

Limiting Cases for Spectrum Closure Results

AARON HUNTER
SIMON FRASER UNIVERSITY
BURNABY, CANADA, V5A 1S6
amhunter@cs.sfu.ca

Received by Greg Restall

Published October 26, 2004

<http://www.philosophy.unimelb.edu.au/ajl/2004>

© 2004 Aaron Hunter

Abstract: The spectrum of a first-order sentence is the set of cardinalities of its finite models. Given a spectrum S and a function f , it is not always clear whether or not the image of S under f is also a spectrum. In this paper, we consider questions of this form for functions that increase very quickly and for functions that increase very slowly. Roughly speaking, we prove that the class of all spectra is closed under functions that increase arbitrarily quickly, but it is not closed under some natural slowly increasing functions.

I INTRODUCTION

I.1 SPECTRUM CLOSURE RESULTS

The *spectrum* of a first-order sentence ϕ is the set of cardinalities of its finite models. Let $\text{Sp}(\phi)$ denote the spectrum of ϕ and let SPEC denote the set of all spectra of first-order sentences. Given a set S of natural numbers and a function f , we write $f(S)$ as an abbreviation for $\{f(n) : n \in S\}$. Let $\mathcal{S} \subseteq \text{SPEC}$. We say that \mathcal{S} is *closed under* f if $f(S) \in \mathcal{S}$ whenever $S \in \mathcal{S}$.

In this paper, we study closure results for SPEC under functions that increase very rapidly and functions that increase very slowly. Our main result is that SPEC is closed under functions that increase arbitrarily quickly, but it is not closed under some natural slowly increasing functions. We prove all of our results using elementary model-theoretic constructions.

I.2 DEFINITIONS AND NOTATION

We assume that the reader is familiar with basic model theory as set out, for example, in [2]. We briefly review some basic terminology.

A finite set of non-logical symbols together with an arity for each symbol is a *vocabulary*. All vocabularies will be assumed to be relational. A structure for

a vocabulary \mathcal{L} is a finite set called the *universe*, together with an interpretation of each relation symbol in \mathcal{L} . By convention, M is used to denote the universe of the structure \mathcal{M} and $R^{\mathcal{M}}$ is used to denote the interpretation of R in the structure \mathcal{M} . We use the notation $\mathcal{M} \upharpoonright A$ to denote the restriction of \mathcal{M} to the set $A \subseteq M$.

The notation \bar{a} will be used as an abbreviation for a tuple (a_1, \dots, a_n) when the value of n is understood or unimportant. If $\phi(\bar{x})$ is a formula with free variables \bar{x} , then we write $\mathcal{M} \models \phi[\bar{a}]$ to indicate that \bar{a} satisfies $\phi(\bar{x})$ in \mathcal{M} . For a sentence ϕ , we simply write $\mathcal{M} \models \phi$ and we say that \mathcal{M} is a *model* of ϕ .

A set of first-order structures Σ is an *elementary class* if there is a sentence ψ such that $\Sigma = \{\mathcal{M} : \mathcal{M} \models \psi\}$. Following standard conventions, ω denotes the set of natural numbers and n denotes the set $\{0, \dots, n-1\}$. Given two sets A and B , ${}^A B$ denotes the set of functions from A into B .

1.3 BACKGROUND AND MOTIVATION

It is well-known that, if we think of spectra as sets of binary strings, then SPEC is the set NE of languages that are accepted in exponential time by a non-deterministic Turing machine [3, 6]. Due to the equivalence of SPEC and NE, the study of first-order spectra provides an alternative route from which one can approach difficult problems in complexity theory. One illustrative example of this connection is Asser's problem, which asks if SPEC is closed under complement [1]. If Asser's problem has a negative solution, then it follows by a straightforward argument that $P \neq NP$.

Relatively little is known about the class of spectra, beyond the close connection with NE. Proving closure results for natural subclasses of spectra is one way in which we can improve our understanding of the structure of SPEC. We briefly summarize some of the previous work on closure results.

Some closure results for SPEC follow directly from the complexity-theoretic characterization. For example, NE is clearly closed under the function $n \mapsto n^2$, so SPEC is also closed under this function. However, using the equivalence of SPEC and NE, it is not possible to prove closure results for some natural subclasses of spectra. If a class \mathcal{S} of spectra does not have a nice complexity-theoretic characterization, then we must use some other method for proving closure results for \mathcal{S} .

Many techniques used to prove closure results can be traced back to Fagin's work [4]. Let \mathcal{F}_k denote the class of spectra of sentences involving relation symbols of arity at most k . Fagin proves that, if $S \in \mathcal{F}_k$, then $\{n^k : n \in S\}$ is the spectrum of a sentence involving a single binary relation symbol. This is stronger than a typical closure result, because it actually maps a large class of spectra into a subset of itself. Fagin proves this result by using an explicit transformation from k -ary structures to binary structures.

Explicit transformations on structures have been employed more recently to prove some new closure results. In particular, More proves that \mathcal{F}_2 is closed

under every polynomial with rational coefficients that is asymptotically greater than the identity [7]. Given this result, it is natural to ask what happens if we look at subdiagonal functions. The results in this case are less concrete. For example, it has been shown [5] that \mathcal{F}_2 is closed under $n \mapsto \sqrt{n}$ if and only if $\mathcal{F}_2 = \text{SPEC}$.

In this paper, we ask if there are any limitations on the class of functions for which we can expect to prove closure results for SPEC. We propose that the natural way to approach this problem is from the perspective of model theory, so we prove all of our results using elementary logical techniques.

1.4 SUMMARY OF RESULTS

In §2, we are concerned with the closure of SPEC under rapidly increasing functions. In order to show that SPEC is closed under functions that increase arbitrarily quickly, we first prove that spectra can be very sparse. In particular, given any recursive function f , we prove that there will be a spectrum that is more sparse than $f(\omega)$. It follows from this result that SPEC is closed under functions that increase arbitrarily quickly.

In §3, we turn our attention to slowly increasing functions. It is already known that there is a spectrum S such that $\{n : 2^n \in S\}$ is not a spectrum [3]. We give a new model-theoretic proof of this result. Using our techniques, it may be possible to extend this result to many natural slowly increasing functions, provided that they have some finite iterate that is greater than $n \mapsto 2^n$. In order to prove our main result in this section, we use a Gödel numbering on ordered structures to provide a constructive, diagonalization proof that there is a sentence δ such that $\{n : 2^{2^n} \in \text{Sp}(\delta)\} \notin \text{SPEC}$.

We believe that this approach could have further application in the study of spectra.

2 SPARSE SPECTRA

For the purposes of our first proof, a Turing machine T is a triple (Q, Σ, δ) where

1. $Q = \{q_0, \dots, q_p\}$ is the set of *states*, such that q_0 is the *start state* and q_p is the *final state*.
2. $\Sigma = \{1, B\}$ is the set of *tape symbols*
3. $\delta : Q \times \Sigma \rightarrow Q \times \Sigma \times \{L, R\}$ is the *next move function*.

We emphasize that δ is a total function, so the machines have exactly one instruction for every pair $(q_i, a) \in Q \times \Sigma$. We will only consider Turing machines that have one head and operate on a two-way infinite tape.

A Turing machine is in *standard initial configuration* if it satisfies the following conditions:

1. it is in state q_0
2. for some $k \in \omega, k \neq 0$, there are k adjacent squares on the tape containing the symbol Γ and every other square on the tape contains B
3. the head is scanning the leftmost square on the tape that contains the symbol Γ .

A machine is in *standard final configuration* if it is in state q_p and it satisfies conditions 2 and 3. Given a Turing machine T and a function f over the positive natural numbers, we say that T *computes* f if, whenever T is started in a standard initial configuration with n Γ 's on the tape, T halts in a standard final configuration with $f(n)$ Γ 's on the tape. A function f is *recursive* if there is some Turing machine T that computes f . For technical reasons we require that the Turing machines we use never access any square to the left of the input. It is not hard to see that every recursive function can be computed by such a machine.

