Finite model theory

6. Zero-One Laws

Zero-One Laws

In this section we adopt the convention that a finite model $\mathfrak{M}$ with $|M| = n$ has the domain $\{1, \ldots, n\}$. Thus this assumption will be systematically followed below, mainly without further comment.

Zero-One Laws

Let τ be a finite relational vocabulary. We let $\mathcal{M}_\tau$ denote the class of all finite τ-models; in this section this means τ-models with domains of type $\{1, \ldots, n\}$. A **property of τ-models** is an isomorphically closed class $\mathcal{P} \subseteq \mathcal{M}_\tau$, i.e., a class such that if $\mathfrak{M} \in \mathcal{P}$ and $\mathfrak{N}$ is a τ-model such that $\mathfrak{N} \cong \mathfrak{M}$, then $\mathfrak{N} \in \mathcal{P}$.

We let $\mathcal{M}_\tau(n)$ denote the class of models $\mathfrak{M} \in \mathcal{M}_\tau$ of size n, i.e., the models $\mathfrak{M} \in \mathcal{M}_\tau$ with domain $M = \{1, \ldots, n\}$. If $\mathcal{P}$ is a property of τ-models, then we let $\mathcal{P}(n)$ denote the models $\mathfrak{M} \in \mathcal{P}$ of size n.

Zero-One Laws

Let $\mathcal{P}$ be a property of τ-models. The probability $\mu_n(\mathcal{P})$ of $\mathcal{P}$ over models of size n is given by

$$\mu_n(\mathcal{P}) = \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n)|}.$$

That is, $\mu_n(\mathcal{P})$ is the probability that a randomly chosen τ-model of size n has property $\mathcal{P}$.

Definition 6.1

The **asymptotic probability $\mu(\mathcal{P})$ of $\mathcal{P}$** is given by

$$\mu(\mathcal{P}) = \lim_{n\to\infty} \mu_n(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n)|}.$$

Zero-One Laws

Example 6.2

Let τ be the vocabulary $\{P\}$ with a unary relation symbol P. Consider the property $\mathcal{P}$ that there is precisely one element in P, i.e., $|P| = 1$. The probability of this happening in a τ-model of size n is

$$\mu_n(\mathcal{P}) = \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n)|} = \frac{n}{2^n}.$$

Thus the asymptotic probability $\mu(\mathcal{P})$ of this happening when the model size increases is

$$\mu(\mathcal{P}) = \lim_{n\to\infty} \mu_n(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n)|} = \frac{n}{2^n} = 0.$$

Zero-One Laws

Example 6.3

Let τ be an arbitrary finite relational vocabulary. Consider the property $\mathcal{P}$ that the domain of the model is even. The probability of this happening in a τ-model of size n is

$$\mu_n(\mathcal{P}) = \begin{cases} 1 & \text{if } n \text{ is even,} \\ 0 & \text{if } n \text{ is odd.} \end{cases}$$

Thus the asymptotic probability

$$\mu(\mathcal{P}) = \lim_{n\to\infty} \mu_n(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n)|}$$

of the property $\mathcal{P}$ does not exist.

Zero-One Laws

Consider the vocabulary $\tau = \{P\}$ with one unary relation P. Find a property $\mathcal{P}$ of τ-models for which the asymptotic probability is $\frac{1}{2}$. Prove your solution correct.

Zero-One Laws

One such property is PARITY consisting of those τ-models where $|P|$ is even. (The domain does not have to be even, just $|P|$.) To prove this, define the following bijection $f : \mathcal{M}_\tau(n) \to \mathcal{M}_\tau(n)$ that pairs up models of size n into pairs $(\mathfrak{M}, \mathfrak{N})$ such that $f(\mathfrak{M}) = \mathfrak{N}$ and $f(\mathfrak{N}) = \mathfrak{M}$.

Define $f(\mathfrak{M}) = \mathfrak{N}$ if and only if for all $i \in \{1, \ldots, n\}$,

$$i \in P^{\mathfrak{M}} \setminus P^{\mathfrak{N}} \quad \text{or} \quad i \in P^{\mathfrak{N}} \setminus P^{\mathfrak{M}},$$

That is $\mathfrak{N}$ is obtained from $\mathfrak{M}$ by defining $P^{\mathfrak{N}}$ to be the complement relation of $P^{\mathfrak{M}}$.

