(i+1)/2 + C$, where C is a constant independent of i. In particular, the size of the domain required to encode the prefix of the computation at the point where m is the number on tape 1 is roughly $2m(4m+1)$, and is certainly bounded above by $9m^2$ and below by $7m^2$ for all sufficiently large m. We will use these estimates in describing $\mathbf{M}$.

The machine $\mathbf{M}$ proceeds in phases; each phase ends by writing a symbol on tape 2. At the completion of phase k, for all k large enough, there will exist some number i_k such that

$c_{i_k} < 0.1$ if k is even, and $c_{i_k} > 0.9$ if k is odd. Since we will also show that $i_{k+1} > i_k$, this will prove the theorem.

The first phase consists of one step; at this step, **M** writes 0 on tape 2, and head 2 moves to the right. Suppose the kth phase ends with writing a 1 on tape 2. We now describe the $(k+1)$st phase. (The description if the kth phase ends with writing a 0 is almost identical, and left to the reader.)

Let n_l be the size of the domain required to encode the prefix of the computation up to the end of phase l. Since the value at the end of tape 2 changes only at the end of every phase, and b_j is 1 if and only if the last element on tape 2 is 1, b_j is 0 for $n_1 \leq j < n_2$, b_j is 1 for $n_2 \leq j < n_3$, and so on. **M** begins the $(k+1)$st phase by copying the number m on tape 1 to tape 3 (the work tape). The copying is done using head 1b (head 1a continues to update the number every fourth step). Suppose the number eventually copied is m_k. Clearly, m_k will be greater than the number that was on tape 1 in the computation prefix that was encoded by domain size n_k. Therefore, $n_k < 9m_k^2$ for k sufficiently large.

We now get to the heart of the construction, which is the computation of when to next write a value on tape 2. (Note that this value will be a 0, since we want the values to alternate.) Notice that by R1, R2, and R3 there must be a pair (i_*, j_*) such that:

(a) $i_* > m_k$,

(b) $\sum_{j=1}^{9m_k^2} |a_{i_*j}| < 0.01$,

(c) $\sum_{j=1}^{j_*} a_{i_*j} > 0.99$, and

(d) $\sum_{j=1}^{j_*} |a_{i_*j}| > U - 0.01$.

Moreover, since a_{ij} is computable for all i and j, **M** can effectively find such a pair by appropriate dovetailing. Suppose that in fact $i_* > i_{\min}$. (Since $i_* > m_k$ by part (a), this will be true once k is large enough.) Then we claim that, no matter what the values of $b_0, \ldots, b_{n_k}$ and $b_{j_*+1}, b_{j_*+2}, \ldots$, if $b_{n_k+1} = \cdots = b_{j_*} = 1$, then $c_{i_*} > 0.9$. To see this, note that if $i_* > i_{\min}$, then (by definition of $i_{\min}$) $\sum_{j=1}^{\infty} |a_{i_*j}| < U + 0.01$. Thus, by part (d) above it follows that $\sum_{j=j_*+1}^{\infty} |a_{i_*j}| < 0.02$. Using part (b) and the fact that $n_k < 9m_k^2$, it follows that $\sum_{j=1}^{n_k} |a_{i_*j}| < 0.01$. Now from part (c) we get that $\sum_{j=n_k+1}^{j_*} a_{i_*j} > 0.98$. If $b_{n_{k+1}} = \cdots = b_{j_*} = 1$, then

$$
\begin{aligned}
c_{i_*} &= \textstyle\sum_{j=1}^{\infty} a_{i_*j} b_j \\
&= \textstyle\sum_{j=1}^{n_k} a_{i_*j} b_j + \sum_{j=n_k+1}^{j_*} a_{i_*j} b_j + \sum_{j=j_*+1}^{\infty} a_{i_*j} b_j \\
&\geq \textstyle\sum_{j=n_k+1}^{j_*} a_{i_*j} - \sum_{j=1}^{n_k} |a_{i_*j}| - \sum_{j=j_*+1}^{\infty} |a_{i_*j}| \\
&\geq 0.98 - 0.01 - 0.02 \\
&> 0.9.
\end{aligned}
$$

Thus, it suffices for **M** to add the next 0 to tape 2 so as to guarantee that $n_{k+1} > j_*$, since our choice of φ will then guarantee that $b_{n_k+1} = \cdots = b_{j_*} = 1$. This can be done by waiting to add the 0, until after the number m on tape 1 is such that $7m^2 > j_*$. As we observed above, the size of the domain required to encode the prefix of the computation up to this point is at least $7m^2$. Since this domain size is n_{k+1} by definition, it follows that $n_{k+1} \geq j_*$, as desired.

This completes the description of the $(k+1)$st phase. We can then take $i_{k+1} = i_*$, and guarantee that $c_{i_{k+1}} > 0.9$, as desired. Note that, for every k, $i_{k+1} > m_k$, and (m_k) is a strictly increasing sequence. Thus, we obtain infinitely many indices i at which $c_i > 0.9$ and infinitely many at which $c_i < 0.1$, as desired.

Since $\#world_N^{\{R\}}(\theta) \neq 0$ for all sufficiently large N, this shows that both $\Box\Diamond\Pr^w_\infty(\varphi|\theta)$ and $\Diamond\Box\Pr^w_\infty(\varphi|\theta)$ are well defined, but their A-transform does not converge. The case of random-structures follows immediately, because for every N, $\Pr^w_N(\varphi|\theta)$ is either 0 or 1. Consequently $\Pr^{s,\{R\}}_N(\varphi|\theta)$ has the same value as $\Pr^w_N(\varphi|\theta)$, and the limiting behavior is the same. ∎

We remark that there are a few well-known summability methods which are not, strictly speaking, matrix transforms. Nevertheless, our theorem is applicable to these cases as well (at least, to all cases we are aware of). The best example of this is *Abel convergence*. A sequence (s_j) is said to be Abel convergent if $\lim_{x\to 1^-}(1-x)\sum_{j=1}^\infty s_j\, x^{(j-1)}$ exists. This is not a matrix transform, because we must consider all sequences of x that tend to 1. However, consider any particular sequence of rationals that converges to 1, say

$$\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \ldots, \frac{i}{i+1}, \ldots$$

We can use these to define a matrix variant of the Abel method, by setting

$$a_{ij} = \frac{\left(\frac{i}{i+1}\right)^{(j-1)}}{i} .$$

This is regular and computable, and is strictly weaker than the standard Abel method. More precisely, if the Abel limit converges, then so does this matrix transform. Since our theorem shows that this new summability method does not ensure convergence for conditional probabilities, this is automatically also the case for the Abel limit.