Let f be a function over the positive natural numbers and let $S \subseteq \omega$ where s_0, s_1, \dots is an enumeration of S such that $s_i < s_{i+1}$ for all i . We say that S *majorizes* f if $s_i > f(i)$ for all i .

THEOREM 1 If f is a recursive function, there is a spectrum that majorizes f .

PROOF Suppose that f is recursive, and let $T = (Q, \Sigma, \delta)$ compute f . Given n , let $t(n)$ denote the number of steps before T halts when started in standard initial configuration with n Γ 's on the tape. We choose T such that, for all n , $t(n) \geq f(n)$ and $t(n) \geq n$; clearly if f is recursive then it is computed by some Turing machine with these properties. Let $\mathcal{L} = \{F_1, F_2, \prec, P, Q_0, \dots, Q_p\}$ be a vocabulary where F_1, F_2 , and \prec are binary relation symbols and P, Q_0, \dots, Q_p are unary relation symbols. For each n , we construct an \mathcal{L} -structure \mathcal{M}_n that intuitively represents a computation of T with input n . The universe M_n is the set $t(n) \times t(n)$ and the relations F_1, F_2 and \prec are defined as follows:

$$\begin{aligned} F_1^{\mathcal{M}_n}(a, b)(c, d) &\iff a = b = c, \\ F_2^{\mathcal{M}_n}(a, b)(c, d) &\iff a = b = d, \\ (a, b) \prec^{\mathcal{M}_n} (c, d) &\iff a = b < c = d. \end{aligned}$$

Given $m < t(n)$, we think of $\{(i, m) : i < t(n)\}$ as a representation of the tape of T after m moves when started with input n .¹ For all $a, b < t(n)$, $P^{\mathcal{M}_n}(a, b)$ if and only if, when T is started in a standard initial configuration with n Γ 's on the tape, at the b^{th} step, the a^{th} square (counting from the left) contains the symbol Γ . For all $a, b < t(n)$ and $i \leq t$, $Q_i^{\mathcal{M}_n}(a, b)$ if and only if, when T is started in a standard initial configuration with n Γ 's on the tape, at the b^{th} step, the head is scanning the a^{th} square (counting from the left) and T is in state q_i .

¹We are actually representing the tape segment that begins at the starting square and extends $t(n)$ squares to the right.

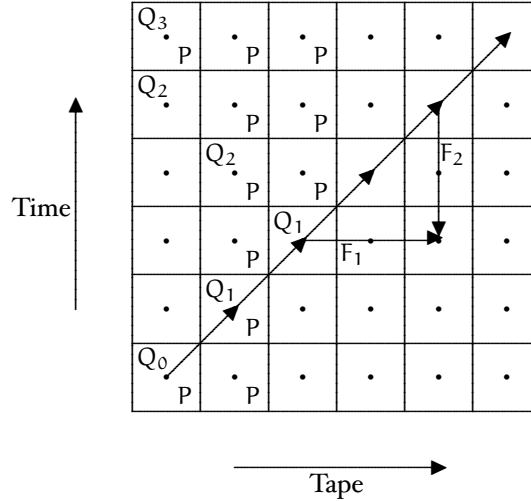


Figure 1: Representing a Computation

Figure 1 provides an illustration of the structure \mathcal{M}_2 which is obtained from a four state Turing machine T which computes $n \mapsto n + 1$ in a natural way. The bottom left point represents $(0, 0)$, the transitive closure of the diagonal arrows is \prec and all unary relations that apply to each point are enclosed in the surrounding square.

We construct a sentence ϕ_T such that $\mathcal{A} \models \phi_T$ if and only if $\mathcal{A} \cong \mathcal{M}_n$ for some n . The sentence ϕ_T is a conjunction of several sentences, the first of which is the sentence Λ defined as follows:

$$\begin{aligned} & \forall w \exists ! x \exists ! y (F_1 x x \wedge F_1 y y \wedge F_1 x w \wedge F_2 y w \\ & \quad \wedge \forall z ((F_1 x z \wedge F_2 y z) \rightarrow w = z)) \\ & \wedge \forall x \forall y ((F_1 x x \wedge F_1 y y) \rightarrow \exists ! z (F_1 x z \wedge F_2 y z \\ & \quad \wedge \forall u \forall v ((F_1 u z \wedge F_2 v z) \rightarrow (u = x \wedge v = y))). \end{aligned}$$

We remark that, in every model of Λ , each element of the universe is uniquely identified with a pair of elements of the set $\{a : F_1 a a\}$. More specifically, if r is an element of a model \mathcal{M} of Λ , then there is a unique pair $(b, c) \in \{a : F_1^{\mathcal{M}} a a\}^2$ such that $F_1^{\mathcal{M}} b r$ and $F_2^{\mathcal{M}} c r$.

Let ϕ_1 be a first-order sentence that asserts that \prec is a strict linear order on $\{a : F_1 a a\}$. For convenience, we let $x \prec_1 y$ abbreviate the following formula:

$$\exists z (F_1 z z \wedge F_1 z x \wedge F_1 z y) \wedge \exists u v (F_1 u u \wedge F_1 v v \wedge F_2 u x \wedge F_2 v y \wedge u \prec v).$$

The intuitive meaning of $a \prec_1 b$ is that, for some m , a and b both represent squares on the tape of T after m moves, and a is to the left of b . Similarly,

we define a formula $x \prec_2 y$ that is satisfied by (a, b) just in case a and b intuitively represent the same square on the tape of T at different stages of the computation, and a represents the square at an earlier stage. Define $x \prec_2 y$ as follows:

$$\exists z(F_1zz \wedge F_2zx \wedge F_2zy) \wedge \exists uv(F_1uu \wedge F_1vv \wedge F_1ux \wedge F_1vy \wedge u \prec_1 v).$$

The formula $x \preceq_i y$ will abbreviate $x \prec_i y \vee x = y$ and the formula $SUC_i xy$ will abbreviate $x \prec_i y \wedge \neg \exists z(x \prec_i z \wedge z \prec_i y)$.

We now define two sentences ϕ_2 and ϕ_3 that intuitively assert that T starts in a standard initial configuration and it halts in a standard final configuration. The sentence ϕ_2 is defined as follows:

$$\begin{aligned} \exists x \forall y (\neg y \prec_1 x \wedge \neg y \prec_2 x \wedge Q_0x \wedge \bigwedge_{i=1}^p \neg Q_i x \wedge (x \prec_1 y \rightarrow \bigwedge_{i=0}^p \neg Q_i y)) \\ \wedge \exists z (x \prec_1 z \wedge \forall v ((v \prec_1 z \rightarrow Pv) \wedge (z \preceq_1 v \rightarrow \neg Pv))). \end{aligned}$$

Let ϕ_3 be the following sentence:

$$\begin{aligned} \exists x \forall y (\neg y \prec_1 x \wedge \neg x \prec_2 y \wedge Q_p x \wedge \bigwedge_{i=0}^{p-1} \neg Q_i x \wedge (x \prec_1 y \rightarrow \bigwedge_{i=0}^p \neg Q_i y)) \\ \wedge \exists z (x \prec_1 z \wedge \forall v ((v \prec_1 z \rightarrow Pv) \wedge (z \preceq_1 v \rightarrow \neg Pv))). \end{aligned}$$

For each $(q_i, a) \in Q \times \Sigma$, we need a sentence $\psi_{(q_i, a)}^T$ that ensures that the immediately succeeding tape configuration is correct whenever the state is q_i and the scanned symbol is a in some row. We illustrate the construction with an example. Suppose that, for some $i \leq t$, $\delta(q_i, 1) = (q_j, 0, L)$. We define $\psi_{(q_i, 1)}^T$ to be the following sentence:

$$\begin{aligned} \forall x ((Q_i x \wedge Px) \rightarrow \exists y z (SUC_2 xz \wedge SUC_1 yz \wedge \neg Pz \wedge Q_j y \wedge \\ \forall v ((v \neq y \wedge (y \prec_1 v \vee v \prec_1 y)) \\ \rightarrow (\bigwedge_{i=0}^p \neg Q_i v \wedge \exists u ((u \prec_1 x \vee x \prec_1 u) \wedge u \prec_2 v \wedge (Pu \leftrightarrow Pv)))))). \end{aligned}$$

Similarly, given $(q_i, a) \in Q \times \Sigma$ one has an appropriate sentence $\psi_{(q_i, a)}^T$.