When n is odd, it is clear that $\mathfrak{M}$ satisfies parity iff $f(\mathfrak{M})$ does not, so exactly half of models of size n satisfy parity. When n is even, we can consider models with the even domain $\{1, \ldots, n-1\}$ and note that half of them satisfy PARITY, and furthermore, each such model can be extended to a model with domain $\{1, \ldots, n\}$ in precisely two ways, one changing the parity status and one keeping it as it is. Thus half of the models with an even domain satisfy PARITY.

Zero-One Laws

Definition 6.4

A logic $\mathcal{L}$ is said to have the **zero-one law** if every property of τ-models definable (with respect to the class of finite τ-models) by a τ-sentence of $\mathcal{L}$ has asymptotic probability of 0 or 1.

Zero-One Laws

Let τ be a finite relational vocabulary. Let $\mathcal{C} \subseteq \mathcal{M}_\tau$ be some subclass of the class $\mathcal{M}_\tau$ of all finite τ-models. Let $\mathcal{P}$ be a property of τ-models and recall $\mathcal{P}(n)$ denoes the subclass of $\mathcal{P}$ containing the models of size n. The probability $\mu_n^{\mathcal{C}}(\mathcal{P})$ of $\mathcal{P}$ with respect to models of size n in $\mathcal{C}$ is given by

$$\mu_n^{\mathcal{C}}(\mathcal{P}) = \frac{|\mathcal{P}(n) \cap \mathcal{C}|}{|\mathcal{M}_\tau(n) \cap \mathcal{C}|}.$$

That is, $\mu_n^{\mathcal{C}}(\mathcal{P})$ is the probability of the property $\mathcal{P}$ being realized if we randomly choose a model of size n from $\mathcal{C}$.

Definition 6.5

The **asymptotic probability $\mu^{\mathcal{C}}(\mathcal{P})$ of $\mathcal{P}$ with respect to $\mathcal{C}$** is given by

$$\mu^{\mathcal{C}}(\mathcal{P}) = \lim_{n\to\infty} \mu_n^{\mathcal{C}}(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n) \cap \mathcal{C}|}{|\mathcal{M}_\tau(n) \cap \mathcal{C}|}.$$

Zero-One Laws

Recall that a graph is a finite $\{E\}$-model where E is an irreflexive and symmetric binary relation. Let $\mathcal{G}$ be the class of graphs. A graph has an isolated node if there exists an element v such that $(v, u) \notin E$ (and thus $(u, v) \notin E$) holds for all elements u. Let $\mathcal{P}$ be the graph property of having an isolated node, i.e., $\mathcal{P}$ is the class of graphs with an isolated node. Figure out the asymptotic probability $\mu^{\mathcal{G}}(\mathcal{P})$ of a graph having an isolated node.

Zero-One Laws

The asymptotic probability $\mu^{\mathcal{G}}(\mathcal{P})$ of this happening when the graph domain size increases is

$$\mu^{\mathcal{G}}(\mathcal{P}) = \lim_{n\to\infty} \mu_n^{\mathcal{G}}(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n) \cap \mathcal{G}|}$$

Recall here that $\mathcal{P}$ was already a class of graphs, so $\mathcal{P}(n) = \mathcal{P}(n) \cap \mathcal{G}$. Let us work this limit out.

Zero-One Laws

Let $I(n)$ be the number of graphs of domain size n with an isolated node. For convenience and simplicity, suppose now that $n \geq 4$. Over a domain of n nodes, there are n ways of choosing an isolated node and $2^{\binom{n-1}{2}}$ ways of choosing the edges between the remaining nodes. This implies that

$$I(n) \leq n \cdot 2^{\binom{n-1}{2}},$$

where we note that $n \cdot 2^{\binom{n-1}{2}}$ of course gives an overestimation, as the same graph more than one isolated node will be counted multiple times. Thereby

$$\mu^{\mathcal{G}}(\mathcal{P}) = \lim_{n\to\infty} \mu^{\mathcal{G}}_n(\mathcal{P}) = \lim_{n\to\infty} \frac{|\mathcal{P}(n)|}{|\mathcal{M}_\tau(n) \cap \mathcal{G}|} \leq \frac{n \cdot 2^{\binom{n-1}{2}}}{2^{\binom{n}{2}}}$$

$$= n \cdot 2^{\frac{(n-1)!}{(n-3)!2!} - \frac{n!}{(n-2)!2!}} = n \cdot 2^{\frac{(n-1)!(n-2)}{(n-2)!2!} - \frac{n!}{(n-2)!2!}}$$

$$= n \cdot 2^{\frac{-2(n-1)}{2}} = \frac{n}{2^{n-1}}.$$

Zero-One Law

As $\lim_{n\to\infty} \dfrac{n}{2^{n-1}} = 0$, the asymptotic probability $\mu^{\mathcal{G}}(\mathcal{P})$ of having an isolated node is zero. This makes sense, as in a large graph, it is indeed unlikely for an individual node u not to be linked to any other one of the (large number of) other nodes.