4 Undecidability results

We have seen that asymptotic conditional probabilities do not always exist. We might hope that at least we can easily decide when they do exist, so that we would know when the random-worlds or random-structures method is applicable. As we show in this section, this hope is not realized. In this section, we show the undecidability of several important problems associated with asymptotic conditional probabilities: deciding whether the limit is well defined, deciding whether the limit exists, and giving some nontrivial approximation to its value (deciding whether it lies in some nontrivial interval). Liogon'kiĭ [Lio69] showed that the problem of computing the asymptotic conditional probability for the random-worlds method is undecidable. He did not consider other problems, nor did he characterize the degree of undecidability of the problem.

We analyze the complexity of these problems in terms of the *arithmetic hierarchy*. This is a hierarchy that extends the notions of r.e. (recursively enumerable) and co-r.e. sets. We briefly review the relevant definitions here, referring the reader to [Rog67, Chapter 14] for further details. Consider a formula ξ in the language of arithmetic (i.e., using $0, 1, +, \times$) having

j free variables. The formula ξ, interpreted over the natural numbers, is said to define a *recursive set* if the set of j-tuples satisfying the formula is a recursive set. We can define more complex sets using quantification. We define a Σ^0_k *prefix* as a block of quantifiers of the form $\exists x_1 \ldots x_h \forall y_1 \ldots y_m \ldots$, where there are k alternations of quantifiers (but there is no restriction on the number of quantifiers of the same type that appear consecutively). A Π^0_k prefix is defined similarly, except that the quantifier block starts with a universal quantifier. A set A of natural numbers is in Σ^0_k if there is a first-order formula $\xi(x) = Q\xi'$ in the language of arithmetic with one free variable x, such that $n \in A$ iff $\xi(n)$ is true, where Q is a Σ^0_k quantifier block and ξ' defines a recursive set. We can similarly define what it means for a set to be in Π^0_k. A set is in Σ^0_1 iff it is r.e. and it is in Π^0_1 iff it is co-r.e. The hierarchy is known to be strict; higher levels of the hierarchy correspond problems which are strictly harder ("more undecidable").

We prove the following results for the random worlds method. We later modify the proofs to apply to the random-structures method.

- Deciding whether $\Box\Diamond\mathrm{Pr}^w_\infty(\varphi|\theta)$ is well defined is Π^0_2-complete.
- Deciding whether $\Diamond\Box\mathrm{Pr}^w_\infty(\varphi|\theta)$ is well defined is Σ^0_2-complete.
- Deciding whether $\Diamond\Box\mathrm{Pr}^w_\infty(\varphi|\theta)$ (resp., $\Box\Diamond\mathrm{Pr}^w_\infty(\varphi|\theta)$) exists is Π^0_3-complete, given that the limit is well defined.
- Deciding whether the limit $\Diamond\Box\mathrm{Pr}^w_\infty(\varphi|\theta)$ (resp., $\Box\Diamond\mathrm{Pr}^w_\infty(\varphi|\theta)$) is in some nontrivial closed interval is Π^0_2-complete, given that the limit exists.

The lower bounds all rely on the construction in Section 3.2, and use a fixed finite vocabulary, consisting of equality and a single binary predicate. Most of them can, in fact, be translated to a language without equality, at the cost of adding two more binary predicates (see Section 4.4).

4.1 Well definedness of the limit

We start with the problem of deciding if the asymptotic probability is well defined; this is certainly a prerequisite for deciding whether the limit exists. Of course, this depends in part on which definition of well definedness we use.

Theorem 4.1: *Let Φ be a vocabulary containing at least one non-unary predicate symbol.*

(a) The problem of deciding whether a sentence in $\mathcal{L}(\Phi)$ is satisfiable for infinitely many domain sizes is Π^0_2-complete.

(b) The problem of deciding whether a sentence in $\mathcal{L}(\Phi)$ is satisfiable for all but finitely many domain sizes is Σ^0_2-complete.

Proof: We start with the upper bounds. First observe that the problem of deciding whether a first-order sentence ξ is satisfiable in some model with domain size N, for some fixed N, is recursive (and with the help of some suitable encoding of formulas as natural numbers, we can encode this problem in the language of arithmetic). Given this, deciding if ξ is satisfiable in infinitely many domain sizes can be encoded using a Π^0_2 block: for all N, there exists $N' > N$

such that ξ holds in some model of domain size N'. Similarly, deciding if ξ is satisfiable for all but finitely many domain sizes can clearly be encoded using a Σ_2^0 block: there exists N such that for all $N' > N$, ξ holds in some model with domain size N'. This proves the upper bounds.

It is well known that the following problem is Π_0^2-complete: "Given a Turing machine $\mathbf{M}$, does $\mathbf{M}$ halt on infinitely many inputs?", and the dual problem—"Given a Turing machine $\mathbf{M}$, does $\mathbf{M}$ halt on only finitely many inputs?"—is Σ_0^2-complete [Rog67, Theorem 13-VIII; Corollary 14-VIII(b)]. We prove the two lower bounds by reducing these problems to intermittent and persistent well definedness, respectively. First, given an arbitrary Turing machine $\mathbf{M}$, we effectively construct another Turing machine $\mathbf{M}'$ that, when started on empty input, starts simulating the computations of $\mathbf{M}$ on all inputs by dovetailing, and enters a special state q_s once for each input on which $\mathbf{M}$ halts. (We leave details of this construction to the reader.) Let $\theta_{\mathbf{M}'}$ be the sentence that forces its models to encode prefixes of the computation of $\mathbf{M}'$ on empty input, as described in Section 3.2, and let φ be the sentence that says, with respect to this encoding, that the last layer is complete, and that $\mathbf{M}'$ is in state q_s in the ID encoded in this last layer. Clearly $\varphi \wedge \theta_{\mathbf{M}'}$ is satisfiable for infinitely many domain sizes N iff $\mathbf{M}$ halts on infinitely many inputs, while $\neg\varphi \wedge \theta_{\mathbf{M}'}$ is satisfiable for all but finitely many domain sizes N iff $\mathbf{M}$ halts on only finitely many inputs. This proves the lower bounds. ∎

Corollary 4.2: *Let Φ be a vocabulary containing at least one non-unary predicate symbol. For $\varphi, \theta \in \mathcal{L}(\Phi)$, the problem of deciding whether $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^{s,\Phi}(\varphi|\theta)$) is well defined is Π_2^0-complete, and the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., $\Diamond\Box\Pr_\infty^{s,\Phi}(\varphi|\theta)$) is well defined is Σ_2^0-complete.*

4.2 Existence of the limit

If deciding well definedness were the only difficulty in computing, then there might still be hope. In many cases, it might be obvious that the sentence we are conditioning on is satisfiable in all (or, at least, in infinitely many) domain sizes. As we are about to show, the situation is actually much worse. Deciding if the limit exists is even more difficult than deciding well definedness; in fact, it is Π_3^0-complete. We prove this result by first showing that the problem of deciding whether an r.e. sequence of rationals converges to 0 is Π_3^0-complete.