The sentence ϕ_T is the conjunction

$$\bigwedge_{i=1}^3 \phi_i \wedge \bigwedge_{(q_i, a) \in Q \times \Sigma} \psi_{(q_i, a)}^T.$$

We now prove that $\mathcal{A} \models \phi_T$ if and only if $\mathcal{A} \cong \mathcal{M}_n$ for some n . It follows immediately from the construction that $\mathcal{M}_m \models \phi_T$ for every m , so clearly $\mathcal{A} \models \phi_T$ whenever there is an n such that $\mathcal{A} \cong \mathcal{M}_n$. Suppose that $\mathcal{A} \models \phi_T$. Since

$\mathcal{A} \models \Lambda$, $\mathcal{A} = m^2$ for some m and we can define the function $g : \mathcal{A} \rightarrow m \times m$ as follows:

$$g(a) = \begin{cases} (i, i) & \text{if } F_1^{\mathcal{A}}aa \text{ and } a \text{ is } i^{\text{th}} \text{ in the order } \prec^{\mathcal{A}}; \\ (i, j) & \text{if } F_1^{\mathcal{A}}ba, F_2^{\mathcal{A}}ca \text{ and } b \text{ and } c \text{ are } i^{\text{th}} \text{ and } j^{\text{th}} \text{ in the order } \prec^{\mathcal{A}}. \end{cases}$$

Clearly g is a bijection. Since $\mathcal{A} \models \phi_2$, there is some n such that $P^{\mathcal{A}}g^{-1}(i, 0)$ if and only if $i < n$. It is easy to see that g can be extended to an isomorphism from \mathcal{A} to \mathcal{M}_n . Therefore, $\text{Sp}(\phi_{\top}) = \{t^2(n) : n \in \omega\}$. It is clear that this set majorizes f . $\#$

As a consequence of Theorem 1, we can show that SPEC is closed under functions that increase as fast as we like.

COROLLARY 1 For any recursive function f , there is a function g such that

- $g(n) > f(n)$ for all n
- $g(S) \in \text{SPEC}$ for all $S \in \text{SPEC}$.

PROOF Let f be computed by T and let $S = \text{Sp}(\psi)$. Let \mathcal{M}_n , t and ϕ_{\top} be as in Theorem 1. Define $\theta(x)$ to be the following formula:

$$\exists y \forall z (\neg z \prec_1 y \wedge \neg z \prec_2 y \wedge y \prec_1 x \wedge Px).$$

Observe that, if $\mathcal{M}_n \models \theta[a]$, then $(0, 0)$ is the only point that can witness the variable y . Intuitively, $\mathcal{M}_n \models \theta[a]$ if and only if a represents a square on the initial tape of T that contains a 1. Observe that, for any n , there are exactly n distinct points a in \mathcal{M}_n such that $\mathcal{M}_n \models \theta[a]$. Therefore, by Theorem 1, $\mathcal{A} \models \phi_{\top} \wedge \psi^{\theta(x)}$ if and only if $\mathcal{A} \cong \mathcal{M}_n$ for some $n \in \text{Sp}(\phi)$. Hence $\text{Sp}(\phi_{\top} \wedge \psi^{\theta(x)}) = t^2(S)$. Clearly $t^2(n) > f(n)$ for all n , so t^2 satisfies the conditions for g in the statement of the corollary. $\#$

3 SLOWLY INCREASING FUNCTIONS

3.1 MOTIVATION

We now turn our attention towards closure results under slowly increasing functions. Our original goal was simply to provide a model-theoretic proof of Fagin's result that there exists $S \in \text{SPEC}$ such that $\{n : 2^n \in S\} \notin \text{SPEC}$ [3]. What we actually prove is a more general theorem, from which Fagin's result follows.

The intuition behind our proof is simple. We construct a Gödel numbering on sentences that associates a natural number $G(\phi)$ with each sentence ϕ . Our Gödel numbering actually assigns numbers to ordered structures, but for the moment we look at a simplified version of the argument. Next, we construct a sentence δ such that every model of δ has size 2^{2^n} for some n . Moreover, the sentence δ is defined such that every model of δ essentially finds the sentence ϕ with $G(\phi) = n$ and it makes sure that ϕ does not have a model of size n . This leads to a diagonalization proof of Fagin's non-closure result.

The key point in our construction is that, if the models of δ are sufficiently large with respect to n and $G(\phi)$, then we can use them to check whether or not ϕ has a model of size n . There is nothing particularly important about the cardinality 2^{2^n} .

In subsequent sections we use $(\forall x_0, \dots, x_t \in A)\psi$ and $(\exists x_0, \dots, x_t \in A)\psi$ as abbreviations for:

$$\begin{aligned} &\forall x_0 \dots x_t ((Ax_0 \wedge \dots \wedge Ax_t) \rightarrow \psi) \\ &\exists x_0 \dots x_t (Ax_0 \wedge \dots \wedge Ax_t \wedge \psi). \end{aligned}$$

We hope this convention improves the readability of some of the longer sentences.

3.2 OUTLINE OF THE PROOF

The usual formulation of first-order logic provides a canonical list v_0, v_1, \dots of variables. Let ψ be a formula involving a relation symbol R of indeterminate arity. We say that ψ is in *extended form* if it satisfies the following conditions:

1. there exists k such that each atomic subformula of ψ involving R has the form $Rv_{ik} \dots v_{(i+1)k-1}$ for some i
2. each atomic subformula of ψ involving R occurs only once in ψ
3. every variable in ψ occurs in some atomic subformula involving R .

Informally, a formula in extended form is simply a formula in which the variable indices are sequential in each atomic subformula involving R . Equality subformulas are still allowed, with no restriction on the variables. We remark that every $\{R\}$ -formula is equivalent to some formula in extended form. We illustrate how to convert a formula into extended form with a simple example.

EXAMPLE 1 Consider the formula

$$\forall x \exists y (Rxy \wedge Ryy).$$

The first step towards creating an equivalent extended formula is to replace the atomic components with Rv_0v_1 and Rv_2v_3 . The variable corresponding to x needs to be universally quantified, and the variables corresponding to y need to be existentially quantified. Moreover, we need to add equality formulas to assure that the variables corresponding to y are witnessed by the same object. We get the following equivalent extended formula:

$$\forall v_0 \exists v_1 v_2 v_3 [(Rv_0v_1 \wedge Rv_2v_3) \wedge (v_1 = v_2 \wedge v_2 = v_3)].$$

We now give a brief outline of the proof of our main result. First, we define a vocabulary \mathcal{L}_s and an injective mapping that associates an ordered \mathcal{L}_s -structure \mathcal{M}_ϕ with each extended $\{R\}$ -formula ϕ , where the arity of R runs through all positive integers. The injectiveness is strong in the sense that, if ϕ and ϕ' are distinct, then $\mathcal{M}_\phi \not\cong \mathcal{M}_{\phi'}$. The variables occurring in ϕ are in 1-1 correspondence with the subset of \mathcal{M}_ϕ picked out by the unary relation symbol V . In a particular formula ϕ , all the occurrences of the relation symbol R have the same arity k . The elements of \mathcal{M}_ϕ corresponding to the first k variables are picked out by the unary relation symbol K of \mathcal{L}_s . From the details of the construction of \mathcal{M}_ϕ , it will be clear that

$$\{\mathcal{M} : \mathcal{M} \cong \mathcal{M}_\phi \text{ for some extended } \{R\}\text{-formula } \phi\}$$

is an elementary class.