Zero-One Laws

Definition 6.6
Let $\mathcal{C}$ be some class of finite τ-models for some finite relational vocabulary τ. A logic $\mathcal{L}$ is said to have the **zero-one law with respect to** $\mathcal{C}$ if every property of τ-models definable with respect to $\mathcal{C}$ by a τ-sentence of $\mathcal{L}$ has asymptotic probability of 0 or 1 with respect to $\mathcal{C}$.

Zero-One Laws

Definition 6.7

Let $k \in \mathbb{Z}_+$ and $\ell, m \in \mathbb{N}$. Let $Ext_{\ell,m}$ denote the first-order $\{E\}$-sentence

$$\forall x_1 \dots \forall x_\ell \forall y_1 \dots \forall y_m \Big(\bigwedge_{1 \le i \le \ell} \bigwedge_{1 \le j \le m} x_i \neq y_j$$
$$\to \exists z \Big(\bigwedge_{1 \le i \le \ell} (z \neq x_i \wedge E(z, x_i)) \wedge \bigwedge_{1 \le j \le m} (z \neq y_j \wedge \neg E(z, y_j)) \Big) \Big).$$

The k-**extension axiom** is the formula

$$Ext_k = \bigwedge_{0 \le \ell \le k} Ext_{\ell,k-\ell}.$$

Note that this is a formula of the finite variable logic FVL^{k+1}.

Zero-One Law

Informally, the extension axiom Ext_k states that for any subsets A and B that do not overlap (i.e., $A \cap B = \emptyset$) and cover at most k nodes (i.e., $|A| + |B| \leq k$), there exists a point $v \notin A \cup B$ that connects via E to every node in A and to no node of B. Thereby v can be seen as an observation point that sees precisely everything in A and absolutely nothing in B. And there exists such an observation point for every two non-overlapping sets A and B that cover at most k points.

Lemma 6.8

Let $\mathbb{G}$ and $\mathbb{H}$ be a graphs that both satisfy Ext_k, i.e, we have $\mathbb{G} \models Ext_k$ and $\mathbb{H} \models Ext_k$. Then the duplicator has a winning strategy in the $(k+1)$-pebble game PG'_{k+1}.

Proof. Suppose the duplicator has survived all the way up to a stage S where the pebble variables $v_{i_1}, \ldots, v_{i_r}$ have been placed on both graphs, $r \leq k+1$. Let the spoiler place the pebble variable v_i (possibly $v_i \in \{v_{i_1}, \ldots, v_{i_r}\}$) in one of the graphs so that in that graph, v_i connects to the pebble variables $u_1, \ldots, u_\ell$ via E and does not connect to $z_1, \ldots, z_m$ via E. Here $\{u_1, \ldots, u_\ell\} \cup \{z_1, \ldots, z_m\} = \{v_{i_1}, \ldots, v_{i_r}\}$ if v_i has not been used in the earlier rounds, and $\{u_1, \ldots, u_\ell\} \cup \{z_1, \ldots, z_m\} = \{v_{i_1}, \ldots, v_{i_r}\} \setminus \{v_i\}$ if v_i has been used in the previous rounds. We have the following two cases.

Zero-One Laws

1. If the spoiler places v_i on a currently already chosen element, the duplicator responds by choosing the corresponding already chosen element in the other graph.
2. Otherwise, because the graphs satisfy the extension axiom Ext_k, the duplicator can respond so that v_i in the other graph connects via E to $u_1, \ldots, u_\ell$ and does not connect to $z_1, \ldots, z_m$.

This conclude the proof.$\square$

We note that the extension axiom has been custom made for enabling the duplicator to win the pebble game PG'_{k+1}.

Zero-One Laws

Lemma 6.9

The asymptototic probability of the extension axiom Ext_k is 1 (with respect to the class of graphs).

Proof. Consider graphs of size n with $n \geq k+1$. Recall the domain of these graphs is $\{1, \ldots, n\}$. The probability that there is an edge between two different elements $a, b \in \{1, \ldots, n\}$ is $\frac{1}{2}$. For Ext_k to fail, there must be some sets $\{a_1, \ldots, a_\ell\}$ and $\{b_1, \ldots, b_m\}$ of elements $(\ell + m \leq k)$ in the domain and so that for all point d outside the sets, we have that

- d fails to link to all elements of the first set, or
- d links to some element of the the second set.