Theorem 4.3: *The problem of deciding whether a recursively enumerable infinite sequence of rational numbers converges to zero is Π_3^0-complete.*

Proof: The following problem is known to be Π_3^0-complete: "Does each of the Turing machines in a given r.e. set of Turing machines diverge on all but finitely many inputs?", where the input to this problem is itself a Turing machine (that generates the encodings for the collection of Turing machines we are asking about). More precisely, taking W_x to be the x^{th} r.e. set in some enumeration of r.e. sets (or, equivalently, the inputs on which the x^{th} Turing machine halts, in some enumeration of Turing machines), the set $\{z : \forall y(y \in W_z \Rightarrow W_y \text{ finite}\}$ is Π_3^0-complete. (The complement of this set is proved to be Σ_3^0-complete in Theorem 14-XV of [Rog67].) For our purposes it is slightly better to consider a variant of this problem, namely "Does each of the Turing machines in a given r.e. set of Turing machines enter some distinguished state, say

q_s, only finitely many times when started on the empty input?" The two problems are easily seen to be equivalent, in that either one can be effectively reduced to the other.

The lower-bound is proved by reducing this problem to the question of whether a sequence converges to zero. We assume, without loss of generality, that our Turing machine generator $\mathbf{G}$ computes a total function, whose values are encodings of other Turing machines. That is, on input i, it is guaranteed to terminate and produce the ith machine (note that the machines produced by $\mathbf{G}$ on different inputs are not necessarily distinct). We now define H_{ij} to have value 1 if the ith machine generated by $\mathbf{G}$ is in state q_s on its jth step after being started on empty input, and value 0 otherwise. Note that H_{ij} is a computable function of i, j, and the encoding of $\mathbf{G}$, because we can simulate $\mathbf{G}$ to obtain the encoding of the ith machine, then simulate this machine for j steps.

We use the numbers H_{ij} to define an r.e. sequence $s_1, s_2, \ldots$ of rational numbers in [0,1], where s_k is defined as $0.\, H_{1k}\, H_{2k}\, \ldots H_{kk}$. The computability of H_{ij} guarantees that this sequence is recursively enumerable. Clearly the sequence $s_1, s_2, \ldots$ converges to 0 iff, for all i, the sequence $H_{i1}, H_{i2}, \ldots$ is eventually 0, i.e., there exists n_i such that $H_{ij} = 0$ for all $j > n_i$. But the sequence $H_{i1}, H_{i2}, \ldots$ is eventually 0 iff the ith Turing machine reaches q_s only finitely often. This proves the lower bound.

For the upper bound, note that the question of whether the limit of $s_1, s_2, \ldots$ exists and equals 0 can be written: "For all M, does there exist N_0 such that for all $N > N_0$, $|s_N| < 1/M$?" The unquantified part of this question is clearly recursive and can be formulated in the language of arithmetic, while the quantifier block is a Π_3^0 prefix. The result follows. ∎

Theorem 4.4: *Let Φ be a vocabulary containing at least one non-unary predicate symbol. For sentences $\varphi, \theta \in \mathcal{L}(\Phi)$, the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$) exists is Π_3^0-complete. The lower bound holds even if we have an oracle that tells us whether the limit is well defined and its value if it exists.*

Proof: To prove the lower bound, we reduce the problem of deciding if an r.e. sequence of rationals converges to 0 to that of deciding if a particular asymptotic conditional probability exists. Suppose $\mathbf{S}$ is a machine that generates an infinite sequence of rational numbers, $s_1, s_2, \ldots$. Without loss of generality, we can assume that the numbers are in $[0, 1]$; if necessary, a new machine $\mathbf{S}'$ such that $s'_i = max(1, |s_i|)$ is easily constructed which clearly has the same properties with respect to convergence to zero. We also assume that the output is encoded in a special form: a rational value a/b is output on the tape as a sequence of a 1's, followed by $(b - a)$ 0's, suitably delimited.

Let R be a binary predicate symbol. (Of course, any non-unary predicate will suffice.) We begin by constructing $\theta_{\mathbf{S}} \in \mathcal{L}(\{R\})$ such that finite models of $\theta_{\mathbf{S}}$ correspond naturally to prefixes of computations of $\mathbf{S}$, as described in Section 3.2. Let c be a constant. Let $\theta'_{\mathbf{S}} \in \mathcal{L}(\{c, R\})$ be the conjunction of $\theta_{\mathbf{S}}$ and sentences asserting that, in the computation-prefix of $\mathbf{S}$ encoded by the domain, the denotation of c corresponds to a cell in that section of the last complete ID that represents the output. Note that for any fixed domain size, $\theta'_{\mathbf{S}}$ has $a + (b - a) = b$ times as many models over $\{c, R\}$ as $\theta_{\mathbf{S}}$ does over $\{R\}$, where a/b is the most recent sequence value generated by $\mathbf{S}$ in the computation simulated so far. According to our discussion at the end of Section 3.2, $\#world_N^{\{R\}}(\theta_{\mathbf{S}}) = N!$, so $\#world_N^{\{c,R\}}(\theta'_{\mathbf{S}}) = b \cdot N!$.

To complete the reduction, consider a sentence φ that says that the simulated computation has just finished writing another sequence element, and the denotation of c corresponds to a cell in that output containing the symbol 1. Assume that the last sequence element written in the prefix corresponding to domain size N is a/b. Note that if there are models of $\varphi \wedge \theta'_{\mathbf{S}}$ of domain size N, then there are in fact $a \cdot N!$ such models over $\{c, R\}$ (corresponding to the a choices for the denotation of c). In this case $\Pr_N^w(\varphi|\theta'_{\mathbf{S}})$ has value a/b. It follows that the sequence $\Pr_N^w(\varphi|\theta'_{\mathbf{S}})$, for increasing N, is precisely the sequence generated by $\mathbf{S}$ interspersed with zeros at domain sizes corresponding to computations that have not just output a new value. Note that both persistent and intermittent limits are well defined for this sequence. If this limit exists at all, it must have value zero, and this will be the case just if the sequence generated by $\mathbf{S}$ has this property. This proves the lower bound. We remark that the use of an extra constant c is not necessary in our proof; it can be eliminated as discussed in Section 3.2.