Second, we define a Gödel numbering G that associates a unique natural number $G(\mathcal{M})$ with each ordered structure \mathcal{M} .

In §3.6, we prove the following theorem.

THEOREM 2 There is a sentence δ such that, for every $\{R\}$ -formula ϕ in extended form,

$$2^{2^{G(\mathcal{M}_\phi)}} \in \text{Sp}(\delta) \text{ if and only if } G(\mathcal{M}_\phi) \notin \text{Sp}(\bar{\phi}),$$

where $\bar{\phi}$ denotes the universal closure of ϕ .

Every model of δ is isomorphic to a model \mathcal{N} which is associated with an extended $\{R\}$ -formula in the following way. Suppose the extended $\{R\}$ -formula is ϕ and the only variables occurring, free or bound, in ϕ are v_0, \dots, v_{t-1} . The vocabulary of δ contains unary relation symbols A, B, C, V and G which have the following significance²:

1. $(\mathcal{N} \upharpoonright A, \mathcal{L}_s) = \mathcal{M}_\phi$
2. V picks out the elements of \mathcal{M}_ϕ corresponding to the variables which occur in ϕ
3. K picks out the elements of \mathcal{M}_ϕ corresponding to the first k variables, where k is the arity of R in ϕ
4. $G = n = G(\mathcal{M}_\phi)$
5. $B = \{f : f \text{ is a function from } V^{\mathcal{N}} \text{ to } n\}$
6. $C = \mathcal{P}(B)$, where $\mathcal{P}(B)$ denotes the power set of B .

²We omit the superscript \mathcal{N} on the interpretations of the relation symbols in the list.




Figure 2: The Basic Picture

Intuitively, we think of B^N as the set of t -tuples of elements of n and of C^N as the set of all sets of t -tuples of n .

The vocabulary of δ contains a ternary relation symbol T_1 such that

$$T_1^N bvg \iff b(v) = g$$

for all $b \in B^N$, $v \in V^N$ and $g \in G^N$. There is also a binary relation symbol T_2 in the vocabulary of δ such that

$$T_2^N cb \iff b \in c$$

for all $c \in C$ and $b \in B$. The basic picture is illustrated in Figure 2.

There is a sentence δ_1 such that any model of δ_1 is isomorphic to a structure which is related to some extended $\{R\}$ -formula ϕ in the manner specified above. Conversely, for each extended $\{R\}$ -formula ϕ , there is an associated model of δ_1 . Further, δ_1 is only concerned with the restriction of the model to the union of the interpretations of A , B , C and G . More precisely, if $\mathcal{N} \models \delta_1$ and

$$\mathcal{N}' \upharpoonright (A^N \cup B^N \cup C^N \cup G^N) = \mathcal{N} \upharpoonright (A^N \cup B^N \cup C^N \cup G^N)$$

then $\mathcal{N}' \models \delta_1$. From the discussion below it will be clear that

$$|A^N \cup B^N \cup C^N \cup G^N| \leq 2^{2^n}.$$

Next, there is a sentence δ_2 involving the unary relation symbol G such that any model of δ_2 has cardinality 2^{2^n} where n denotes the cardinality of the

interpretation of G . We can choose δ_2 so that G is the only symbol that occurs in the intersection of the vocabularies of δ_1 and δ_2 . Furthermore, δ_2 places no restriction on the size of the interpretation of G .

>From 2, there are elements of $a_{v_0}, \dots, a_{v_{t-1}} \in \mathbb{N}$ corresponding to the variables v_0, \dots, v_{t-1} which occur in ϕ . In fact, $V = \{a_{v_0}, \dots, a_{v_{t-1}}\}$ and $K = \{a_{v_0}, \dots, a_{v_{k-1}}\}$. With each point $c \in C^{\mathbb{N}}$ define R_c as follows:

$$(g_0, \dots, g_{k-1}) \in R_c \iff \text{there is } f \in c \text{ such that } f(a_{v_i}) = g_i, \text{ for each } i < k.$$

Clearly, R_c runs through all k -ary relations on n as c runs through C . The structure \mathcal{M}_ϕ is defined in such a way that we can include relations in the vocabulary of δ which allow us to find a sentence τ and a formula $\sigma(x)$ such that, uniformly for all models \mathcal{N} of $\delta_1 \wedge \tau$ and $c \in C^{\mathbb{N}}$, $\mathcal{N} \models \sigma[c]$ if and only if $\langle n, R_c \rangle \models \phi_{\mathcal{N}}$.³ The sentence δ whose existence is asserted in Theorem 2 will be the conjunction

$$\delta_1 \wedge \delta_2 \wedge \tau \wedge \forall x \neg \sigma(x).$$

Intuitively this sentence says

“ G picks out a set whose cardinality is the Gödel number n of the structure representing some extended formula ϕ , the whole model has cardinality 2^{2^n} and $n \notin \text{Sp}(\bar{\phi})$ ”

We proceed as follows:

1. define the mapping $\phi \mapsto \mathcal{M}_\phi$ (§3.3)
2. define the Gödel numbering G and state a basic lemma (§3.4)
3. construct δ_1 and δ_2 (§3.5)
4. prove the main result (§3.6).

After reading §3.3, one who is willing to accept the existence of suitable G , δ_1 and δ_2 may choose to skip ahead to §3.6.

3.3 STRUCTURES REPRESENTING FORMULAS

Let ϕ be an $\{R\}$ -formula in extended form involving the variables v_0, \dots, v_{t-1} and let $\mathcal{L}_s = \{V, K, N, L, <, E, R_v, R_\wedge, R_\neg, R_\exists\}$ where N, L , and V are unary relation symbols and all of the other relation symbols are binary. We construct an \mathcal{L}_s -structure \mathcal{M}_ϕ where the universe M_ϕ is given by the following conditions:

1. for each $i < t$, there is a unique point $a_{v_i} \in M_\phi$
2. for each subformula $\psi(\bar{v})$ of ϕ , there is a unique point $a_\psi \in M_\phi$
3. for any v_i and ψ , $a_{v_i} \neq a_\psi$

³We have added a subscript \mathcal{N} to ϕ to emphasize that ϕ depends on \mathcal{N} whereas σ does not.

4. nothing else is in \mathcal{M}_ϕ .