The two sets should not have points in common, and d should indeed be external to both of the two sets. Certainly there are less than $n^\ell \cdot n^m$ ways of choosing the sets $\{a_1, \ldots, a_\ell\}$ and $\{b_1, \ldots, b_m\}$ (we are not aiming at minimal bounds, but instead are satisfied with gross overestimations). And there are no more than $n - (\ell + m)$ ways of choosing d. The probability that a fixed d links to the two sets so that Ext_k fails is $(1 - (\dfrac{1}{2})^\ell (\dfrac{1}{2})^m)$.

Zero-One Laws

Thus the probability that there exist such sets of sizes ℓ and m with no suitable d to satisfy Ext_k is surely not greater than

$$\begin{aligned}
&n^{\ell} \cdot n^{m} \cdot (1 - (\frac{1}{2})^{\ell}(\frac{1}{2})^{m})^{n-(\ell+m)} \\
=&\left(n^{\ell} \cdot n^{m} \cdot (1 - (\frac{1}{2})^{\ell}(\frac{1}{2})^{m})^{-(\ell+m)}\right) \cdot (1 - (\frac{1}{2})^{\ell}(\frac{1}{2})^{m})^{n} \\
=&P(n) \cdot a^{n}
\end{aligned}$$

where

$$P(n) = \left(n^{\ell} \cdot n^{m} \cdot (1 - (\frac{1}{2})^{\ell}(\frac{1}{2})^{m})^{-(\ell+m)}\right)$$

is a polynomial in n and a is the term $(1-(\frac{1}{2})^{\ell}(\frac{1}{2})^{m})$. Clearly $0 \leq a < 1$, so the term a^n (which is exponential in n and thus dominates the polynomial term $P(n)$) ensures this probability goes to zero as n approaches infinity.

We fixed ℓ and m so that $\ell + m \leq k$. Since k is a fixed constant (unlike n which we let go to infinity), there is a fixed number of ways we can choose ℓ and m, and in all of those cases, the above probability goes to zero. Thus the axiom Ext_k has asymptotic probability zero of failing.$\square$

Zero-One Law

Theorem 6.10

The finite variable logic FVL has the zero-one law over the class of graphs.

Proof. Let φ be a formula of FVL. Thus φ contains a finite number of variables. Suppose there are (at most) $k+1$ variables in φ, so φ is a formula of FVL^{k+1}. We discuss two cases.

Suppose that there exists a graph $\mathbb{G}$ that satisfies Ext_k and φ, i.e., $\mathbb{G} \models Ext_k \wedge \varphi$. As $\mathbb{G} \models Ext_k$, for every graph $\mathbb{H}$ such that $\mathbb{H} \models Ext_k$, the duplicator wins the pebble game $\mathrm{PG}'_{k+1}(\mathbb{G}, \mathbb{H})$ by Lemma 6.8. Now, by Lemma 6.9, the asymptotic probability of Ext_k is 1, so the proportion of those graphs $\mathbb{H}$ such that the duplicator wins $\mathrm{PG}'_{k+1}(\mathbb{G}, \mathbb{H})$ approaches 1 as the domain size of $\mathbb{H}$ approaches infinity. When the duplicator wins $\mathrm{PG}'_{k+1}(\mathbb{G}, \mathbb{H})$, then $\mathbb{G}$ and $\mathbb{H}$ must agree on formulae of FVL^{k+1}, including φ. Thus the asymptotic probability of φ is 1.

Zero-One Law

Suppose then that there exists no graph $\mathbb{G}$ such that $\mathbb{G} \models \mathit{Ext}_k \wedge \varphi$. Therefore, for all graphs $\mathbb{H}$ that satisfy Ext_k, we have $\mathbb{H} \not\models \varphi$. Thus the proportion of graphs that do not satisfy φ approaches 1 as the graph size grows. Thus the asymptotic probability of φ is 0.$\square$

Zero-One Law

Corollary 6.11

FO *and* LFP *have the zero-one law with respect to the class of graphs.*

Proof. Recall that LFP (and thus FO) translates into FVL. $\square$

Zero-One Law

The general zero-one law of FVL is proved similarly to the case restricticted to graphs. One just needs to take into account the general relational vocabulary. The details are not different from the case for graphs in any significant way. We omit the full details.

Theorem 6.12
FVL *has the zero-one law.*

Corollary 6.13
LFP *and thus* FO *have the zero-one law.*

Zero-One Law

Problem: Show that PARITY cannot be defined in LFP. Solution: LFP has the zero-one law and the asymptotic probability of parity is one half.