To prove the upper bound, note that the question of existence for $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ can be stated as: "Is it true that for all integers M, there exist rational numbers $r_1 \leq r_2$ and integers N_0 and $N_1 > M$ such that for all $N \geq N_0$, (1) $\#world_{N_1}^{\Phi}(\theta) \neq 0$, (2) if $\#world_N^{\Phi}(\theta) \neq 0$, then $\Pr_N^w(\varphi|\theta) \in [r_1, r_2]$, and (3) $r_2 - r_1 \leq 1/M$?" The unquantified part is clearly recursive, showing that the problem of deciding whether $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ exists is in Π_0^3. We can state the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ exists as follows: "Is it true that for all integers M, there exist rational numbers $r_1 \leq r_2$ and an integer N_0 such that for all $N \geq N_0$, (1) $\#world_N^{\Phi}(\theta) \neq 0$, (2) $\Pr_N^w(\varphi|\theta) \in [r_1, r_2]$, and (3) $r_2 - r_1 \leq 1/M$?" Thus, the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ exists is also in Π_3^0. ∎

4.3 Computing the limit

Even if we have an oracle that will tell us whether the conditional probability is well defined and whether it exists, it is difficult to compute the asymptotic probability. Indeed, given any nontrivial interval (one not of the form $[0, 1]$), it is even difficult to tell whether the asymptotic probability is in the interval.

Theorem 4.5: *Let Φ be a vocabulary containing at least one non-unary predicate symbol, and let $r, r_1, r_2 \in [0, 1]$ be rational numbers such that $r_1 \leq r_2$. For sentences $\varphi, \theta \in \mathcal{L}(\Phi)$, given an oracle for deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$) exists,*

(a) the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta) = r$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) = r$) is Π_2^0-complete,

(b) if $[r_1, r_2] \neq [0, 1]$, then the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta) \in [r_1, r_2]$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) \in [r_1, r_2]$) is Π_2^0-complete,

(c) if $r_1 \neq r_2$, then the problem of deciding if $\Diamond\Box\Pr_\infty^w(\varphi|\theta) \in (r_1, r_2)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) \in (r_1, r_2)$) is Σ_2^0-complete.

Proof: We start with part (a). Just as with our earlier results, the upper bound is the easier part. This problem can be stated as "For all M, does there exist an $N > M$ such that $\#world_N^{\Phi}(\theta) > 0$, and $|\Pr_N^w(\varphi|\theta) - r| < 1/M$?" It is easy to see that this sentence has the appropriate form for Π_2^0. Furthermore, it is true just if there is some subsequence of domain sizes such that the asymptotic probability, when restricted to these sizes, has value r. If the sequence as a whole has a limit at all (and we can check this with the oracle) then this limit must also be r.

To prove the lower bound, we proceed just as in the proof of Theorem 4.1 by reducing the problem "Does a Turing machine reach a specified state q_s infinitely often?" to the problem of deciding whether the asymptotic probability is r. Let $\mathbf{M}$ be an arbitrary Turing machine. As discussed in Section 3.2, we can find a sentence $\theta_{\mathbf{M}} \in \mathcal{L}(\{R\})$ such that finite models of $\theta_{\mathbf{M}}$ correspond naturally to prefixes of computations of $\mathbf{M}$.

Our next step is to construct sentences φ_r and θ_r such that $\Pr_N^w(\varphi_r|\theta_r) = r$, for all N. Suppose $r = a/b$. Choose k such that $2^k > b$. We can easily construct propositional formulas α_r and β_r using k primitive propositions $p_1, \dots, p_k$ such that β_r has exactly b satisfying assignments and $\alpha_r \wedge \beta_r$ has exactly a satisfying assignments. Let φ_r and θ_r be the sentences that result by replacing occurrences of the primitive proposition p_i in α_r or β_r by $P_i(c)$, where P_i is a unary predicate symbol and c is a constant symbol. It is ea sy to see that $\Pr_N^w(\varphi_r|\theta_r) = r$ for all N.

Let Q be a unary predicate not among $\{P_1, \dots, P_k\}$, and let θ' be a sentence asserting that there is exactly one domain element satisfying Q, and that this element corresponds to one of the tape cells representing the head position when the machine is in state q_s. Define θ to be $\theta_{\mathbf{M}} \wedge \theta_r \wedge (\theta' \vee \forall x\, Q(x))$. For any domain size N, let t_N denote the number of times the machine has reached q_s in the computation so far. The sentence θ has $t_N + 1$ times as many models over $\{R, P_1, \dots, P_k, Q, c\}$ as the sentence $\theta_{\mathbf{M}} \wedge \theta_r$ has over $\{R, P_1, \dots, P_k, c\}$. We now consider two cases: $r < 1$ and $r = 1$. If $r < 1$, let φ be simply $\varphi_r \vee (\neg\varphi_r \wedge \forall x\, Q(x))$. It is easy to see that $\Pr_N^w(\varphi|\theta)$ is $r + (1-r)/(t_N + 1)$. If $\mathbf{M}$ reaches q_s finitely often, say t' times, the limit as $N \to \infty$ is $r + (1-r)/(t'+1)$, otherwise the limit is r. The limit always exists, so our oracle is not helpful. This proves the required lower bound if $r < 1$. If $r = 1$, then we can take θ to be $\theta_{\mathbf{M}} \wedge (\theta' \vee \forall x\, Q(x))$ and φ to be $\neg\forall x\, Q(x)$. In this case, $\Pr_N^w(\varphi|\theta)$ is $t_N/(t_N+1)$; therefore, the limit is 1 if $\mathbf{M}$ reaches q_s infinitely often, and strictly less than 1 otherwise. Again, the lower bound follows. Note that, as discussed in Section 3.2, we can avoid actually using new unary predicates and constants by encoding them with the binary predicate R.

For part (b), the upper bound follows using much the same arguments as the upper bound for part (a). For the lower bound, we also proceed much as in part (a). Suppose we are given an interval $[r_1, r_2]$ with $r_2 < 1$, and a Turing machine $\mathbf{M}$. Using the techniques of part (a), we can construct sentences φ and θ such that $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ and $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ are both well defined, and such that the asymptotic probability is r_2 if $\mathbf{M}$ reaches state q_s infinitely often, and strictly greater than r_2 otherwise. This proves the lower bound in this case. If $r_2 = 1$, we use similar arguments to construct sentences φ and θ such that the asymptotic conditional probability is r_1 if $\mathbf{M}$ reaches state q_s infinitely often, and is strictly less than r_1 otherwise. Again, the lower bound follows.

Finally, for part (c), observe that the asymptotic probability is in (r_1, r_2) iff it is not in $[0, r_1] \cup [r_2, 1]$. The arguments of part (b) showing that checking whether the asymptotic probability is in a closed interval is Π_2^0-complete can be extended without difficulty to dealing with the union of two closed intervals. Thus, the problem of deciding whether the asymptotic probability is in an open interval is Σ_2^0-complete. ∎

It is easy to see that analogues to Theorems 4.4 and 4.5 hold for the random-structures method as well. The upper bounds hold with no change in proof. The same is true for the lower bounds as well, since, as we observed in Section 3.2, the sentences constructed to show that the lower bounds hold are true only in rigid structures (and therefore, random worlds and random structures agree).