The relation $<_1^{\mathcal{M}_\phi}$ is a total linear ordering. For definiteness, we define it to be the ordering where $a_\theta <_1^{\mathcal{M}_\phi} a_\gamma$ if and only if the string of symbols that denotes θ precedes the string of symbols that denotes γ in the lexicographic ordering obtained from the following ordering of symbols:

$$\langle v_0 \rangle < \langle v_1 \rangle < \dots < \langle v_{t-1} \rangle < \langle (\rangle < \langle) \rangle < \langle R \rangle < \langle = \rangle < \langle \wedge \rangle < \langle \neg \rangle < \langle \exists \rangle$$

The interpretations of the other symbols in \mathcal{L}_1 are defined as follows⁴:

$$\begin{aligned} \forall a &\iff a = a_{v_i} \text{ for some } i < t \\ \text{Ka} &\iff a = a_{v_i} \text{ for some } i < k \\ \text{Na} &\iff a = a_\psi \text{ where } \psi \text{ is } R\bar{v} \text{ for some } \bar{v} \\ \text{La} &\iff a = a_\psi \text{ where } \psi \text{ is } v_i = v_j \text{ for some } i, j < t \\ \text{Eab} &\iff a = a_{v_{kp+i}} \text{ and } b = a_{v_i}, R \text{ is } k\text{-ary and } i < k \\ \text{R}_v \text{ab} &\iff a = a_{v_i} \text{ and } b = a_\psi \text{ where } \psi \text{ is atomic, and } v_i \text{ occurs in } \psi \\ \text{R}_\wedge \text{ab} &\iff a = a_\gamma \text{ and } (b = a_{(\gamma \wedge \theta)} \text{ or } b = a_{(\theta \wedge \gamma)}) \text{ for some } \gamma \text{ and } \theta \\ \text{R}_\neg \text{ab} &\iff a = a_\gamma \text{ and } b = a_{(\neg \gamma)} \text{ for some } \gamma \\ \text{R}_\exists \text{ab} &\iff a = a_\gamma \text{ and } b = a_{(\exists v_i(\gamma))} \text{ for some } v_i \text{ and } \gamma \\ &\text{or } a = a_{v_i} \text{ and } b = a_{(\exists v_i(\gamma))} \text{ for some } v_i \text{ and } \gamma. \end{aligned}$$

As an illustration, Figure 3 shows \mathcal{M}_ϕ for the sentence $\exists v_0 v_1 (Rv_0 v_1 \wedge v_0 = v_1)$. All edges in Figure 3 should be interpreted as arrows from the higher point to the lower point. The unary relations and the edges of $<_1^{\mathcal{M}_\phi}$ have been omitted for clarity.

We note that, given the isomorphism type of \mathcal{M}_ϕ , it is easy to recover ϕ .

It is tedious, but not difficult, to construct a sentence μ_0 such that $\mathcal{M} \models \mu_0$ if and only if $\mathcal{M} \cong \mathcal{M}_\phi$ for some $\{R\}$ -formula ϕ in extended form. The details of the construction of μ_0 are left to the reader.

3.4 A GÖDEL NUMBERING OF ORDERED STRUCTURES

Fix a vocabulary $\mathcal{L} = \{<, R_1, \dots, R_t\}$ where $<$ is binary and R_i is k -ary for each i . Let \mathcal{M} denote a structure where $<^{\mathcal{M}}$ is a total linear ordering. Define $(R_i^{\mathcal{M}})^* \subseteq |M|^k$ as follows:

$$j \in (R_i^{\mathcal{M}})^* \iff \text{the } j^{\text{th}} \text{ element of } M^k \text{ in the lexicographic order obtained from } <^{\mathcal{M}} \text{ is in } R_i^{\mathcal{M}}.$$

The *Gödel number* $G(\mathcal{M})$ of \mathcal{M} is defined to be

$$2^{k|M|} \left(1 + \sum_{i \in (R_1^{\mathcal{M}})^*} 2^{1+i} + \sum_{i \in (R_2^{\mathcal{M}})^*} 2^{|M|^k+1+i} + \dots + \sum_{i \in (R_t^{\mathcal{M}})^*} 2^{(t-1)|M|^k+1+i} \right).$$

⁴For ease of readability, the superscript \mathcal{M}_ϕ has been omitted.




Figure 3: A Structural Representation of a Formula

We remark that G is essentially injective in the sense that, whenever $\mathcal{M}_1 \not\cong \mathcal{M}_2$, it follows that $G(\mathcal{M}_1) \neq G(\mathcal{M}_2)$. This is true because, given a sum of distinct powers of 2, it is always possible to determine which powers occur in the sum.

We make the following observation about the Gödel numbering G over \mathcal{L} . The proof is straightforward and has been omitted.

LEMMA 1 Let Γ be an elementary class over \mathcal{L} such that each $\mathcal{A} \in \Gamma$ assigns a strict linear ordering of A to the symbol $<$. Let A and G be new unary relation symbols. There is an $\mathcal{L}^* \supseteq \mathcal{L} \cup \{A, G\}$ and an elementary class Γ^* over \mathcal{L}^* such that

1. if $\mathcal{M}^* \in \Gamma^*$, then
 - a. \mathcal{M}^* is a disjoint union of $A^{\mathcal{M}^*}$ and $G^{\mathcal{M}^*}$
 - b. for each $P \in \mathcal{L}$, $P^{\mathcal{M}^*} \subseteq (A^{\mathcal{M}^*})^k$
 - c. $(\mathcal{M}^* \upharpoonright A^{\mathcal{M}^*}, \mathcal{L}) \in \Gamma$
 - d. $|G^{\mathcal{M}^*}| = G(\mathcal{M}^* \upharpoonright A^{\mathcal{M}^*}, \mathcal{L})$
2. if $\mathcal{M} \in \Gamma$, then there exists $\mathcal{M}^* \in \Gamma^*$ such that $\mathcal{M} = (\mathcal{M}^* \upharpoonright A^{\mathcal{M}^*}, \mathcal{L})$

If $\mathcal{M}^* \in \Gamma^*$, then \mathcal{M}^* has the form pictured in Figure 4.

Notice that we have stated Lemma 1 for a vocabulary in which every relation symbol has the same arity, but we intend to apply it to a sentence in the vocabulary \mathcal{L}_s in which there are unary and binary relation symbols. This is




Figure 4: The Structure \mathcal{M}^*

not a problem because there is a natural way in which unary relation symbols can be simulated by binary relation symbols, and hence we can think of \mathcal{L}_s as a vocabulary in which all of the relation symbols are binary.

3.5 CONSTRUCTION OF δ_1 AND δ_2

The sentence δ_1 is a conjunction of three sentences, the first of which is obtained from Lemma 1 and the others are constructed explicitly in this section.

Recall that every model of μ_0 is isomorphic to \mathcal{M}_ϕ for some formula ϕ , so the set

$$\{\mathcal{M} : \mathcal{M} \cong \mathcal{M}_\phi \text{ for some extended } \{R\}\text{-formula } \phi\}$$

is an elementary class over \mathcal{L}_s . By Lemma 1, there is a vocabulary $\mathcal{L}_s^* \supseteq \mathcal{L}_s$ and an \mathcal{L}_s^* -sentence μ_1 that satisfies the following conditions:

1. if $\mathcal{M}^* \models \mu_1$, then for some extended $\{R\}$ -formula ϕ
 - a. \mathcal{M}^* is a disjoint union of $A^{\mathcal{M}^*}$ and $G^{\mathcal{M}^*}$
 - b. $(\mathcal{M}^* \upharpoonright A^{\mathcal{M}^*}, \mathcal{L}_s) \cong \mathcal{M}_\phi$
 - c. $|G^{\mathcal{M}^*}| = G(\mathcal{M}_\phi)$
2. if $\mathcal{M} \cong \mathcal{M}_\phi$ for some extended $\{R\}$ -formula ϕ , then there exists $\mathcal{M}^* \in \Gamma^*$ such that $\mathcal{M}_\phi = (\mathcal{M}^* \upharpoonright A^{\mathcal{M}^*}, \mathcal{L}_s)$.

Next, we define a $\{B, G, V, T_1\}$ -sentence μ_2 that intuitively asserts that B is the set of functions from V to G and $T_1 bvg$ if and only if $b \in B$, $v \in V$, $g \in G$ and $b(v) = g$. We construct μ_2 so that it makes the following informal assertions:

1. B is non-empty
2. the relation T_1 is a subset of $B \times V \times G$
3. each $b \in B$ maps each $v \in V$ to a unique point in G via the relation T_1
4. if $b, b' \in B$ and $b \neq b'$, then there is some $v \in V$ such that $b(v) \neq b'(v)$
5. if $b \in B, v \in V, g \in G$ and $b(v) \neq g$ then there is a function $f \in B$ which is the same as b except that $f(v) = g$.