4.4 Eliminating Equality

At first glance, it seems that the proofs of all the above results make heavy use of equality. As we now show, we can eliminate the use of equality from most of these results, at the price of adding two more binary predicate symbols to the vocabulary.

Theorem 4.6: *Suppose G and E are binary predicate symbols not appearing in Φ, and $\varphi, \theta \in \mathcal{L}(\Phi)$ are such that $\#world_N^{\Phi}(\theta)$ is a non-decreasing function of N. Then we can find sentences $\varphi', \theta' \in \mathcal{L}^-(\Phi \cup \{G, E\})$ such that*

$$\lim_{N\to\infty} (\Pr_N^w(\varphi|\theta) - \Pr_N^w(\varphi'|\theta')) = 0 \ .$$

Proof: The idea of the proof is somewhat similar to that used in [KV90] to eliminate equality. Let φ and θ be as in the hypotheses of the theorem. Define θ^E to be the result of replacing all subformulas of θ of the form $t_1 = t_2$ by $E(t_1, t_2)$; we define φ^E similarly. Thus, we are using E to represent equality. Let η be a conjunction of formulas that force E to be an equivalence relation, as well as a congruence on G and on all symbols in Φ. Thus, a typical conjunct of η (which in fact forces E to be a congruence on G) has the form:

$$\forall x\, y\, z (E(x, y) \Rightarrow ((G(x, z) \Leftrightarrow G(y, z)) \wedge (G(z, x) \Leftrightarrow G(z, y)))).$$

Let θ' be $\theta^E \wedge \eta$, and φ' be φ^E.

As we now show, there are many more models of θ' of size N where E is true equality than there are where E is some equivalence relation other than equality. To simplify the notation, we write w_N instead of $\#world_N^{\Phi}(\theta)$. It is easy to see that there are precisely $w_N \cdot 2^{N^2}$ models of size N of θ' over $\Phi \cup \{G, E\}$ where E is equality: for every model of size N of θ over Φ, there are 2^{N^2} models of θ', because the choice of G is unrestricted.

Now we must get an estimate on the number of models of θ' where E is an equivalence relation, but not equality. It turns out that the crucial factor is the number of equivalence classes into which E partitions the domain. Let $\left\{ {N \atop k} \right\}$ be the number of ways of partitioning N elements into exactly k equivalence classes. ($\left\{ {N \atop k} \right\}$ is known as a *Stirling number of the second kind*; see [GKP89].) It is easy to see that there are $w_k \cdot \left\{ {N \atop k} \right\} \cdot 2^{k^2}$ models of θ' where E partitions the domain into k equivalence classes, since for each such way, there are 2^{k^2} choices for G, and w_k choices for the denotations of the predicates in Φ that make θ^E true. Thus, our goal is to show that $(\sum_{k=1}^{N-1} w_k \cdot \left\{ {N \atop k} \right\} \cdot 2^{k^2})/w_N \cdot 2^{N^2}$ asymptotically converges to 0.

To do this, we need a good estimate on $\left\{ {N \atop k} \right\}$. We begin by showing that $\binom{N}{k} N!$ is an overestimate for $\left\{ {N \atop k} \right\}$. To see this, consider any partition, order the equivalence classes by the minimal elements appearing in them, and order the elements in an equivalence class in increasing order. This gives us an ordering of the N elements in the domain. Suppose the equivalence classes (listed in this order) have size $n_1, \ldots, n_k$. This corresponds to choosing elements $n_1, n_1 + n_2, \ldots, n_1 + \cdots + n_k$ from the domain. Thus, with each partition into k equivalence classes, we can associate a unique pair consisting of a permutation and a choice of k elements out of N.

This estimate suffices for values of k which are relatively small compared to N. We use a finer estimate for $\left\{ {N \atop k} \right\}$ if $k \geq N - \log N$. In this case, at least $k - \log N$ equivalence classes must have size 1. The remaining $\log N$ equivalence classes partition at most $N - (k - \log N) \leq 2\log N$ elements. Thus, a bound on $\left\{ {N \atop k} \right\}$ in this case is given by

$$\begin{aligned}
\binom{N}{k-\log N}\left\{ {N-(k-\log N) \atop \log N} \right\} &\leq \binom{N}{N-2\log N}\left\{ {2\log N \atop \log N} \right\} \\
&\leq \binom{N}{N-2\log N}\binom{2\log N}{\log N}(2\log N)! \\
&\leq \binom{N}{N-2\log N} 2^{2\log N}\,(2\log N)! \\
&= \frac{N!}{(N-2\log N)!}\,2^{2\log N} \\
&\leq N^{2\log N}\,2^{2\log N} \\
&= 2^{2\log^2 N + 2\log N}\ .
\end{aligned}$$

Thus, we have that

$$\begin{aligned}
\sum_{k=1}^{N-1}\left\{ {N \atop k} \right\}\cdot 2^{k^2} &= \sum_{k=1}^{N-\log N}\left\{ {N \atop k} \right\}\cdot 2^{k^2} + \sum_{k=N-\log N+1}^{N-1}\left\{ {N \atop k} \right\}\cdot 2^{k^2} \\
&\leq N!\;2^{(N-\log N)^2}\left(\sum_{k=1}^{N-\log N}\binom{N}{k}\right) + 2^{2\log^2 N+2\log N}\sum_{k=N-\log N+1}^{N-1} 2^{k^2} \\
&\leq 2^{N\log N}\;2^{(N-\log N)^2}2^N + 2^{2\log^2 N+2\log N}2^{(N-1)^2+1} \\
&\leq 2^{N^2-N\log N+N+\log^2 N} + 2^{N^2-2N+2\log^2 N+2\log N+2} \\
&\leq 2^{N^2-\Omega(N)}\ .
\end{aligned}$$

Let σ be the formula $E(x,y) \Leftrightarrow x = y$, which says that E is true equality. (Note that σ is not in $\mathcal{L}^-(\Phi \cup \{G, E\})$, since it mentions $=$, but that is not relevant to the discussion below.) It now easily follows that for any $\epsilon > 0$, we can choose N_0 large enough, so that for any $N > N_0$,

$$\begin{aligned}
\Pr_N^w(\neg\sigma|\theta') &\leq \frac{\sum_{k=1}^{N-1} w_k \cdot \left\{ {N \atop k} \right\}\cdot 2^{k^2}}{w_N \cdot 2^{N^2}} \\
&\leq \frac{w_N \sum_{k=1}^{N-1}\left\{ {N \atop k} \right\}\cdot 2^{k^2}}{w_N \cdot 2^{N^2}} \\
&\leq \frac{2^{N^2-\Omega(N)}}{2^{N^2}} = 2^{-\Omega(N)} < \epsilon/2\ .
\end{aligned}$$

Therefore, since $\Pr_N^w(\varphi'|\theta' \wedge \sigma) = \Pr_N^w(\varphi|\theta)$, it follows that

$$\begin{aligned}
|\Pr_N^w(\varphi'|\theta') - \Pr_N^w(\varphi|\theta)| &= |[\Pr_N^w(\varphi'|\theta' \wedge \sigma) \cdot \Pr_N^w(\sigma|\theta') + \Pr_N^w(\varphi'|\theta' \wedge \neg\sigma) \cdot \Pr_N^w(\neg\sigma|\theta')] - \Pr_N^w(\varphi|\theta)| \\
&\leq |\Pr_N^w(\varphi|\theta)(1 - \Pr_N^w(\sigma|\theta'))| + |\Pr_N^w(\neg\sigma|\theta')| \\
&\leq \epsilon/2 + \epsilon/2 = \epsilon \ ,
\end{aligned}$$

thus completing the proof. ∎

Using Theorem 4.6, we can show analogues to most of our results for the language with equality. First, we can immediately deduce the following corollary to Theorem 3.3.

Corollary 4.7: *Let A be any computable regular matrix transform, and let Φ be a vocabulary containing at least three non-unary predicate symbols. There exist $\varphi, \theta \in \mathcal{L}(\Phi)$ such that the A-transform of the sequence $\Pr_N^w(\varphi|\theta)$ (resp., $\Pr_N^{s,\Phi}(\varphi|\theta)$) exists, but does not converge.*

Proof: It is easy to verify that for the θ used in the proof of Theorem 3.3, it is indeed the case that $\#world_N^\Phi(\theta)$ is a non-decreasing function of N. ∎

We now show that similar analogues to most of the complexity results of this section also hold. The exceptions are Theorem 4.1 and Corollary 4.2.

For a language with no equality, $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ is well defined iff $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ is well defined iff θ is satisfiable for some model. This is true because if θ is satisfied in some model of size N, then it is also satisfied in some model of size N' for every $N' > N$. As a consequence, we can show:

Theorem 4.8: *Let Φ be a vocabulary containing at least two non-unary predicate symbols. For $\varphi, \theta \in \mathcal{L}^-(\Phi)$, the problem of deciding if $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ (resp., $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$) is well defined is r.e.-complete.*

Proof: We can state the problem of deciding whether $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$ is well defined as follows: Does there exist an $N > 0$ for which $\#world_N^\Phi(\theta) > 0$. The unquantified part is clearly recursive, thus proving the upper bound. For the lower bound, we proceed as before. For a given Turing machine $\mathbf{M}$, we let $\theta_\mathbf{M}$ encode a prefix of the computation of $\mathbf{M}$ on empty input which is a complete prefix currently in an accepting state. Let $\theta_\mathbf{M}^E$ be the same formula, but with equality replaced by the binary predicate E, as in the proof of Theorem 4.6. Let η be the formula forcing E to be an equivalence relation and a congruence on R. The sentence $\theta_\mathbf{M}^E \wedge \eta$ is satisfiable in infinitely many domain sizes iff it is satisfiable for some domain size iff $\mathbf{M}$ halts. Note that we did not need the additional predicate G in this proof. ∎

We now show that the remaining complexity results do carry over. It is clear that all our upper bounds hold trivially for the language without equality. We consider the lower bounds, one by one.

Theorem 4.9: *Let Φ be a vocabulary containing at least three non-unary predicate symbols. For sentences $\varphi, \theta \in \mathcal{L}^-(\Phi)$, the problem of deciding if $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$) exists is Π_3^0-complete. The lower bound holds even if we have an oracle that tells us whether the limit is well defined.*

Proof: The sentence $\theta'_{\mathbf{S}}$ in Theorem 4.4 does not satisfy the requirement of Theorem 4.6, since $\#world_N^{\Phi}(\theta'_{\mathbf{S}}) = N! \cdot b$, where a/b is the is the most recent sequence value generated by $\mathbf{S}$ in the computation so far. The values of b do not necessarily form a non-decreasing sequence. However, it is easy to transform $\mathbf{S}$ to an equivalent Turing machine $\mathbf{S}'$, that outputs the rationals in a non-reduced form satisfying the constraint. Using this transformation, the result follows from Theorems 4.4 and 4.6. ∎

Theorem 4.10: *Let Φ be a vocabulary containing at least three binary predicates, and let $r, r_1, r_2 \in [0,1]$ be rational numbers such that $r_1 \le r_2$. For sentences $\varphi, \theta \in \mathcal{L}^{-}(\Phi)$, given an oracle for deciding if $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta)$) exists,*

(a) the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta) = r$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) = r$) is Π_2^0-complete,

(b) if $[r_1, r_2] \ne [0,1]$, then the problem of deciding whether $\Diamond\Box\Pr_\infty^w(\varphi|\theta) \in [r_1, r_2]$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) \in [r_1, r_2]$) is Π_2^0-complete,

(c) if $r_1 \ne r_2$, then the problem of deciding if $\Diamond\Box\Pr_\infty^w(\varphi|\theta) \in (r_1, r_2)$ (resp., $\Box\Diamond\Pr_\infty^w(\varphi|\theta) \in (r_1, r_2)$) is Σ_2^0-complete.

Proof: It can be verified that the sentences constructed in the proof of Theorem 4.5 satisfy the constraints of Theorem 4.6. ∎

We can trivially obtain analogues to the results in this section for the random-structures method by adding one more binary predicate to the language (not used in the relevant formulas), and using Proposition 2.6. For example, we can show that if Φ contains at least four non-unary predicate symbols, then for $\varphi, \theta \in \mathcal{L}^{-}(\Phi)$, the problem of deciding if $\Diamond\Box\Pr_\infty^{s,\Phi}(\varphi|\theta)$ (or $\Box\Diamond\Pr_\infty^{s,\Phi}(\varphi|\theta)$) exists is Π_3^0-complete. Similar analogues to the other results in that section also hold. Details are left to the reader.

5 Is there any hope?

These results show that most interesting problems regarding asymptotic probabilities are badly undecidable in general. Are there restricted sublanguages for which these questions become tractable, or at least decidable?

All of our negative results so far depend on having at least one non-unary predicate symbol in the vocabulary. In fact, it clearly suffices to have the non-unary predicate symbols appear only in θ. However, as we indicated in the introduction, this additional expressive power of θ is essential. If we restrict θ to refer only to unary predicates and constants, many of the problems we encounter in the general case disappear. This holds even if φ can refer to arbitrary predicates. In the companion paper [GHK93] we focus on this important special case. Here, we consider one other case.