The key point is that, given a non-empty set B of functions from V to G , if B contains the “immediate neighbours” of each of its elements, then $B = {}^V G$. Let μ_2 denote the following sentence:

$$\begin{aligned} & \exists x(Bx) \wedge \forall xyz(T_1xyz \rightarrow (Bx \wedge Vy \wedge Gz)) \\ & \wedge (\forall x \in B)(\forall y \in V)(\exists!z \in G)(T_1xyz) \\ & \wedge (\forall x, y \in B)(x \neq y \rightarrow (\exists u \in V)(\exists v, w \in G)(T_1xuv \wedge T_1xuw \wedge v \neq w)) \\ & \wedge (\forall x \in B)(\forall y \in V)(\forall z \in G)(\exists w \in B)(\forall u \in V)(\forall v \in G) \\ & ((u \neq y \vee v \neq z) \rightarrow (T_1xuv \leftrightarrow T_1wuv)) \wedge (T_1xyz \leftrightarrow \neg T_1wyz)). \end{aligned}$$

It is easy to see that $\mathcal{M} \models \mu_2$ if and only if there exists \mathcal{A} such that

1. $\mathcal{A} \cong (\mathcal{M} \upharpoonright (B^{\mathcal{M}} \cup G^{\mathcal{M}} \cup V^{\mathcal{M}}), \{B, G, V, T_1\})$
2. $B^{\mathcal{A}}$ is the set of functions from $V^{\mathcal{A}}$ to $G^{\mathcal{A}}$
3. $T_1^{\mathcal{A}}bvg$ if and only if $b \in B^{\mathcal{A}}, v \in V^{\mathcal{A}}, g \in G^{\mathcal{A}}$ and $b(v) = g$.

The third sentence required is a $\{B, C, T_2\}$ -sentence that intuitively asserts that C is the power set of B and T_2cb if and only if $c \in C, b \in B$ and $b \in c$. We construct μ_3 so that it so that it makes the following assertions:

1. $C^{\mathcal{M}}$ is non-empty
2. the relation T_2 is a subset of $C \times B$
3. if $c, c' \in C$ and $c \neq c'$, then there exists b such that b is an element of exactly one of c and c'
4. if $c \in C$ and $b \in B$, then there exists $c' \in C$ which has the same elements as c except that $b \in c'$ if and only if $b \notin c$.

Let μ_3 denote the following sentence:

$$\begin{aligned} & (\exists x \in C)(\forall y \in B)(\neg T_2xy) \\ & \wedge \forall xy(T_2xy \rightarrow (Cx \wedge By)) \\ & \wedge (\forall x, y \in C)(\exists z \in B)(T_2xz \leftrightarrow \neg T_2yz) \\ & \wedge (\forall x \in C)(\forall y \in B)(\exists z \in C)(\forall w \in B) \\ & ((w \neq y \rightarrow (T_2xw \leftrightarrow T_2zw)) \wedge (T_2xy \leftrightarrow \neg T_2zy)). \end{aligned}$$

It is easy to see that $\mathcal{M} \models \mu_3$ if and only there is some n such that

$$\langle \mathcal{P}(n), n, \in \rangle \cong \left(\mathcal{M} \upharpoonright (B^{\mathcal{M}} \cup C^{\mathcal{M}}), \{B, C, T_2\} \right).$$

Let δ_1 be the conjunction $\mu_1 \wedge \mu_2 \wedge \mu_3$. It is clear that $\mathcal{M} \models \delta_1$ if and only if \mathcal{M} is isomorphic to a structure \mathcal{N} which is related to an extended $\{R\}$ -formula in the manner described in §3.2.

Using the same idea that was used to define μ_3 , it is easy to define a sentence δ_2 such that the only relation symbol from \mathcal{L}_s which occurs in δ_2 is G and, whenever $\mathcal{M} \models \delta_2$, the cardinality of M is $2^{2^{|G^{\mathcal{M}}|}}$.

3.6 THE MAIN RESULT

Let ϕ be an extended $\{R\}$ -formula involving t variables, and let \mathcal{N} be a structure that is related to ϕ in the way described in §3.2. Recall that R_c denotes the relation

$$\{(g_0, \dots, g_{k-1}) : \text{for some } f \in c, f(a_{v_i}) = g_i \text{ for each } i < k\}.$$

We construct a formula $\sigma(x)$ such that

$$\mathcal{N} \models \sigma[c] \iff \langle n, R_c \rangle \models \phi$$

where n denotes the size of the set picked out by the predicate G .

It is convenient to think of elements of B as t -tuples of elements of G , so for the rest of this section we let (g_0, \dots, g_{t-1}) denote the function $f : V \rightarrow G$ such that $f(a_{v_i}) = g_i$ for each $i < t$.

Let P be a ternary relation symbol that is interpreted by \mathcal{N} and does not occur in δ_1 or δ_2 . The role of P is critical, so we take a moment to explain the underlying intuition. We think of each $c \in C$ as a k -ary relation over n , which we have denoted by R_c . Recall that each element $a \in A$ corresponds either to a variable in ϕ , or it corresponds to a subformula of ϕ . Hence, we can think of each element $a \in A - V$ as a subformula a_ψ of ϕ . Finally, we have identified each element $\bar{b} \in B$ with a t -tuple over n . The relation P intuitively holds for a triple (R_c, a_ψ, \bar{b}) if and only if the subformula a_ψ is satisfied by \bar{b} when R_c is the interpretation of R . Slightly more formally, we want to define P on $C \times (A - V) \times B$ such that

$$Pca_\psi(g_0, \dots, g_{t-1}) \iff \langle n, R_c \rangle \models \psi[g_0, \dots, g_{t-1}]. \quad (i)$$

We remark that it follows from this condition that $Pca_\psi b$ holds for some c, b if and only if ϕ has a model of size n . We construct a first-order sentence τ that will ensure that P satisfies condition (i).

We will proceed by cases, based on the middle argument of P . The middle argument is an element of $A - V$, so it corresponds to a subformula of ϕ . Hence, we will first construct a sentence τ_N that ensures that (i) holds whenever the

middle argument corresponds to an atomic $\{R\}$ -formula. Next, we construct similar sentences for equality formulas, conjunctions, negations and existentials. In the end, we let τ be the conjunction of the sentences corresponding to each case, and we can prove by induction that (i) must hold for every model of τ . We will demonstrate the construction of τ_N in detail to illustrate the idea, and we will sketch the remaining cases with less explanation.

We now construct τ_N . Let ψ be the formula $R(v_{ik}, \dots, v_{(i+1)k-1})$. In this case, $\langle n, R_c \rangle \models \psi[\bar{g}]$ if and only if R_c holds for the subsequence of \bar{g} with indices from ik to $(i+1)k-1$. In other words,

$$\langle n, R_c \rangle \models \psi[g_0, \dots, g_{t-1}] \iff (g_{ik}, \dots, g_{(i+1)k-1}) \in R_c.$$

Hence, we will define τ_N such that

$$Pc a_\psi(g_0, \dots, g_{t-1}) \iff (g_{ik}, \dots, g_{(i+1)k-1}) \in R_c. \quad (2)$$

The idea is simply to use the vocabulary already present in N to ensure that this condition must hold. Recall that N picks out the set of points in A which correspond to atomic $\{R\}$ -formulas. Thus, we would like a sentence giving necessary and sufficient conditions for $Pc a_\psi$, quantifying over all $c \in C$, $a \in N$, and $b \in B$. The sentence intuitively asserts that there is a tuple d in c that agrees with b at all indices of variables present in the formula corresponding to a . We define τ_N as follows.

$$\begin{aligned} & (\forall c \in C)(\forall a \in N)(\forall b \in B) \\ & (Pc a_\psi \leftrightarrow \\ & \quad \exists d(T_2 c d \wedge (\forall v, u \in V)((R_v v a \wedge E v u) \\ & \quad \rightarrow (\exists w \in G)(T_1 b v w \wedge T_1 d u w))). \end{aligned}$$

Essentially τ_N says that there is some relation $d \in c$ that agrees with b on each pair (u, v) satisfying $(R_v v a \wedge E v u)$. Pairs that satisfy this condition are pairs of variables where u represents v_i for some $i < k$ and v represents some variable with an index that is equivalent to i modulo k . This variable translation is required to assure that the indices of the variables in ψ don't matter, but it is easier to read the sentence if we ignore this condition. If we suppose that ψ is $Rv_0 \dots v_{k-1}$, the sentence simply asserts that the tuple b is equivalent to one in c . It is straightforward to verify that this sentence is true in N if and only if (2) holds.