A close look at our proofs in the previous sections shows that we typically started by constructing sentences of low quantification depth, that use (among other things) an unbounded number of unary predicates. For example, the original construction of the sentences encoding computations of Turing machines used a unary predicate for every state of the machine. We then explained how to encode everything using only one binary predicate. In the process of doing this encoding, we had to introduce additional quantifiers (for example, an existential quantifier for every unary predicate eliminated). Thus, our undecidability results seem to require one of two things: an unbounded vocabulary (in terms of either the number of predicates or of their arity), or unbounded quantification depth. Do we really need both? It is actually easy to show that the answer is yes.

Definition 5.1: Define $d(\xi)$ to be the *depth of quantifier nesting* in the formula ξ:

- $d(\xi) = 0$ for any atomic formula ξ,
- $d(\neg\xi) = d(\xi)$,
- $d(\xi_1 \wedge \xi_2) = \max(d(\xi_1), d(\xi_2))$,
- $d(\forall y\, \xi) = d(\xi) + 1$. ∎

Let $\mathcal{L}_d(\Phi)$ consist of all sentences $\varphi \in \mathcal{L}(\Phi)$ such that φ has quantification depth at most d.

Theorem 5.2: *For all d, there exists a Turing machine $\mathbf{M}_d$ such that for all $\varphi, \theta \in \mathcal{L}_d(\Phi)$, $\mathbf{M}_d$ decides in time linear in the length of φ and θ whether $\Diamond\Box\Pr^w_\infty(\varphi|\theta)$ (resp., $\Box\Diamond\Pr^w_\infty(\varphi|\theta)$, $\Diamond\Box\Pr^{s,\Phi}_\infty(\varphi|\theta)$, $\Box\Diamond\Pr^{s,\Phi}_\infty(\varphi|\theta)$) is well defined, if so whether it exists, and if it exists computes an arbitrarily good rational approximation to its value.*

Proof: Let $\mathcal{L}_i^d(\Phi)$ consist of all formulas (not necessarily sentences) of quantification depth at most i that mention only the variables $x_1, \ldots, x_d$. Notice that there is an algorithm that runs in linear time that effectively converts a sentence in $\mathcal{L}_d(\Phi)$ to a sentence in $\mathcal{L}_d^d(\Phi)$. We now show that (a) we can effectively find a finite set Σ_i^d of formulas such that every formula in $\mathcal{L}_i^d(\Phi)$ is equivalent to a formula in Σ_i^d, and (b) there is a linear time algorithm that effectively converts a formula in $\mathcal{L}_i^d(\Phi)$ to an equivalent formula in Σ_i^d. This is sufficient to show that any problem—including all those relating to conditional probabilities—whose inputs are formulas in $\mathcal{L}_i^d(\Phi)$ and whose outputs only depend on the semantics of formulas, is solvable in linear time. This is because there exists a constant-time algorithm—essentially a lookup table—that, given a formula in Σ_i^d, outputs the correct response. So, given any formula, we can find the equivalent formula in Σ_i^d, and use this algorithm to obtain the appropriate output. Note that we cannot necessarily give an effective construction that produces the lookup table.

We first prove the existence of Σ_i^d for each fixed d by induction on i. For the base case $i = 0$, observe that our assumptions imply that there are only finitely many distinct "literals" consisting of a predicate symbol, followed by the appropriate number of arguments drawn from the constants in Φ and $x_1, \ldots x_d$. (For the purpose of this proof, we treat equality just like any other binary predicate.) Every formula in $\mathcal{L}_0^d(\Phi)$ is a Boolean combination of these literals, and there are only finitely many non-equivalent Boolean combinations of formulas in a finite set. We can effectively construct a set Σ_0^d consisting of one representative of each equivalence class of equivalent formulas. For later ease of exposition, we assume that if the equivalence class includes a literal, then that is the representative chosen to be in Σ_0^d.

For the inductive step, suppose that we have constructed Σ_i^d. Every formula in $\mathcal{L}_{i+1}^d(\Phi)$ is equivalent to a Boolean combination of formulas of the form $Qx_j\,\psi$, where $j \le d$, ψ has depth at most i, and Q is either $\exists, \forall$, or is absent altogether. By the inductive hypothesis, we can replace ψ by an equivalent formula $\sigma_\psi \in \Sigma_i^d$. Therefore, every formula in $\mathcal{L}_{i+1}^d(\Phi)$ is equivalent to a Boolean combination of formulas of the form $Qx_j\,\sigma_\psi$, where $j \le d$ and $\sigma_\psi \in \Sigma_i^d$. Since Σ_i^d is finite and $j \le d$, this is a Boolean combination of formulas in a finite set. Using the fact that there are only finitely many inequivalent Boolean combinations of formulas in a finite set, we can again construct a finite set Σ_{i+1}^d extending Σ_i^d for which the result follows.

To complete the proof, we need to show how to determine the appropriate $\sigma \in \Sigma_i^d$ given a sentence $\xi \in \mathcal{L}_i^d(\Phi)$. We assume that ξ is fully parenthesized. First, it is clear that there exists a constant time algorithm (a lookup table) such that: given a formula of the form $\sigma_1 \wedge \sigma_2$, $\neg\sigma_1$, or $\exists x_j\,\sigma_1$, for $\sigma_1, \sigma_2 \in \Sigma_i^d$, it finds an equivalent formula in Σ_i^d. This is easy to see because, as Σ_i^d is finite, there are only a finite number of possible inputs.

We now proceed by reading ξ from left to right, doing the following:

1. push all literals and operators (Boolean connectives and quantifiers) on a stack as they are encountered,
2. when we encounter a right parenthesis, pop the immediately preceding symbols off the stack, so as to obtain a subformula of the form $\sigma_1 \wedge \sigma_2$, $\neg\sigma_1$, or $\exists x_j\,\sigma_1$,
3. find the formula $\sigma \in \Sigma_i^d$ which is equivalent to the popped subformula,
4. push σ back onto the stack.

It is straightforward to prove by induction that in Step 2, the formulas σ_1 and σ_2 are both in Σ_i^d. The base case follows by our assumption about Σ_i^d containing all literals. The inductive step follows by the construction of the lookup table algorithm. Moreover, the subformula σ pushed onto the stack in Step 4 is logically equivalent to the formula it replaces. It follows that after ξ is read, there is exactly one formula on the stack, which is equivalent to ξ.