Next we construct a sentence τ_L that ensures that (i) holds for equality formulas. Suppose that ψ is the formula $v_i = v_j$. In this case, we want

$$Pc a_\psi(g_0, \dots, g_{t-1}) \iff g_i = g_j.$$

Let τ_L denote the sentence

$$\begin{aligned} & (\forall c \in C)(\forall a \in L)(\forall b \in B) \\ & (Pc a_\psi \leftrightarrow \forall v w((R_v v a \wedge R_v w a) \rightarrow \exists u(T_1 b v u \wedge T_1 b w u))). \end{aligned}$$

If ψ is $\theta \wedge \gamma$, then the set of t -tuples that satisfy ψ in \mathcal{N} is just the intersection of the sets of t -tuples that satisfy θ and γ respectively. Hence, in order to ensure that (i) will hold for a_ψ , we construct a sentence τ_\wedge that is true in \mathcal{N} if and only if the following condition holds for all points in A corresponding to subformulas of the form $\theta \wedge \gamma$:

$$Pca_{\theta \wedge \gamma} \bar{g} \iff Pca_\theta \bar{g} \text{ and } Pca_\gamma \bar{g}.$$

Let τ_\wedge denote the sentence

$$\forall a(\exists xy(R_\wedge xa \wedge R_\wedge ya) \rightarrow (\forall c \in C)(\forall b \in B)(Pcab \leftrightarrow (Pcxb \wedge Pcyb))).$$

Similarly, we define τ_\neg and τ_\exists , respectively, as follows:

- $\forall a(\exists x(R_\neg xa) \rightarrow (\forall c \in C)(\forall b \in B)(Pcab \leftrightarrow \neg Pxab))$
- $\forall a(\exists xy(R_\exists xa \wedge R_\exists ya \wedge \forall y) \rightarrow (\forall c \in C)(\forall b \in B)(Pcab \leftrightarrow \exists d(Pcxd \wedge (\forall w \in V)(w \neq y \rightarrow \exists u(T_1 dwu \wedge T_1 bwu))))).$

It is easy to see that τ_\neg and τ_\exists intuitively make the appropriate assertions.

Let τ denote the conjunction $\tau_N \wedge \tau_L \wedge \tau_\wedge \wedge \tau_\neg \wedge \tau_\exists$. The following lemma can be proved by a straightforward induction on ψ .

LEMMA 2 Suppose $\mathcal{N} \models \delta_1 \wedge \delta_2 \wedge \tau$ and ϕ is the extended $\{R\}$ -formula associated with \mathcal{N} . For each subformula ψ of ϕ , for all $c \in C^{\mathcal{N}}$ and all $b = (g_0, \dots, g_{t-1}) \in B^{\mathcal{N}}$,

$$P^{\mathcal{N}}ca_\psi(g_0, \dots, g_{t-1}) \iff \langle G^{\mathcal{N}}, R_c \rangle \models \psi[g_0, \dots, g_{t-1}].^5$$

We construct the formula $\sigma(x)$ mentioned in §3.2. The formula $\sigma(x)$ intuitively states that $Pxa_\phi b$ for all $b \in B$. Define $\sigma(x)$ as follows:

$$(\exists u \in A)(\neg \exists y(R_\vee uy \vee R_\wedge uy \vee R_\neg uy \vee R_\exists uy) \wedge (\forall b \in B)(Pxab)).$$

Observe that, if $\mathcal{N} \models \delta_1 \wedge \sigma[c]$ and ϕ is the associated $\{R\}$ -formula, then the only point that can witness the existentially quantified variable u in σ is the point a_ϕ .

We are now in a position to prove the main result of this section.

THEOREM 2 There is a sentence δ such that, for every $\{R\}$ -formula ϕ in extended form,

$$2^{2^{G(\mathcal{M}_\phi)}} \in \text{Sp}(\delta) \text{ if and only if } G(\mathcal{M}_\phi) \notin \text{Sp}(\bar{\phi}),$$

where $\bar{\phi}$ denotes the universal closure of ϕ .

⁵The use of “ a_ψ ” is predicated upon \mathcal{N} being a model such that its restriction to A, \mathcal{L}_s is actually equal to \mathcal{M}_ϕ , not just isomorphic.

PROOF Let δ be the conjunction $\delta_1 \wedge \delta_2 \wedge \tau \wedge \forall x \neg \sigma(x)$ and let ϕ be an extended $\{R\}$ -formula. Say R has arity k and ϕ involves t variables v_0, \dots, v_{t-1} .

Suppose $2^{2^{G(\mathcal{M}_\phi)}} \in \text{Sp}(\delta)$. Say $\mathcal{M} \models \delta$ and $|\mathcal{M}| = 2^{2^{G(\mathcal{M}_\phi)}}$. Since $\mathcal{M} \models \delta_1$, there is an extended $\{R\}$ -formula ψ and a structure \mathcal{N} related to ψ as described in §3.2 such that $\mathcal{M} \cong \mathcal{N}$ and $G^{\mathcal{N}} = G(\mathcal{M}_\psi) = \{0, \dots, G(\mathcal{M}_\psi) - 1\}$. Since $\mathcal{M} \models \delta_2$, $|\mathcal{M}| = 2^{2^{G(\mathcal{M}_\psi)}}$. Therefore,

$$G(\mathcal{M}_\psi) = |G^{\mathcal{N}}| = |G^{\mathcal{M}}| = G(\mathcal{M}_\phi).$$

Since $\mathcal{A} \mapsto G(\mathcal{A})$ and $\gamma \mapsto \mathcal{M}_\gamma$ are both injective, ψ and ϕ are the same formula.

Since $\mathcal{N} \models \forall x \neg \sigma(x)$, for every $c \in C^{\mathcal{N}}$ there is some $\bar{g} \in B^{\mathcal{M}}$ such that not $P^{\mathcal{N}} c a_\phi \bar{g}$. Hence, by Lemma 2, for all $c \in C^{\mathcal{N}}$, $\langle G(\mathcal{M}_\phi), R_c \rangle \not\models \phi[\bar{g}]$ for some \bar{g} . By the definition of the satisfaction relation, it follows that $\langle G(\mathcal{M}_\phi), R_c \rangle \not\models \phi$. Since R_c runs through all k -ary relations on $G(\mathcal{M}_\phi)$, it follows that $G(\mathcal{M}_\phi) \notin \text{Sp}(\bar{\phi})$.

Suppose that $G(\mathcal{M}_\phi) \notin \text{Sp}(\bar{\phi})$. It is clear that we can construct a model \mathcal{N} that is related to ϕ in the manner specified in §3.2. By construction, $\mathcal{N} \models \delta_1$. Moreover, if $n = |G^{\mathcal{N}}|$ then

$$|A^{\mathcal{N}} \cup B^{\mathcal{N}} \cup C^{\mathcal{N}} \cup G^{\mathcal{N}}| \leq n + n^t + 2^{n^t} + n \leq 2^{2^n}.$$

We pause to explain the second inequality. The smallest possible value for $|\mathcal{M}_\phi| = |A^{\mathcal{N}}|$ is 2 and the smallest possible value for t is 1 because the smallest extended formula is Rv_0 . Therefore, from the definition of G , $n = G(\mathcal{M}_\phi) \geq 16$. Hence, since $n^t \geq 16$,

$$n + n^t + 2^{n^t} + n \leq 2^{n^t} + 3n^t < 2^{n^t+1}.$$

So it is sufficient to show that $2^{n^t+1} \leq 2^{2^n}$, which is the same as $n^t < 2^n$. Clearly this is equivalent to $t \log_2 n < n$, which is what we prove. From the definition of G , $n \geq 2^{2^{t+2}}$ since $|\mathcal{M}_\phi| \geq t+1$. It follows that $\log_2 n \geq 2t+2$ and hence $t < t+1 < \frac{\log_2 n}{2}$. Therefore, it is sufficient to show that $(\frac{\log_2 n}{2})(\log_2 n) < n$. This is equivalent to showing that $(\log_2 n)^2 < 2n$. By taking the derivative, it is clear that $2x - (\log_2 x)^2$ is increasing for $x \geq 4$. Furthermore, if $x = 16$ then $(\log_2 x)^2 < 2x$. Therefore, since $n \geq 16$, $(\log_2 n)^2 < 2n$ which establishes the desired inequality.

By the preceding inequality, we can choose \mathcal{N} such that $\mathcal{N} \models \delta_2$. Since P does not occur in δ_1 or δ_2 , we can choose \mathcal{N} so that it assigns any interpretation we like to the symbol P . Choose \mathcal{N} such that, for each subformula ψ of ϕ , (1) holds. It follows that $\mathcal{N} \models \tau$. Since

$$G(\mathcal{M}_\phi) = |G^{\mathcal{N}}| \notin \text{Sp}(\bar{\phi}),$$

it follows from Lemma 2 that $\mathcal{N} \models \forall x \neg \sigma(x)$. #

As a consequence of Theorem 2, we get an alternative proof of a well-known result of Fagin from [3]. Our proof bears little resemblance to the original.

COROLLARY 2 There exists $S \in \text{SPEC}$ such that $\{n : 2^n \in S\} \notin \text{SPEC}$.

PROOF Suppose that $\{n : 2^n \in S\} \in \text{SPEC}$ whenever $S \in \text{SPEC}$. It follows that

$$S \in \text{SPEC} \Rightarrow \{n : 2^{2^n} \in S\} \in \text{SPEC}.$$

Let $T = \text{Sp}(\delta)$ and suppose that $T^* = \{n : 2^{2^n} \in T\} \in \mathcal{F}_k$. It is straightforward to show that any spectrum in \mathcal{F}_k is also the spectrum of a sentence involving a single $k + 1$ -ary relation. Hence $T^* = \text{Sp}(\phi)$ for some sentence ϕ in extended form involving one relation symbol $\{R\}$. By Theorem 2, $G(\mathcal{M}_\phi) \notin \text{Sp}(\phi) = T^*$ if and only if $2^{2^{G(\mathcal{M}_\phi)}} \in \text{Sp}(\delta) = T$. This contradicts the definition of T^* . $\#$

The original proof of this result uses a machine-based argument that is significantly shorter than the proof of Theorem 2. The advantage of our proof is that it makes it clear that the only thing that is really important about the function $n \mapsto 2^n$ is the fact that 2^n is large with respect to n . Roughly speaking, any function f such that some finite iterate of f is larger than $n \mapsto 2^n$ should be sufficiently large to attempt our construction.

We remark that it may not be possible to work out the details for all functions that increase sufficiently quickly. Basically, the problem is that we need to define an analogue of δ_2 ; in other words, we need a sentence that relates the size of the universe to the size of the interpretation of G . Creating this kind of sentence is straightforward for many natural functions. In this paper, for example, we have shown that there is a sentence ψ involving a unary relation symbol G such that every finite model \mathcal{M} of ψ has cardinality $2^{|G|^\mathcal{M}}$. It is straightforward to construct similar sentences for many natural functions, such as $n \mapsto n^k$. However, it is not obvious that this is so easy for complex functions that do not follow a simple pattern. It would be nice to give a simple characterization of the functions which can be used in our construction.

4 CONCLUSION

In this paper, we have studied general closure properties for SPEC using model-theoretic techniques. The use of model-theoretic techniques in the study of spectra has been explicitly advocated in [5] and [7]. Unfortunately, the transformations used in those papers appear to be somewhat limited in their potential. Our intention in the present paper was to use elementary logical techniques to prove some spectra-theoretic results with a different flavour.

In the discussion following Theorem 2, we indicated that our methods would be suitable for proving non-closure results for the inverses of functions that have finite iterates greater than $n \mapsto 2^n$. However, it is not clear if our methods would be useful for proving non-closure results when the function in

question is sub-exponential, even if it is asymptotically greater than every polynomial over ω . It would be interesting to see what we can say about functions in this category.

There are two obvious questions that could be the subject of follow-up work. First of all, it would be interesting to know if we can establish any bounds on the closure properties of any natural, restricted classes of spectra. Second, we would like to see if the Gödel numbering introduced in this paper can be employed to study any other interesting problems in the theory of spectra. We are currently working on such applications.⁶

REFERENCES

- [1] G. Asser, “Das repräsentantenproblem in prädikatenkalkül der ersten stufe mit identität,” *Zeitschrift für math. Logik Grundlagen* 1 (1955) 252–263.
- [2] C. Chang, H. Keisler, *Model Theory*, Elsevier, Amsterdam: North-Holland, 1973.
- [3] R. Fagin, “Generalized first-order spectra and polynomial time recognizable sets,” in: R. Karp (Ed.), *Complexity of Computation*, American Mathematical Society, Providence, R.I., 1974, pp. 43–73.
- [4] R. Fagin, “A spectrum hierarchy,” *Zeitschrift für math. Logik Grundlagen* 21 (1975) 123–134.
- [5] A. Hunter, “Spectrum hierarchies and subdiagonal functions,” in: *Eighteenth International Symposium on Logic in Computer Science (LICS 03)*, 2003.
- [6] N. Jones, A. Selman, “Turing machines and the spectra of first-order formulas with equality,” *Journal of Symbolic Logic* 39 (1974) 139–150.
- [7] M. More, “Investigation of binary spectra by explicit polynomial transformations of graphs,” *Theoretical Computer Science* 124 (1994) 221–272.

⁶The author would like to thank Jim Delgrande for ongoing support. Special thanks are also due to Alistair Lachlan for many useful discussions and insightful suggestions about the results presented in this paper.

The *Australasian Journal of Logic* (ISSN 1448-5052) disseminates articles that significantly advance the study of logic, in its mathematical, philosophical or computational guises. The scope of the journal includes all areas of logic, both pure and applied to topics in philosophy, mathematics, computation, linguistics and the other sciences.

Articles appearing in the journal have been carefully and critically refereed under the responsibility of members of the Editorial Board. Only papers judged to be both significant and excellent are accepted for publication.

The journal is freely available at the journal website at

<http://www.philosophy.unimelb.edu.au/ajl/>.

All issues of the journal are archived electronically at the journal website.

SUBSCRIPTIONS Individuals may subscribe to the journal by sending an email, including a full name, an institutional affiliation and an email address to the managing editor at ajl-editors@unimelb.edu.au. Subscribers will receive email abstracts of accepted papers to an address of their choice. For institutional subscription, please email the managing editor at ajl-editors@unimelb.edu.au.

Complete published papers may be downloaded at the journal's website at <http://www.philosophy.unimelb.edu.au/ajl/>. The journal currently publishes in pdf format.

SUBMISSION The journal accepts submissions of papers electronically. To submit an article for publication, send the \LaTeX source of a submission to a member of the editorial board. For a current list of the editorial board, consult the website.

The copyright of each article remains with the author or authors of that article.