Given Φ and d, it is easy to construct Σ_i^d and a Turing machine that, for each pair of formulas $\varphi, \theta \in \mathcal{L}_i^d(\Phi)$, finds the equivalent formulas $\sigma_\varphi, \sigma_\theta \in \Sigma_i^d$. Given that, it remains only to construct a lookup table that tells us, for any formulas $\sigma_\varphi, \sigma_\theta \in \Sigma_i^d$, the behavior of $\Diamond\Box\Pr_\infty^w(\varphi|\theta)$ ($\Box\Diamond\Pr_\infty^w(\varphi|\theta)$, $\Diamond\Box\Pr_\infty^{s,\Phi}(\varphi|\theta)$, $\Box\Diamond\Pr_\infty^{s,\Phi}(\varphi|\theta)$). We can easily construct a finite set of linear-time Turing machines, corresponding to the different possible lookup tables. One of these will allow us to correctly determine the behavior of the asymptotic probability (well definedness, existence, and value of limit). ∎

The proof of the previous theorem says that, for each d, there exist lookup tables that effectively determine the behavior of the asymptotic probability for sentences in $\mathcal{L}_d(\Phi)$. Moreover, it shows that we can effectively construct a finite set of lookup tables, one of which is bound to be the right one. Unfortunately, we cannot effectively determine which one is the right one, for if we could, we could effectively construct $\mathbf{M}_d$ given Φ and d, and this would contradict our earlier undecidability results. Thus, even for this extremely restrictive sublanguage we cannot effectively construct algorithms for computing asymptotic conditional probabilities.

Acknowledgments: We would like to thank Fahiem Bacchus, with whom we started working on this general area of research. We would also like to thank Ron Fagin and Moshe Vardi for useful comments on a previous draft of this paper. Ron's comments encouraged us to discuss both the random structures and the random worlds methods in this paper. We would also like to thank Moshe for pointing out the reference [Lio69], and for suggesting the technique used to prove Theorem 4.6.

References

[BGHK94] F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of belief. Technical Report 9855, IBM, 1994. To appear, *Artificial Intelligence*. Available by anonymous ftp from logos.uwaterloo.ca/pub/bacchus or via WWW at http://logos.uwaterloo.ca. A preliminary version of this work appeared in *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)*, 1993, pages 563–569.

[BGHK95] F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller. Reasoning with noisy sensors in the situation calculus. In *Proc. Fourteenth International Joint Conference on Artificial Intelligence (IJCAI '95)*, pages 1933–1940, 1995.

[Car50] R. Carnap. *Logical Foundations of Probability*. University of Chicago Press, Chicago, 1950.

[Car52] R. Carnap. *The Continuum of Inductive Methods*. University of Chicago Press, Chicago, 1952.

[Che83] P. C. Cheeseman. A method of computing generalized Bayesian probability values for expert systems. In *Proc. Eighth International Joint Conference on Artificial Intelligence (IJCAI '83)*, pages 198–202, 1983.

[Com88] K. Compton. 0-1 laws in logic and combinatorics. In I. Rival, editor, *Proc. 1987 NATO Adv. Study Inst. on algorithms and order*, pages 353–383. Reidel, Dordrecht, Netherlands, 1988.

[DD85] K. G. Denbigh and J. S. Denbigh. *Entropy in Relation to Incomplete Knowledge*. Cambridge University Press, Cambridge, U.K., 1985.

[DG79] B. Dreben and W. D. Goldfarb. *The Decision Problem: Solvable Classes of Quantificational Formulas*. Addison-Wesley, Reading, Mass., 1979.

[Fag76] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58, 1976.

[Fag77] R. Fagin. The number of finite relational structures. *Discrete Mathematics*, 19:17–21, 1977.

[Gai60] H. Gaifman. Probability models and the completeness theorem. In *International Congress of Logic Methodology and Philosophy of Science*, pages 77–78, 1960. This is the abstract of which [Gai64] is the full paper.

[Gai64] H. Gaifman. Concerning measures in first order calculi. *Israel Journal of Mathematics*, 2:1–18, 1964.

[GHK92] A. J. Grove, J. Y. Halpern, and D. Koller. Asymptotic conditional probabilities for first-order logic. In *Proc. 24th ACM Symp. on Theory of Computing*, pages 294–305, 1992.

[GHK93] A. J. Grove, J. Y. Halpern, and D. Koller. Asymptotic conditional probabilities: the unary case. Research Report RJ 9563, IBM, 1993. To appear in *SIAM Journal on Computing*.

[GHK94] A. J. Grove, J. Y. Halpern, and D. Koller. Random worlds and maximum entropy. *Journal of A.I. Research*, 2:33–88, 1994.

[GKLT69] Y. V. Glebskiĭ, D. I. Kogan, M. I. Liogon'kiĭ, and V. A. Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Kibernetika*, 2:17–28, 1969.

[GKP89] R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete Mathematics—A Foundation for Computer Science*. Addison-Wesley, Reading, Mass., 1989.

[Jay78] E. T. Jaynes. Where do we stand on maximum entropy? In R. D. Levine and M. Tribus, editors, *The Maximum Entropy Formalism*, pages 15–118. MIT Press, Cambridge, Mass., 1978.

[Key21] J. M. Keynes. *A Treatise on Probability*. Macmillan, London, 1921.

[Kri86] J. von Kries. *Die Principien der Wahrscheinlichkeitsrechnung und Rational Expectation*. Freiburg, 1886.

[KV90] Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws and decision problems for fragments of second-order logic. *Information and Computation*, 87:302–338, 1990.

[Lap20] P. S. de Laplace. *Essai Philosophique sur les Probabilités*. 1820. English translation is *Philosophical Essay on Probabilities*, Dover Publications, New York, 1951.

[Lew79] H. R. Lewis. *Unsolvable Classes of Quantificational Formulas*. Addison-Wesley, New York, 1979.

[Lio69] M. I. Liogon'kiĭ. On the conditional satisfiability ratio of logical formulas. *Mathematical Notes of the Academy of the USSR*, 6:856–861, 1969.

[Lyn85] J. Lynch. Probabilities of first-order sentences about unary functions. *Trans. American Mathematical Society*, 287:543–568, 1985.

[MP92] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*, volume 1. Springer-Verlag, Berlin/New York, 1992.

[PS72] R. E. Powell and S. M. Shah. *Summability Theory and Applications.* Van Nostrand Reinhold, 1972.

[PV89] J. B. Paris and A. Vencovska. On the applicability of maximum entropy to inexact reasoning. *International Journal of Approximate Reasoning*, 3:1–34, 1989.

[Rog67] H. Rogers, Jr. *Theory of Recursive Functions and Effective Computability.* McGraw-Hill, New York, 1967.

[Tra50] B. A. Trakhtenbrot. Impossibility of an algorithm for the decision problem in finite classes. *Doklady Akademii Nauk SSSR*, 70:569–572, 1950.